Eidgenössische Technische Hochschule Zürich
Swiss Federal Institute of Technology Zurich

Institut für
Technische Informatik und
Kommunikationsnetze

DaimlerChrysler

Master's Thesis

5.9GHz Dedicated Short Range Communication

Design of the Vehicular Safety Communication Architecture

Issued by Andreas Meier

Supervisor: Daniel Jiang
Professor: Dr. Lothar Thiele

Start Date: 2nd February 2005
Issue Date: 1st August 2005

Abstract

Wireless short to medium range communication in a vehicular environment has the potential to improve safety. In particular, a dedicated frequency band, referred to as 5.9GHz DSRC, has been allocated in the US for this exact purpose.

The exchange of safety messages among vehicles and with infrastructure devices poses major challenges. Specifically, safety messages have to be adaptively distributed within a certain range of a basically unbounded system. However, a traditional communication architecture is woefully inadequate for safety communications in such dynamic environments. For instance, the DSRC's non-deterministic channel characteristic requires functionality that increases the reception probability for safety messages. Furthermore, DSRC has prerequisites, such as channel switching, in order to offer non-safety communication on a separate channel.

This Master's Thesis presents a design of a vehicular safety communication architecture that meets demanding safety application requirements and also supports non-safety communication. In particular, an effective broadcast message distribution scheme is introduced, a channel-switch protocol is presented, and a communication stack is proposed.

Index Terms – Dedicated Short Range Communication, DSRC, 5.9GHz DSRC, IEEE 802.11p, Communication Architecture, Communication Stack, Vehicular Safety, Message Echoing, Independent Channel Switch Protocol, ICS

Acknowledgements

A number of people contributed to this thesis it their own unique ways and I would like to express my gratitude to them.

First and foremost I would like to thank Prof. Dr. Lothar Thiele for supporting my Master's Thesis.

I would like to express my gratitude to my supervisor Dan Jiang. His expertise in inter-vehicle communication and DSRC in particular provided me a most valuable source of information. Dan always directed me in the right way whenever I got lost in detail and was running risk to loose sight of the main objective.

I would like to thank my colleagues at the DaimlerChrysler Research and Technology North America. In particular to Svetoslav Yankov for the close collaboration, to Vikas Taliwal for the inspiring discussions, and to Chi Chen for all the DSRC related matters he helped me with.

My gratitude is given to Marc Torrent-Moreno for the interesting dialogue about vehicular communication architectures and his inspiring comments on the communication stack.

My gratitude is also given to the DaimlerChrysler Scholarship Program in Research and Technology, especially to Katja, Petra, Carola and Dr. Reinhold Eberhardt, for their support in general and for establishing the fruitful relationship with Dan.

I would also like to thank to Brian, Simon, Emily, Christoph and Rico for proofreading and all their helpful comments.

Finally, I want to express my special thanks to my parents for their long lasting support.

Contents

List of Figures

List of Tables

Abbreviations and Acronyms

Acronym	Description
AEVW	Approaching Emergency Vehicle Warning
AIFS	Arbitration Inter-Frame Space
BDTP	Bulky Data Transfer Protocol
BPSK	Binary Phase Shift Keying
CAN	Central Autonomic Network
CSMA/CA	Carrier Sense Multiple Access With Collision Avoidance
CTS	Clear To Send
dBm	Decibels Relative to One Milliwatt
DCF	Distributed Coordination Function
DC RTNA	DaimlerChrysler Research and Technology North America
DGPS	Differential Global Positioning System
DSRC	Dedicated Short Range Communications
EDCF	Enhanced Distributed Coordination Function
EEBL	Extended Electronics Brake Lights
EIRP	Effective Isotropic Radiation Power
FCC	Federal Communications Commission
FIPS	Federal Information Processing Standard
GMT	Greenwich Mean Time
GPRS	General Packet Radio System
GPS	Global Positioning System
HSM	Hardware Security Module
IP	Internet Protocol
IPv4	Internet Protocol version 4
IPv6	Internet Protocol version 6
LT	Luby Transform
LLC	Logic Link Control (Sublayer)
MAC	Medium Access Control (Sublayer)
MCC	Safety Message Creation and Control Layer
OFDM	Orthogonal Frequency-Division Multiplexing
PHY	Physical (Layer)
PQC	Packet Quality and Control (Layer)
QAM	Quadrature Amplitude Modulation
QoS	Quality of Service
QPSK	Quadrature Phase Shift Keying
RFC	Request For Comments
RSU	Road Side Unit
RTK	Real Time Kinematics
RTS	Request To Send

Acronym	Description
SAF	Safety (Layer)
SAP	Service Access Point
SHS	Single Hop Switch (Layer)
TCP	Transmission Control Protocol
UDP	User Datagram Protocol
UMTS	Universal Mobile Telecommunication System
V2I	Vehicle-to-Infrastructure
V2V	Vehicle-to-Vehicle
VSC	Vehicle Safety Communication Consortium
VANET	Vehicular Ad-Hoc Network
WAAS	Wide Area Augmentation System
WAVE	Wireless Access in Vehicular Environment

Chapter 1

Introduction

1.1 Motivation

In 1999 the Federal Communications Commission [46] authorised a 75MHz spectrum from 5.850 to 5.925GHz—referred to as 5.9GHZ DSRC—to be used for vehicle-to-vehicle and infrastructure-to-vehicle communication in the United States. Its main purpose is to implement applications that reduce accidents and improve traffic flow. In order to make the system more attractive for deployment, commercial applications can use the spectrum with certain restrictions as well. Not only in the United States, but also many other countries, particularly countries in Europe, are about to allocate a frequency spectrum for vehicular safety communication.

There are various projects addressing safety in combination with vehicular ad-hoc networks. This includes Network On Wheels [40] and its predecessor FleetNet [41] in Germany, IP PReVENT [42] and the Car-to-Car Communication Consortium [43] in Europe, and the Vehicular Safety Consortium and its successor VSC2 in the United States. Concluding, it can be said that there is a lot of effort all around the globe to promote this new technology.

Many of the current projects address specific safety applications such as intersection collision warning or extended electronic brake lights. Such applications do not require a sophisticated communication architecture to exchange a few messages. However, this cannot be assumed for the fully deployed system that provides all kinds of safety applications involving all vehicles on the street. It is therefore necessary to design a communication architecture that can meet all safety application demands. In addition, there are non-safety applications, such as toll collection or infotainment, that will also need transfer data.

Currently an IEEE working group defines a set of standards for a vehicular communication architecture [9, 10, 11, 12, 13]. Recently Füssler et al. proposed a new design approach for a vehicular ad-hoc networks [16] as well. However, all these activities neither address the specific demands for safety in detail, nor deal with the question of how non-safety communication can be achieved without jeopardising safety.

It is the objective of this Master's Thesis to design a vehicular safety communication architecture that meets all the demands of the safety applications and offers the possibility of non-safety communication.

1.2 Structure

This thesis is organised as follows:

Chapter 2 provides information about DSRC. Includes most important technical details about 5.9GHZ DSRC and presents the channel allocation scheme in the United States. Highlights that a non-deterministic channel scheme must be assumed for the communication and that the channel load must be limited due to the unavoidable hidden terminal effect. Furthermore, discusses the boundaries of a vehicular ad-hoc network.

Chapter 3 discusses related work. Presents ongoing projects focussing on vehicular ad-hoc networks to improve safety. Furthermore, analyses communication architectures and channel-switch schemes suggested to be used with DSRC.

Chapter 4 provides an overview of the potential safety and non-safety applications.

Chapter 5 analyses the demands of the safety applications and their impact upon the messages. Emphasis that a combination of routine and event-based messages is necessary to provide safety. Highlights that safety applications do not transfer messages application based and discusses security-related issues.

Chapter 6 deals with the implications on the communication architecture the requirements discussed in Chapter 5 have. Analyses how the communication channel can be prevented of breaking down, despite the need to send messages on a regular basis. Introduces an acknowledgement scheme that is not required to generate additional packets. Further discusses the event message distribution, presents the echo mechanism and introduces a publish/subscribe mechanism for safety messages.

Chapter 7 discusses channel switching. The proposed channel-switch scheme ensures safety based on a timely delivery of events and continuously updates the context. With this ensured, the non-safety throughput of a single station and the overall system is maximised.

Chapter 8 proposes a communication stack to provide all the necessary services in accordance with the requirements presented in the preceding chapters.

Chapter 9 discusses the communication stack. Emphasis that the proposed architecture provides the necessary functionality and discusses its advantages compared with other approaches.

Chapter 10 summarises the thesis. Furthermore, lists the main contributions and discusses further work.

Appendix A presents security related issues of vehicular safety communication. **Appendix B** presents the PHY and MAC frame and proposes a shortened alternative for the latter. **Appendix C** proposes a safety message structure and discusses the most important V2V safety information units.

Chapter 2

Dedicated Short Range Communication (DSRC)

Dedicated Short Range Communication commonly refers to short to medium range wireless communications that offers data transfer in a vehicular ad-hoc network.

A vehicular ad-hoc network with the purpose to do safety communication has two major differences compared to a traditional ad-hoc network: first, the stations are positioned much wider apart resulting in a non-deterministic channel characteristic, and second, the hidden terminal effect cannot be avoided due to the unbounded system and the broadcast nature of safety communication. Consequently a new standard had to be defined to meet the vehicular ad-hoc network's specific requirements. This new standard, referred to as *IEEE 802.11p* [9], is based on 'IEEE 802.11a' [7] and defines the functionality of the physical and the medium access control layer. It should be noted that DSRC devices are assumed to support both, the 'IEEE 802.11p' and the 'IEEE 802.11a' standard.[1]

The remainder of this chapter highlights the characteristic of DSRC and is structured as follows. A brief technical overview is given and the channel concept for safety and non-safety communication is presented. It is pointed out that DSRC show everything but a deterministic channel characteristic due to fading and path loss. It is shown that the hidden terminal collision cannot be avoided. This results in an even worse reception characteristic if the channel load is not limited in order to restrain the number of collisions. It is further pointed out that the channel can handle the load caused by the protocols proposed in this work. This chapter closes discussing the boundaries of a vehicular ad-hoc networks.

It should be noted that the spectrum and channel specific information presented in this chapter are in accordance with the regulations of the Federal Communication Commission in the United States.

[1] The frequency ranges of 'IEEE 802.11a' and 'IEEE 802.11p' are just next to each other. So the implementation of a radio supporting both frequency ranges is not assumed to cause difficulties. If other wireless standards are supported, in particular 'IEEE 802.11g' and 'IEEE 802.11n', is up to the manufacturer.

2.1 DSRC – Technical Overview

- **Bandwidth**

  75MHz (5.850 – 5.925GHz)

- **Channels**

  There are seven non-overlapping 10MHz channels. There is the option to combine two of them into one 20MHz channel.

- **Guard Channel**

  5MHz are reserved at the lower end as a so-called "guard channel".

- **Modulation**

  BPSK OFDM, QPSK OFDM, 16-QAM OFDM, 64-QAM OFDM. The first 128 bits are always BPSK coded.

- **OFDM symbol duration**

  8.0 microseconds

- **Data Rate**

  6, 9, 12, 18, 24, and 27Mbps with 10MHz Channels (3Mbps preamble) (or 6, 9, 12, 18, 24, 36, 48, and 54 Mbps with 20MHz Channel option) (6Mbps preamble)

- **Power**

  Usually less than 33dBm (2W) but up to 44.8dBm (30W) for qualified public safety applications on the Control Channel.

2.2 Channels

The channel scheme for DSRC in the United States is illustrated in Figure 2.1. DSRC supports seven non-overlapping 10MHz channels with the option to join two channels to double the bandwidth. This is in contrast with the 'IEEE 802.11b/g' standards, providing several more but overlapping channels. See Table 2.1 for details.

It is necessary that all safety messages are transmitted on one designated channel only in order to ensure that all vehicles listen to the proper one for such messages. This channel is referred to as *Control Channel* and corresponds to the channel number 178 in the United States. The communication on the Control Channel is principally used for safety related communication only and non-safety data exchange is strictly limited in terms of transmission time and interval as listed in Table 2.2.

The other six channels—eight, if you take the two 20MHz channels into account—are referred to as *Service Channels*. Two of them, namely 172 and 184, are reserved for safety-related applications. However, these two channels are not meant to be an option for the regular safety communication on the Control Channel. The remaining six channels, namely 174, 175, 176, 180, 181 and 182 can be used for non-safety communication.

It cannot be expected that manufacturers equip their DSRC devices with two radios. This implies that channel switching is necessary to do both safety and non-safety communication. Such a channel switch can be performed in less than one millisecond.[2]

[2]This period is in accordance with Atheros [47], the manufacturer of DSRC radio prototypes.

The maximum allowed effective isotropic radiated power (EIRP) is regulated as shown in Table 2.3. The EIRP of 44.8dBm (30W) on the Control Channel would theoretically allow to broadcast safety messages up to 1000 metres. In reality this range can hardly ever be achieved due to fading and interference. It should be noted that, unlike most mobile communication systems, energy consumption is not a major issue in a vehicle.

2.3 Channel Characteristic

Wireless communication is in general much less reliable than a wired one. In a wired communication system packets are usually lost due to congestion and are not likely to be corrupted because of bit errors.

This is different in a wireless communication system. There are two natural effects, namely path loss and fading, resulting in a reception degradation of the signal. These effects are discussed in the remainder of this section to provide a general idea of the reception characteristic of a single DSRC broadcast.

Path loss quantifies the decrease of the signal strength in accordance with the distance to the sender. The signal strength can attenuate faster or slower than it does in free space. So the reception range for a signal can vary greatly in a fast changing environment. In addition to the path loss, the fading effect adds a lot of variation to the signal strength due to minor changes in the environment. This is because of the multi-path propagation of the signal, resulting in constructive or destructive interference.

The combination of the path loss and the fading is depicted in Figure 2.2(a), showing an exemplary, rather non-deterministic, channel characteristic of a DSRC test run.[3] It should be noted that this is everything but the deterministic channel characteristic, as illustrated in Figure 2.2(b) with a randomly chosen reception range.

2.4 Message Collisions

In an ad-hoc network, different stations may want to send data at the same time. In order to prevent a collision, wireless networks can use the DCF scheme to solve the "multi station access contention problem". This scheme minimises the probability of a collision as long as the involved stations are in sensing range. However, this cannot be assumed for an ad-hoc network in which the hidden terminal problem can arise.

The hidden terminal problem is illustrated in Figure 2.3. In this situation vehicle A is broadcasting data. Vehicle B is in reception range but not vehicle C. Since the latter is not aware of the ongoing communication, it may start transmitting data as well. This could result in a message collision between the two sending stations, namely at vehicle B.

In order to tackle the hidden terminal problem, the DCF scheme can be combined with the RTS/CTS handshake mechanism. However, the broadcast nature of safety messages—there is no distinct station to address in the RTS frame—does not allow such a handshake. And so far no other scheme has been proposed to solve the hidden terminal problem efficiently in an unbounded system doing broadcast communication. Therefore, it has to be assumed that collisions occur frequently if a lot of messages are sent on the channel.

[3] This characteristic may vary a lot depending on the current environment of the broadcasting station.

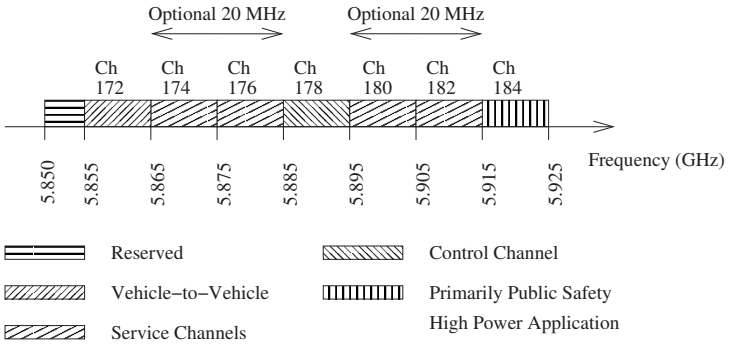

Figure 2.1: The channel scheme assigned in the United States.

Standard	IEEE 802.11a	IEEE 802.11b	IEEE 802.11g	IEEE 802.11p (DSRC)
Modulation	OFDM	DSSS	OFDM	OFDM
Frequency [GHz]	5.725−5.850	2.400−2.485	2.400−2.483	5.850−5.925
Channel Bandwidth [MHz]	20	22	22	10/(20)
Nr of Channels/ non-overlapping	12/8	14/3	14/3	7/7
Max Rate [MBit/s]	54	11	54	27/(54)

Table 2.1: Comparison of the most common wireless standards with IEEE 802.11p.

	RSU	Vehicle
Maximum Data Transmission Duration	750μs	580μs
Minimum Interval between Data Transmissions	100ms	750ms

Table 2.2: Control Channel usage limits for non-safety transmission in the United States.

Channel Number	Frequency [GHz]	Max EIRP [dBm]	
		Public Safety	Private Usage
172	5.855 – 5.865	33.0	33.0
174	5.865 – 5.875	33.0	33.0
175	5.865 – 5.885	23.0	23.0
176	5.875 – 5.885	33.0	33.0
178	5.885 – 5.895	44.8	33.0
180	5.895 – 5.905	23.0	23.0
181	5.895 – 5.915	23.0	23.0
182	5.905 – 5.915	23.0	23.0
184	5.915 – 5.925	40.0	33.0

Table 2.3: Transmitter power limitations in the United States.

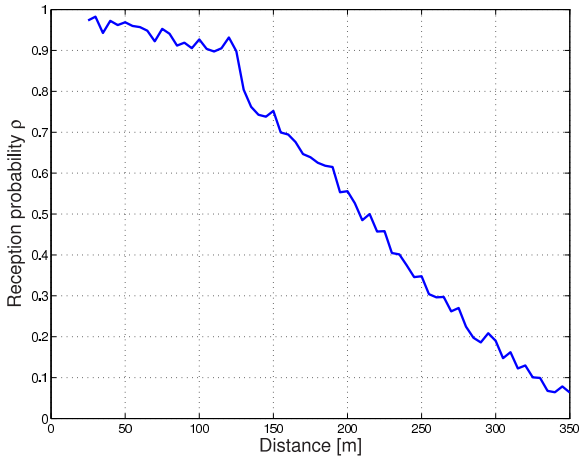

(a) Real, non-deterministic channel characteristic for DSRC.

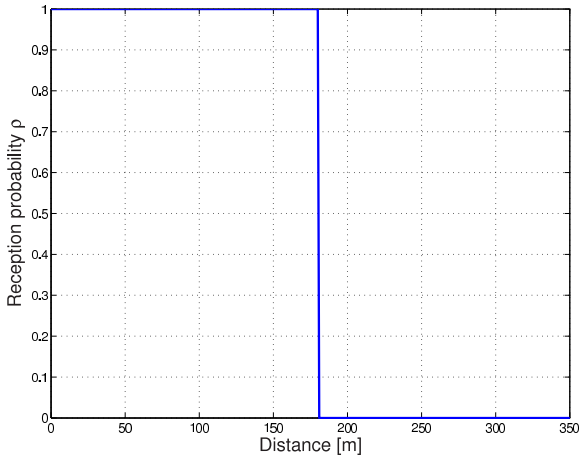

(b) Deterministic channel characteristic with a randomly chosen reception range.

Figure 2.2: The reception probability against the distance to the sending station.

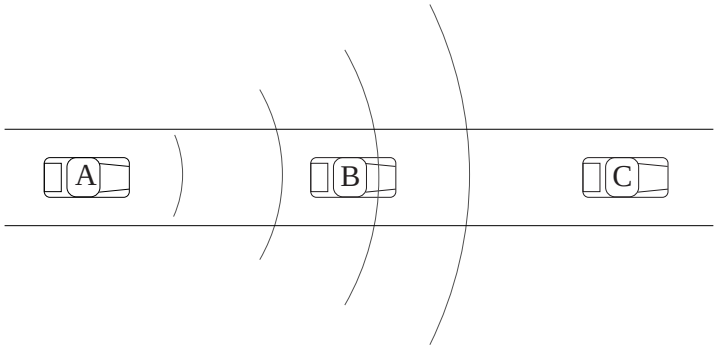

Figure 2.3: Hidden terminal problem: Vehicle C is unaware of the fact that vehicle A is broadcasting data and might start a broadcast as well. This may result in a message collision between the two vehicles.

2.5 Channel Capacity

The channel load has to be limited to restrain the number of packet collisions. However, current research cannot provide a number, such as bandwidth or channel occupation time, to specify the channel's capacity.

Some of the protocols proposed in this work require a certain message density in order to show their full potential. In particular the channel-switch scheme and the event distribution are based on the assumption that the channel can handle a message every $2 - 4$ milliseconds.

An average safety message is assumed to be about 300 bytes long and is likely to be QPSK modulated since the signal-to-noise ratio of the higher order modulations is pretty low. These, rather conservative, assumptions result in an average transmission time of a safety message of about 400 microseconds. It is therefore reasonable to assume that the channel does not break down if a safety message is broadcast every few milliseconds.

2.6 Boundaries

Most wireless systems are based on a cellular structure. Such a cell has its natural boundaries in terms of the distance it covers. This does not imply that the system is limited to a single cell, as different cells can be interconnected. And it might be possible that a station changes the cell during the communication using some kind of hand over mechanism.

However, this is different for an ad-hoc network that has in general no fixed boundaries. In particular, the boundary of a vehicular ad-hoc network is highly dynamic due to the rapid movement of the stations. This makes the estimation of the current cell size rather difficult, especially if the non-deterministic broadcast characteristic is taken into account. However, a vehicle will never switch from one cell to another, but two cells combine or a single one splits up.

Chapter 3

Related Work

3.1 DSRC Related Projects

There are various projects addressing the different issues of inter-vehicle communication. Many of them are government funded and have the main purpose to improve safety and traffic flow. The most important of these projects are presented below.

FleetNet

'FleetNet – Internet on the Road' [41] was set up by a consortium of six companies and three universities in Europe to promote the development of vehicular ad-hoc networks. The project started in September 2000 and ended three years later.

The inter-vehicle communication was based on 'IEEE 802.11 a/b' using IPv6 to exchange data. The project focussed on geographic addressing and multi-hop routing along the street. In particular, ten cars were equipped with communication devices to make test runs in different routing scenarios.

Network on Wheels

Networks on Wheels (NOW) [40] is an ongoing project that will end in 2008, superseding its predecessor FleetNet and is supported by the German Federal Ministry of Education and Research.

The main goal of the project is to specify a communication system to transmit sensor data and general information about the vehicle using inter-vehicle communication. A communication system has not been specified yet, but recently a first approach for a protocol architecture for a vehicular ad-hoc network was proposed. This approach is discussed in detail in Section 3.2.3.

IP PReVENT

The Integrated Project PReVENT [42] is a consortium of European automotive manufacturers, co-funded by the European Commission, promoting road safety. This project is not only focussed on inter-vehicle communication. It also considers sensing and positioning technologies in combination with digital-maps.

In regards to inter-vehicle communication, the IP PReVENT's main field of research is to design an intersection safety application. This application detects crossing traffic based on a sophisticated sensor network and uses vehicle-to-infrastructure communication to exchange additional information. However, there is no vehicle-to-vehicle communication involved.

Vehicle Safety Communication Consortium

The Vehicle Safety Communication (VSC) Consortium is a research program initiated by the United States Department of Transportation. It consists of seven automotive manufacturers and has the objective of faciliting the advancement of vehicular safety based on DSRC. In particular, the VSC published a study [1] about possible safety applications to be used with DSRC. Currently the VSC is investigating specific technical issues related to these applications.

3.2 Communication Architectures

Research about inter-vehicle communication is getting more and more popular. However, research about designing a vehicular safety communication architecture is not an advanced research topic at all. Currently there are only two communication architectures proposed that claim to provide safety in combination with 5.9Ghz DSRC. These two approaches and the traditional layered design are presented subsequently.

3.2.1 Layered Approach

The 'traditional' layered approach for a communication architecture is illustrated in Figure 3.1(a). Such a layered architecture has been proven to be very successful since its first implementation about thirty years ago. The basic idea is that each layer provides services for the adjacent upper layer, accessible by a well-defined interface—referred to as "service access point" (SAP). The service itself uses a protocol that is implemented in the specific layer and can be freely designed without considering the overall architecture. For instance, an upper layer does not care what physical medium is used to transport the bit stream as long as the service 'transport raw bits' is provided.

It should be pointed out that the different layers are supposed to be independent modules. Hence a protocol must not interact with another one in a different layer. In particular, a protocol shall not access meta-data delivered in the header of another protocol, or exchange protocol-state information.

3.2.2 IEEE P1609.3

Currently an IEEE working group is working on a set of standards—referred to as 'IEEE 802.11p' [9], 'IEEE P1609.1' [10], 'IEEE P1609.3' [11], 'IEEE P1609.4' [12] and 'IEEE 1556' [13]—specifying a WAVE (Wireless Access in Vehicular Environment) communication system. One of the standards, 'IEEE P1609.3', specifies the overall communication architecture; the other ones focus on the architecture's details. These standards exist in early draft versions and are therefore due to changes.

The communication architecture proposed in 'IEEE P1609.3' is illustrated in Figure 3.1(b). The most evident part is its dual stack: on the one hand there is the well-known TCP/IP stack and on the other there is the WAVE Short Message stack. The function of the latter one is to provide a connectionless transport protocol—similar to UDP, but on a single-hop basis. The safety applications are supposed to use this stack only, while non-safety applications can use both.

It should be noted that the design of this approach is focussed on non-safety applications and considers safety as a black box that should fit in the current design.

3.2.3 Staircase Approach

Füssler et al. have proposed an approach for a communication architecture designed to meet the demands of a vehicular ad-hoc network. Their so-called *staircase approach* [16], is depicted in Figure 3.1(c) and provides the following four key features:

- **Layered Design**

  There are still layers existent in the design. The layers have the same core functions as the 'traditional' TCP/IP or ISO/OSI design, but different names have been assigned to provide a better understanding of their functions:

  - The single-hop layer corresponds to the data link layer and incorporates all functionality dealing with communication with direct radio neighbours. This layer should include a stability element that ensures that the single-hop communication adapts to channel and load conditions.
  - The multi-hop layer corresponds to the network layer and contains the necessary functionality to forward packets to non-neighboured nodes. Geocast is the main protocol in this layer but traditional multi-hop protocols should be supported as well.
  - The data-link corresponds to the transport layer and provides a reliable byte stream on top of the multi-hop layer.

- **Staircase Approach**

  The applications can directly access the different layers. For instance, if only a single-hop broadcast should be performed, the application can bypass the datagram and the multi-hop layer. In order to directly access the different layers, a packet header element such as a port/protocol number is required to allow multiple applications.

- **Information Connector**

  The Information Connector provides a common interface that should allow an efficient exchange of data between the different layers. Such data could be sensor update information, data extracted from packets, or state information of protocol layers and devices.

  The Information Connector is accessible using a publisher/subscriber scheme. In particular, the Information Connector shall not change the state of a protocol, but the protocol might react to the Information Connector's notifications.

- **Management Plane**

  The management plane controls long-term system settings. In particular, the plane is not involved in the dynamic self-organisation motivated by the different network conditions.

The staircase approach proposes two uncommon ideas. First, there is the Information Connector. According to Füssler et al. the necessity for an information exchange between layers is due to the nature of vehicular ad-hoc networks. Unfortunately there is a lack of understanding of exactly why this is the case. This does not imply that the Information Connector is not a good idea. Second, there is a staircase access of the applications to the different layers. The various access points attempt to deal with the different requirements that the different safety and non-safety applications have on the message distribution.

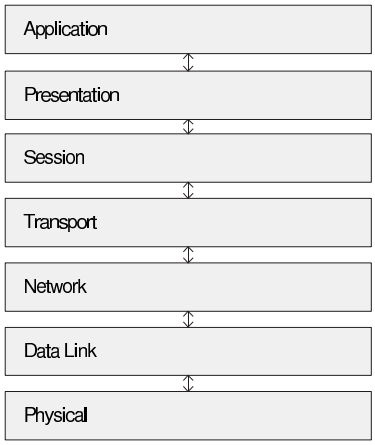

(a) Layered design according to the ISO/OSI reference model.

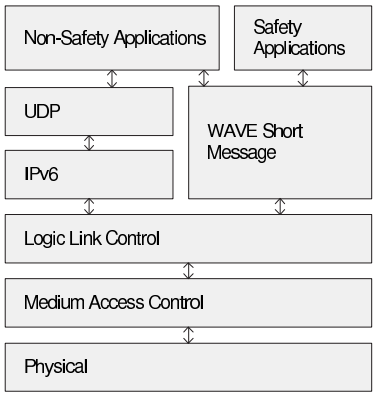

(b) Architecture proposed in 'IEEE P1609.3'.

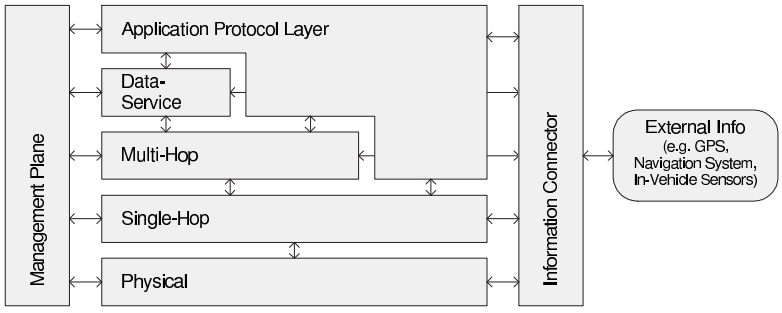

(c) Staircase approach

Figure 3.1: Possible communication stacks for DSRC.

The staircase approach is only meant to provide thoughts of a vehicular ad-hoc network and therefore does not provide a lot of details. Important issues such as congestion control and the distribution of safety messages are not addressed thoroughly. Other, DSRC specific features—such as channel switching—are not addressed at all.

3.3 Channel Switching

Safety messages are transferred on the Control Channel. Non-safety communication on the other hand requires switching to a Service Channel. Hence a vehicle cannot do safety and non-safety communication at the same time. The following two schemes propose tackling the problem in a synchronised manner.

3.3.1 Global Synchronisation

The global synchronised channel switching has been proposed by Scott Andrew, an independent contractor in the automotive industry, and Steve Tengler from Nissan. Channel synchronisation means that the vehicles split up the time into safety and non-safety time slots in a synchronised manner. The safety time slot is used to do safety related communication on the Control Channel exclusively, while the non-safety time slots can be used to exchange non-safety data on a Service Channel.

The idea of the global synchronisation scheme was proposed recently and no details have been revealed yet. In order to get an idea of the potential of this approach, i.e. to estimate the feasibility of this scheme and to benchmark it, the general idea of a global synchronisation is explored.

Technical Details

The global synchronisation is achieved based on the GPS signal and is illustrated in Figure 3.2. The time is split up into time slots of fixed length, $T_{\text{GlobalSync}}$, separated by so-called *GPS synchronisation points*. These synchronisation points are not meant to be a signal sent by the satellite, but are calculated using the global clock provided with the GPS signal.

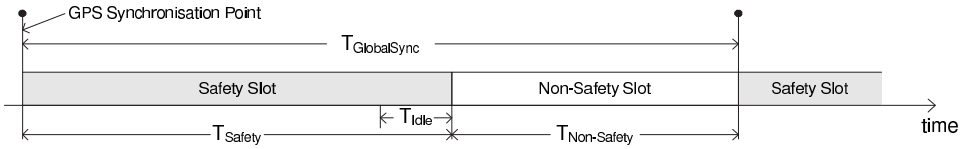

Figure 3.2: Global synchronisation: The time is split up into intervals of fixed length $T_{\text{GlobalSync}}$. The beginning of the interval is always used for safety communication. If the Control Channel is idle for T_{Idle}, the remaining time of the interval can be used for non-safety related communication on other channels.

The GPS synchronisation point indicates the beginning of the safety time slot in which safety related messages are sent. As soon as the channel is idle for T_{Idle} ($\sim$ 5ms – 10ms)—i.e. no more safety messages are sent—the non-safety time slot starts and ends with the next synchronisation point.

The time between two GPS synchronisation points $T_{\text{GlobalSync}}$ is fixed while the safety time slot T_{Safety} is not restricted in terms of its length. Therefore, the time

slot for non-safety communication is determined by the length of the safety slot and can be calculated as

$$T_{\text{Non-Safety}} = T_{\text{GlobalSync}} - T_{\text{Safety}}, \qquad \text{where } T_{\text{Safety}} \leq T_{\text{GlobalSync.}}$$

This implies that the time to do non-safety communication might be zero if there are a lot of safety related messages to be sent.

It cannot be assumed that all stations are listening to safety messages during the non-safety time slot. Therefore, if an event is triggered during the non-safety time slot, the station has to wait to the next GPS synchronisation point before the message can be sent. This fact limits the maximal time that can be spent for non-safety communication and determines the upper bound between two GPS synchronisation points:

$$T_{\text{GlobalSync}} \leq T_{\text{MaxLatency}} + T_{\text{Idle}} \approx 100\text{ms}$$

The maximum tolerable latency $T_{\text{MaxLatency}}$ must be appropriate for all traffic conditions, in particular a high speed and high traffic density situation, and is on the order of 100 milliseconds [3].

Safety Communication

Safety must be ensured and is therefore the most important issue regarding the feasibility of this channel-switch scheme. The following concerns about safety should be considered:

- **Maximum Latency**

  The maximum latency is determined by the time $T_{\text{GlobalSync}}$ between two GPS synchronisation points. This time is chosen based on the maximum allowable latency and therefore fulfils the requirement if a reliable safety communication can be ensured within one safety time slot.

- **Early Switch**

  Whenever the Control Channel is idle for T_{Idle}, the stations are free to switch to another channel. A problem occurs if the switch is done too early—i.e. some stations still send safety messages. Such a situation can arise due to the unbounded nature of the system, i.e. stations out of reception range of a vehicle are sending messages while stations in between are waiting for their channel access.

- **GPS Leakage**

  The synchronisation is done based on the global time provided by the GPS signal. GPS leakages occur once in a while and could jeopardise the synchronisation. However, the GPS synchronisation points can be extrapolated using an internal clock if the GPS signal is not available. It should be noted that this extrapolation of the time can be assumed to be more accurate then the one of the position that is needed for safety communication.

- **Channel Access Time Distribution**

  The channel access strategy used with DSRC is referred to as enhanced DCF. The principal assumption behind the channel-access scheme, is that the message generation is uniformly distributed over the time. However, this assumption does not hold for a synchronised channel access scheme, in which safety messages are aggregated during the non-safety time slot. This aggregation is likely to result in a concentration of message collisions in the beginning of the safety time slot.

It should be noted that the safety aspect is further discussed in Section 7.8.

Non-Safety Communication

The available bandwidth for non-safety applications depends on the traffic generated by the safety applications. An average channel usage scenario is depicted in Figure 3.3(a): a part of the available time is spent to do safety communications, but a distinctive part can be spent for non-safety communication as well. The proportion of safety and non-safety time slots depends mainly on the traffic scenario.

Whenever a lot of vehicles are in broadcast range, the number of safety messages increases and therefore the length of the non-safety time slots decreases. Such a situation is shown in Figure 3.3(b). There is still time available to do non-safety communication, but the available time is fragmented resulting in very short communication slots, that might be too short to be of much use.

It might even happen that time for non-safety communication is not available at all. This is shown in Figure 3.3(c). One might argue that safety is the main purpose and non-safety communication should only be allowed if the situation allows it. However, there are important non-safety applications as well—such as electronic toll collection—that need to exchange little data in order not to jeopardise the traffic flow.

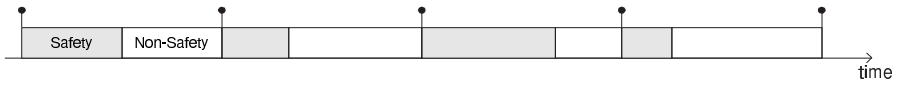

(a) Average throughput.

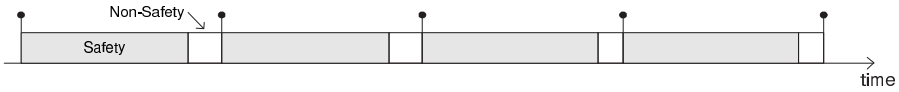

(b) Fragmentation.

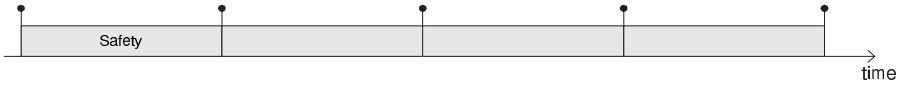

(c) No non-safety throughput at all.

Figure 3.3: Global Synchronisation: Different Throughput Scenarios.

It should be emphasised that the global synchronisation mechanism synchronises the non-safety time slots as well. This provides the possibility to do non-safety communication that implies interaction between most of the stations, such as routing.

3.3.2 *i*-Channel

The *i*-Channel scheme has been introduced by B. Wells and J. Hunzinger from DENSO International America Incorporated [48]. This mechanism is best described as a distributed synchronisation scheme. Instead of having a global synchronisation, based on a superior point, the stations synchronise themselves in a distributed manner.

Technical Details

The *i*-Channel claims to synchronise all vehicles in a so-called *cluster*. Such a cluster is defined as a set of vehicles that are connected, i.e. in broadcast range, with each

other. This connection can be indirect, i.e. via multiple hops, as well.[1] For instance, a 10 mile long highway strip, with a vehicle every 50 metres, would make one sparsely connected cluster, but the cluster falls apart if there is a longer highway stretch without any vehicles.

The synchronisation of a cluster is done as illustrated in Figure 3.4. As a starting point, lets assume that all vehicles in the cluster are synchronised and do therefore switch to the Control Channel at the same time. After a short period—the receive-only time T_{RX} ($\sim 1\mathrm{ms} - 3\mathrm{ms}$)—the stations send their safety related messages. As soon as the Control Channel is idle for T_{idle} ($\sim 5\mathrm{ms} - 10\mathrm{ms}$) or the channel was busy for $T_{\mathrm{SafetyMax}}$ ($\sim 300\mathrm{ms}$) the non-safety time slot starts. This non-safety time slot has a fixed length, which is chosen according to the maximal tolerable latency for whatever traffic condition:

$$T_{\text{Non-Safety}} \lesssim T_{\mathrm{MaxLatency}} \approx 100\mathrm{ms}$$

As soon as this fixed period is over, the vehicles shall switch back to the Control Channel. In order to deal with minor synchronisation shifts, the receive-only time T_{RX} should make it most likely that no safety messages are sent before all vehicles are listening to the Control Channel.

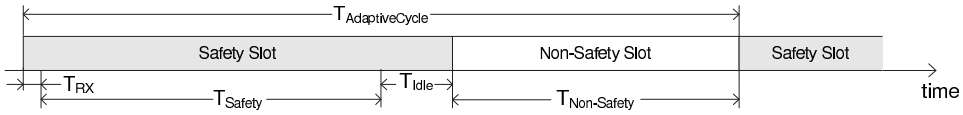

Figure 3.4: i-Channel: The time is split up into a safety time slot of variable length and a non-safety time slot of fixed length.

Two clusters combining to a single one are most likely out of sync. In order to tackle that problem, the stations are expected to swap to a so-called "high awareness" mode regularly (every $500\mathrm{ms} - 1000\mathrm{ms}$). In this mode, the station keeps listening to the Control Channel during the non-safety time slot. If safety communication of another cluster is detected, the station broadcasts a so-called "follow me flag" indicating the other cluster should synchronise.

Safety

There are mostly the same concerns about safety as with the global synchronisation scheme. However, there are a few differences:

- **Maximum Latency**

  As long as the vehicles in the cluster are synchronised the argument of the maximum latency is the same as with the global synchronisation scheme.

  However, if a station is out of sync, the maximum latency could increase significantly. Assuming the high awareness mode is entered once in a second, the maximum latency is just about this time. It should be noted that this worst-case latency usually does not occur between vehicles very close by as they are not likely to be out of sync.

- **Fragmentation**

[1] In mathematical terms a cluster is defined as a connected graph where the vehicles are the nodes and the connections the edges.

The cluster can be widely stretched. In that case, it is likely that the beginning of the idle time is not the same for all vehicles within one cluster. Hence a fragmentation of the cluster is likely.

- **Channel Access Time Distribution**

  The i-Channel suffers the problem of accumulating safety messages during the non-safety time slot as well.

- **Connection**

  The i-Channel claims to synchronise all vehicles in a cluster—i.e. all vehicles directly or indirectly connected. However, such a connection is based on a deterministic broadcast range and cannot be assumed with DSRC.

  First simulations of this channel-switch scheme, using a rather deterministic broadcast characteristic, have been shown to be very promising. However, further simulations are required to see the effect on the synchronisation a more realistic channel characteristic has.

Non-Safety Communication

The average non-safety throughput is comparable with the one of the global synchronisation scheme as the same length of time is spent to do safety communication. However, the i-Channel does not suffer the effect of the fragmentation due to the fixed length of the non-safety time slot. In addition, the safety time-slot's limited period ($T_{\text{SafetyMax}}$) guarantees a minimal bandwidth for non-safety communication and can be estimated as:

$$B_{\text{Min}} = \frac{T_{\text{Non-Safety}}}{T_{\text{Non-Safety}} + T_{\text{SafetyMax}}} \cdot B_{\text{Max}} \approx 0.25 \cdot B_{\text{Max}},$$

where B_{Max} is the bandwidth of the Service Channel.

Chapter 4

Applications

This chapter presents applications that can be used in combination with DSRC. These applications can be divided into three categories: There are two groups of safety related applications—namely the vehicle-to-vehicle safety applications and the infrastructure-to-vehicle safety applications—and there is a group of non-safety applications. The following lists should give a rough idea about the applications belonging to each:

Vehicle-to-Vehicle Safety Applications

- Extended Electronic Brake Lights
- V2V Hazard Warning
- Approaching Emergency Vehicle Warning
- Highway Merge Assistant
- Lane Change Warning

Infrastructure-to-Vehicle Safety Applications

- Traffic Signal Violation Warning
- Left Turn Assistant
- Stop Sign Movement Assistance / Stop Sign Violation Warning
- Hazard Warning (Construction Zone, Curve, Bridge, Tunnel...)

Non-Safety Applications

- Electronic Toll Collection
- Real Time Traffic Information
- Map Update
- Drive Through Payment
- Short Time Internet Access
- Wireless Transfer of Digital Entertainment

4.1 Vehicle-to-Vehicle Safety Applications

The different vehicle-to-vehicle (V2V) safety applications exist only as a first draft [1]. Two of them—namely the 'Extended Electronic Brake Lights' and the more general 'V2V Hazard Warning'—are currently in first stage of development in the Daimler-Chrysler Research and Technology North America (DC RTNA). These two similar applications are considered to have the biggest impact on safety in the near future and are analysed in detail in this section. Additionally, the 'Approaching Emergency Vehicle Warning' is discussed to present another type of V2V safety application.

4.1.1 Extended Electronic Brake Lights (EEBL)

Whenever the brake pedal is tapped in a vehicle, the brake lights light on to warn the vehicles driving behind of this potential source of danger. However, if the driver's sight is limited—e.g. sharp turn, other vehicles, or bad weather conditions—the warning cannot be seen. Such a scenario is illustrated in Figure 4.1. In this example, two cars (A and B) and a truck in between are travelling in the same lane. Suddenly the leading car A breaks hard, but driver B is unlikely to see the lit up brake lights and is unaware of the potential danger.

The 'Extended Electronic Brake Light' (EEBL) is the DSRC counterpart of the traditional brake lights. It intends to improve safety by broadcasting a warning message whenever a hard brake is detected.[1]

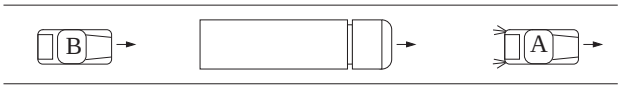

Figure 4.1: Extended Electronic Brake Lights: Vehicle A is braking, but vehicle B is not aware of the hazardous situation ahead as its line of sight is blocked by the truck. Safety can be improved if vehicle A broadcasts a warning message.

A vehicle receiving an EEBL message can integrate this information into its adaptive cruise control system in order to notify the driver about the potential danger. However, not every single received EEBL message should be announced to the driver as the following example illustrates.

In this more complex scenario, illustrated in Figure 4.2, vehicle A is braking and broadcasts a warning message. Vehicle B is driving behind the braking one and should notify its driver about the possible danger. All other vehicles can neglect the warning message for different reasons. Vehicles C has passed the breaking car, vehicle D is driving on another lane and vehicle E and F are heading in opposite direction.

The braking vehicle has to provide all the information that is required by vehicles receiving the announcement to decide about the message's relevancy. If all vehicles have a very accurate positioning system and additional map data, the information of the position would be sufficient to decide about the messages relevancy.[2] However, it cannot be assumed that all vehicles have accurate maps and positioning systems. Therefore, the information of the vehicle's heading should be added to the EEBL message to clarify the situation—e.g. vehicles driving in opposite direction are unlikely to drive in the same lane.

[1] This detection is based on the vehicle's sensor data.
[2] A thorough discussion about position accuracy is held in Appendix C.2.2.

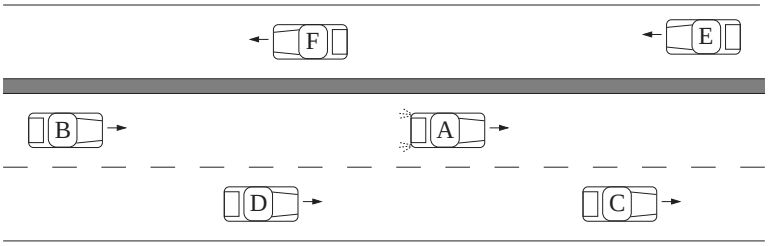

Figure 4.2: Extended Electronic Brake Lights: Vehicle A is braking and broadcasts a warning message. Vehicle B is the only vehicle that should inform its driver about the hazardous situation ahead.

4.1.2 Vehicle-to-Vehicle Hazard Warning (V2VHW)

The V2V Hazard Warning (V2VHW) safety application is the sophisticated DSRC counterpart of the traditional warning light in a vehicle.

A vehicle broadcasts a warning message if it is a potential danger for the other traffic participants—e.g. a vehicle stopped in the middle of the street or the driver lost control over the vehicle. Such a scenario is depicted in Figure 4.3 showing vehicle A being out of control and heading towards the oncoming traffic. In this situation not only the vehicles driving behind car A are in danger, namely C and D, but also vehicle E heading in the opposite direction. The other two cars (B and F) are not at risk since they have already passed vehicle A.

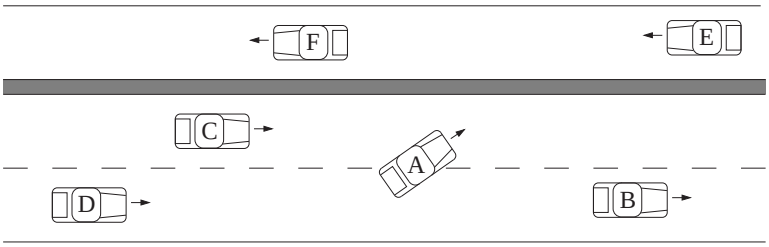

Figure 4.3: Vehicle-to-Vehicle Hazard Warning: Vehicle A is out of control and broadcasts a warning message.

Another example of the V2VHW application is the announcement of a hazardous street condition the vehicle's sensors have detected—e.g. aquaplaning or gliding on ice.

4.1.3 Approaching Emergency Vehicle Warning (AEVW)

Approaching Emergency Vehicle Warning (AEVW) messages are broadcasted by public safety vehicles such as ambulances or police vehicles. This application is the DSRC counterpart to the traditional siren wailing and blue light flashing, providing vehicles driving ahead with the information to yield the right of way.

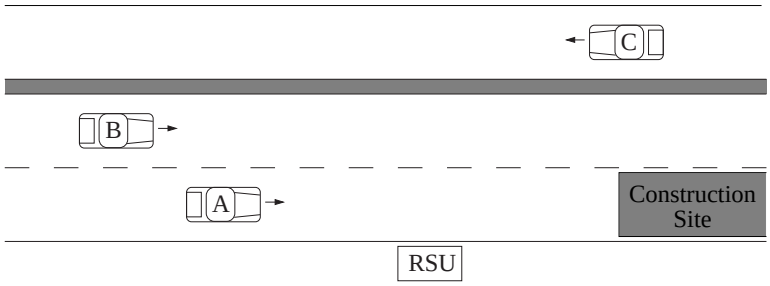

Figure 4.4: V2I Hazard Warning: A construction site is blocking one lane. The RSU is broadcasting warning messages to inform the approaching vehicles about the blocked lane.

4.2 Vehicle-to-Infrastructure Safety Applications

There are various possible V2I safety applications as indicated in the introduction of this chapter. However, the exact functionality of these applications is not known. Therefore, the general 'hazardous situation ahead' application is presented here.

The V2I Hazard Warning message is broadcasted by a so-called *road side unit* (RSU) to inform approaching vehicles about a hazardous situation ahead—such as a construction zone, a traffic light, or a tight curve. Such a situation is illustrated in Figure 4.4 showing a lane that is blocked by a construction site. A RSU, preferably positioned far ahead of the beginning of the construction site, is broadcasting a message containing all necessary information about it—such as details about the construction site, or a new speed limit. This information should be provided to the drivers in vehicle A and B.

It should be noted that it is not meant that RSUs are placed in front of each hazardous situation and road sign. There is rather one dedicated RSU that provides information about road signs and possible hazards close by. This is referred to as "In-Vehicle Signage". This does not imply that mobile RSUs could not set up to indicate a temporary construction zone.

4.3 Non-Safety Applications

Non-safety applications have to transmit the data on a Service Channel. Therefore, a channel switch is required that results in missing safety related messages on the Control Channel. This makes it important to distinguish whether the vehicle is on the street or not—there are no concerns about safety if the vehicle is parked.

It should be noted that no V2V non-safety applications are discussed here. This is due to the fact that no applications have been proposed yet requiring a link between two vehicles.

4.3.1 Internet Access

Nowadays technologies like the 'General Packet Radio System' (GPRS) or the 'Universal Mobile Telecommunication System' (UMTS) are used to access the Internet in a vehicle. These technologies provide good reception coverage, and ideally a continuous connection, even though the station is moving. The main drawback of these technologies is their limited bandwidth—a comparison of the data rates is shown in Table 4.1—and the costs for transferring data. Hence a mobile phone is a good choice

Technology	Details	Bandwidth
GPRS	2.5G	140.8 kbit/s
UMTS	3G	1920 kbit/s
DSRC	Single Channel	27 Mbit/s
DSRC	Dual Channel	54 Mbit/s

Table 4.1: Mobile phones can usually provide a reliable connection to the Internet, but their bandwidth is in general too limited to exchange large amount of data.

to access limited amount of data, but is in general not likely to be used to transfer large data chunks.

DSRC is meant to be a complement to these technologies and not a replacement. In particular, DSRC is not assumed to provide a continuous access to the Internet while driving—the necessary routing and cell handover is assumed to be not realisable.[3]

This should not imply that Internet access is not possible. For instance, a RSU can provide access to the Internet for the time the vehicle is driving in reception range. During that short period, the vehicle can synchronise emails and execute other short time tasks. However, for a continuous connection a technology as mentioned above has to be used.

The situation is different if the vehicle is not on the road. There is no need to exchange safety messages and routing is not an issue either. DSRC can therefore be used similar to 'IEEE 802.11a': A gas station can provide a wireless access point—often referred to as "Wi-Fi hot spot"—to provide access to the Internet, or the vehicle connects to the home network to update its music database.

4.3.2 Applications while Driving

Whenever the vehicle is driving the main purpose of DSRC is safety. Nevertheless, non-safety applications can use the radio to transfer data as well. This can be achieved by switching back and forth between the Control Channel and the Service Channel.

The non-safety applications can be divided into two main categories: 'Short Time Interactions' and 'Drive by Information Fuelling'.

Short Time Interactions

A short time interaction is a data transfer between a vehicle and a RSU that requires some kind of interaction. This data transfer is characterised by the little amount of transmitted data and the short transaction time.

- **Electronic Toll Collection**

  Different technologies exist to automatically pay the toll for designated bridges, tunnels or roads. These technologies have in common that an additional transponder is required to use the system.

  The different toll systems—such as E-PASS, Fastrak or Tolltag—all have their proprietary transponders. DSRC is assumed to standardise the electronic toll collection.

[3]The main purpose of DSRC is to improve safety. This implies that a substantial part of the available time is required to exchange data on the Control Channel. In addition, a reliable connection between two nodes cannot be assumed due to the non-deterministic channel characteristic. Furthermore, it cannot be assumed that RSUs are located all along the roads. These facts imply that a continuous connection to the Internet is most likely not achievable.

- **Local Information**

  The vehicle's passengers are looking for a restaurant or a place to stay overnight. DSRC can be used to request such specific information.

Drive by Information Fuelling

Drive by information fuelling distinguishes if the information is public, i.e. for free, or commercial. The following list should provide an idea about the information that could be provided:

- **News**

  News, such as sports results or the weather forecast, can be broadcasted on a Service Channel. Such a service is likely to be either commercial or ad-funded.

- **Local Announcements**

  This is a typical public service. Places of interest or exhibitions can be announced before entering a city.

- **Real Time Traffic Information**

  A RSU can broadcast the current traffic information ahead. The vehicles route planning software can adapt the current route based on this information. It should be noted that this information is not meant to be provided on the Control Channel.

- **Map Update**

  A road might not be accessible due to construction work. Such information can be provided by a RSU and allows it to adapt the vehicle's map data.

- **Infotainment**

  DSRC can provide all kind of information that is considered to be entertaining—such information is often referred to as "Infotainment". Infotainment can be ad-funded, but is most likely to be commercial. A typical application is to buy a song that was just on air.

The data is usually provided by a RSU and is therefore accessible for a very limited period only—e.g. for about fifteen seconds while driving on a highway. Therefore, if a lot of data is provided, it might be necessary that the data is sent by several consecutive RSUs.

Public information is assumed to be broadcasted continuously on a Service Channel and all interested vehicles can receive the data. This is different with commercial information. Such a commercial data transfer is assumed to consist of three steps:

1. **Requesting the Information**

   The information can be requested at a RSU or in the Internet using a technology like UMTS. This request is likely to contain information about the RSU that will provide the data and a transaction number.

2. **Download the Data**

   As soon as the vehicle arrives in the data-providing RSU's reception range, the vehicle identifies itself with the transaction number. This will initiate the exchange of the (encrypted) data.

3. **Paying**

   As soon as the data is received—this can be checked with a hash code—the transaction can be completed. In order to do that, the key to decrypt the data has to be bought. This ensures that only a completed data transfer has to be paid.

It should be noted, that the first step is not necessarily involved in the transaction. For instance, data of general interest—e.g. sports results—are provided without a request, but have to be paid to access them. Alternatively, the payment can be done by subscription.

Chapter 5

Message Requirements

This chapter deals with the demands the applications have on the safety messages.

The first part of this chapter analyses how DSRC can improve safety. It is shown that accidents are not only caused by a sudden change in the behaviour of a vehicle but occur as well if the involved vehicles are driving in their normal way. This implies that safety messages are not only required to be sent in the case of an event, but also on a regular basis in order to provide the surrounding traffic with the vehicle's current *status*. This status should allow the prediction of the vehicle's behaviour in the near future. The second part analyses two message schemes. A strict periodic message scheme will cause a channel breakdown. A message scheme is therefore presented that contains two message types: so-called *event messages* are sent whenever the vehicle is behaving unpredictably and so-called *routine messages* which are sent periodically to provide the status. The interval of the latter one needs to be adapted in such a way that the channel will not congest for any reason. The third part points out that safety applications do not exchange messages directly with each other. This fact and the important issue of compatibility require a flexible and extendable data structure. The fourth part analyses the message content requirement for the most common V2V applications in order to estimate the message size. The safety data, without any header overhead, is shown to be only a few tens of bytes. The chapter closes discussing security.

5.1 Safety and DSRC

The last chapter introduced some of the intended safety applications based on DSRC technology. In order to design a vehicular safety communication architecture, a clear understanding is required about the hazardous situations that can be avoided using wireless communication. However, DSRC is not meant to be a universal remedy. For instance, DSRC cannot prevent a drunken driver to steer into a ditch, but the surrounding traffic can be alerted about the danger.

DSRC does not exclude safety applications that are meant to be used in combination with autonomous operation of the vehicle per se. However, the communication architecture proposed in this work is meant for safety applications that require the interaction of the driver. Therefore, the proposed architecture is not meant to prevent an accident that is about to occur in a time window, not allowing a human being to react to an alert.

There are two types of scenarios in which safety can be improved using DSRC:

- A hazardous situation does not necessarily imply the involvement of a second vehicle. This is the typical V2I communication that allows the interaction with a traffic light or to advise the driver about a potentially dangerous street condition.

- There are hazardous situations implying the direct or indirect involvement of other vehicles. For instance, a vehicle sensing a hazardous situation on the street—such as aquaplaning—broadcasts this information. However, the most common V2V safety communication has the goal to prevent a collision between two vehicles from happening.

The majority of collisions are based on an event—e.g. a harsh braking car causes a rear-end collision accident. In order to define appropriately what such an event is, it is analysed first how vehicles can collide while driving in a normal and predictable way.

5.1.1 Context Based Hazards

Two vehicles are not supposed to be at the same position at the same time. This can occur even though all traffic participants are driving in a normal, hence predictable, manner. The following three examples present such situations:

Example 1 A vehicle has stopped right behind a tight curve and the hazard lights are not turned on for whatever reasons. The vehicle might have sent safety messages while braking, but the braking-event was over after a few seconds and there is no reason to send further safety messages. After a while, a second car is approaching neither seeing the stopped vehicle nor receiving messages from it. Hence an accident is very likely even though both vehicles are behaving normally.[1]

Example 2 A vehicle is running a red light and collides with another one. If the traffic light is equipped with DSRC, the incorrect driving vehicle knows about the misbehaviour of its driver and broadcasts a warning message. If not, both vehicles are driving normal—from the vehicle's point of view—and there is no reason to broadcast a warning.

Example 3 Figure 5.1(a) illustrates a situation that occurs frequently on the street in the United States. This situation shows two fast vehicles, namely A and B, approaching the slow vehicle C. Driver B is fully aware of the obstacle ahead and intends to pass the slow vehicle. Driver A on the other hand is not aware of the oncoming danger. Since all vehicles are driving straight ahead no event-based messages are sent. Figure 5.1(b) shows the same situation a little bit later. Vehicle B has started the passing manoeuvre at the very last instant and suddenly an obstacle emerges in front of driver A who has to react immediately to avoid an accident.

These three examples have shown that safety messages are required to be sent on a regular basis to inform the surrounding about the vehicle's current status. These messages should contain all the necessary information to predict the vehicle's position over the next few seconds. Therefore, information—such as the current position, speed, acceleration, and heading—of the vehicle's status should be provided.

[1] From a human point of view, a stopped car in the middle of a street is definitely not behaving normal. However, it is the vehicle that has to make the decision.

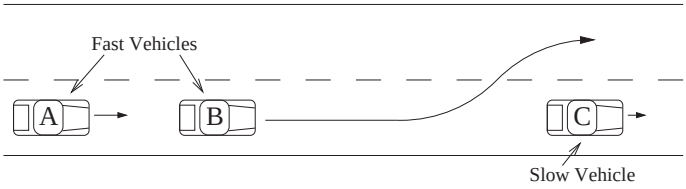

(a) Initial situation.

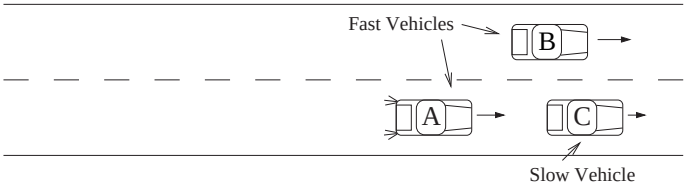

(b) Hazardous situation.

Figure 5.1: A hazardous situation occurs even though all vehicles are driving in a normal way. This situation could have been avoided sending messages on a regular basis.

Receiving status messages from all surrounding vehicles allows the vehicle to get a good impression about the current traffic environment—this collectivity of information about the surrounding vehicles is referred to as *context* in this work. The context can be very helpful in combination with a received event in order to decide if the driver should be notified as the following example points out:

Example 4 An event-based message is received from a hard braking car 200 metres ahead. If there is no other vehicle in between, this is probably not a hazardous situation yet. But this might look different in a heavy traffic situation if a couple of vehicles are in between since these vehicles are likely to brake soon. In such a situation, it might be advisable to alert the driver about the possible danger already.

5.1.2 Event Based Hazards

The occurrence of most collisions is based on the fact that one or more vehicles are not behaving in their usual way and become a danger for other traffic participants. Such a sudden change in the behaviour is usually related with one or more events—such as loss of control over the vehicle, a sharp brake, or intending to change lanes. Generally speaking an event is always related to the fact that the vehicle is not predictable anymore. Hence, a vehicle should always trigger an event whenever its predicted position, based on the regular sent status messages, is jeopardised.

It should be noted that human beings are driving vehicles and they do not operate in a very fast manner. This implies that an event usually lasts for a few hundred milliseconds and sometimes even for seconds. For instance, a driver slamming the brakes does not push the brake pedal for only a few tens of milliseconds.

5.2 Periodic Message Scheme

The V2V safety communication requires alerting the surrounding vehicles with as little delay as possible if an event is triggered. In addition, the vehicles should send message regularly to provide the other vehicles with the current status. Sending safety messages strictly at regular intervals can fulfill these two requirements. These

messages contain information about the status of the vehicle, but whenever an event is triggered, additional event-based information is added to the periodic message.

A periodic message scheme for safety communication has two strong demands on the timing of the message transmission:

1. **Message frequency**

   The information about a triggered event is sent in the next scheduled message. Hence, the maximum delay corresponds to the interval of the periodic messages. This time span must be chosen small enough, i.e. on the order of 100 milliseconds, to ensure that the message will not be delayed for too long.

2. **Message delay**

   A message delay would be tolerable if the vehicle is driving normal—i.e. no event is triggered. However, the periodic message might contain event-based information. And since the event information can already be delayed due to the message frequency, an additional delay is not acceptable.

These two demands may have a severe impact on the channel load. Different channel usage scenarios have been simulated at the DC RTNA. The result can be found in 'VSC Task 12' [4], in which scenario 11 is of particular interest in conjunction with channel load of a periodic message scheme:

This scenario shows a divided highway, four lanes in each direction as illustrated in Figure 5.2(a). A heavy traffic situation is simulated, i.e. 80 vehicles per lane mile, in which each vehicle broadcasts a 200-byte message every 100 milliseconds. The result of this simulation is depicted in Figure 5.2(b), showing that the channel cannot handle the load and breaks down. Within a range of about 80 metres the probability of a successful reception drops to about fifty percent. In a range of 200 metres—this is still an important range for safety communication—the probability of a successful reception is less than twenty percent. A periodic message scheme is therefore not suitable to be used with DSRC.

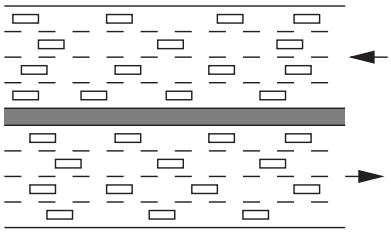

(a) Simulation Scenario: A divided highway, with four lanes in each direction, in a heavy traffic situation (80 vehicles per lane mile). Each vehicle sends a 200-byte message every 100ms.

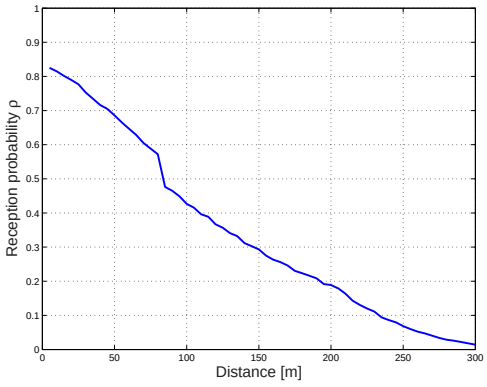

(b) Simulation Result: The channel breaks down under the heavy load of the periodic message scheme.

Figure 5.2: Simulation of the periodic message scheme in a heavy traffic situation.

5.3 Combined Message Scheme

The periodic message scheme requires a short message interval to ensure a short latency in the case of an event occurrence. However, vehicles behave normally most of the time and safety messages therefore rarely contain time-critical information.

In order to prevent a channel breakdown, the message frequency must be adapted according to the current channel load. Such an increase message interval implies that the strict demands of the event message's latency cannot be fulfilled anymore. Therefore, if an event is triggered, a message should be broadcasted immediately instead of adding the event information to the next scheduled message.

This message scheme introduces two different message types, namely event message and routine message, having different demands on the message delay.

5.3.1 Event Messages

A vehicle triggering an event is not predictable for the surrounding traffic any longer. Hence the surrounding, i.e. all vehicles in a certain range, must be signalled about that fact with as little delay as possible.

Maximum Latency

This delay, referred to as *latency*, is the time that passes from the moment the event is triggered to the moment the event is received. The maximum tolerable latency depends a lot on the current situation and depends mainly on the travel speed and the traffic density.

Unfortunately no thorough analysis about the allowable latency is available yet that considers all possible safety applications. In several DSRC related documents—such as 'VSC Task 11' [3]—the maximum allowable latency is assumed to be on the order of $100 - 150$ milliseconds. Compared to the driver's perception time—that is on the order of $500 - 1000$ milliseconds [27]—this seems to be a reasonable value. It should be noted that a vehicle driving 160km/h (∼100mph) travels 4.4 metres in 100 milliseconds.

Distribution Range

An event message has to be distributed within a certain range only—e.g. all vehicles within 200 metres.

In accordance with 'VSC Task 11' [3], most safety applications require the message to be distributed within a range of about 300 metres. However, there are certain safety applications requiring a message delivery up to 1000 metres. It should be noted that the distribution range is typically not symmetrical—safety message are usually meant to be distributed either behind or in front of the driving vehicle.

5.3.2 Routine Messages

Routine messages are sent on a regular basis—at least once or twice every second—to inform the surrounding vehicles about the vehicle's current status. Based on this information, the position of the vehicle is predictable in the near future as long as no event is triggered—at least for a few seconds. This implies that some of these messages can be missed without jeopardising safety.

The necessity of sending messages on a regular basis raises two concerns. One is about privacy and the other is about the additional load that is generated:

Vehicular safety communication raises major concerns about privacy. Sending messages on a regular basis, containing information about the current state of the vehicle, is a delicate issue. For instance, listening to these messages might allow the tracking of a vehicle, or a speeding ticket could be issued due to the vehicles own betrayal. It is therefore necessary that safety messages cannot be allocated to a specific vehicle. Note, this does not necessarily imply that safety messages from the same vehicle are not linkable.

Routine messages add a large amount of additional load to the channel. As discussed previously, the channel can break down under the heavy load of a periodic message scheme. However, routine messages do not have such a strict demand on the message interval as the periodic message scheme and can be adapted to the current channel load.

5.3.3 Message Priority

Event messages have very strict demands on the latency. Routine messages on the other hand are important as well, but shall never jeopardise an effective distribution of an event. In order to tackle that problem, the messages should have a priority assigned to prioritise them accordingly.

It is of paramount interest that all events are distributed to the surrounding vehicles in an efficient way. Nevertheless there are events that are more critical than other ones—e.g. a moderate tapping of the brake pedal compared to a harsh slam on them. It is therefore reasonable to assume that event messages should have different priority levels among themselves.

5.4 Application Communication

In a traditionally layered communication architecture, as introduced in the ISO/OSI [28] reference model, messages are exchanged between specific applications—e.g. a mail client exchanges messages with a mail server. This cannot be assumed for the vehicular safety communication and is discussed below.

There are different automotive manufacturers and many of them offer a wide range of different models. All of these models have their own characteristics, in particular their specific sensor data and dynamic vehicle-behaviour. For that reason, the different vehicles are likely to provide their own adapted safety applications. For instance, a basic class model might only have a general hazard warning application running while a luxury class model provides many different applications taking benefit of the vehicle's sophisticated sensor network. Nevertheless, all of these different safety applications need to be able to understand each other.

Safety messages are broadcasted to all surrounding stations. This unidirectional communication—i.e. no feedback—implies that the sending unit does not know what kind of safety application is processing the data in the receiver unit. Actually, it is likely that different safety applications will analyse the incoming message. Some of them find enough information in the message to process the data; others do not as they require more detailed or different information.

An example of this scenario is shown in Figure 5.3; The sending unit has detected an emergency brake and broadcasts a safety message using DSRC. One of the receiving units has three different safety applications running: *Emergency Brake*, *Hazard Warning* and *Lane Change Warning*. The emergency brake and the hazard warning applications are most likely to find useful information in the message. The lane change warning application on the other hand will ignore the message because there

is most likely no relevant information provided. This example has illustrated that the message cannot be assumed to be of use for one particular application only.

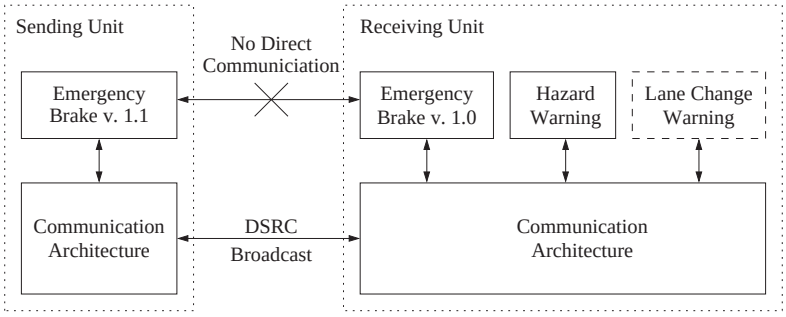

Figure 5.3: The safety applications do not interact using a traditional application-to-application communication. The received information might be of interest for various safety applications.

Another crucial point is that in a few years after launching DSRC technology, second and third generation models will be released running enhanced or new applications. These new applications may need additional data to show their full potential. Hence additional data will be added to the messages that is unknown to the older applications. Nevertheless, it must be ensured that the older applications can pick out the data they are interested in.

For all of these reasons, it is essential to have a flexible data structure that allows every safety application to gather the required data. Such a flexible data structure is presented in Appendix C.1.

5.5 Message Content Requirement

The vast majority of safety messages are assumed to be exchanged between vehicles.[2] It is therefore of interest to estimate the size of these messages, in order to predict the channel impact of the different distribution schemes presented in the following chapters.

Table 5.1 shows a summary of the data that is likely to be required for the V2V applications presented in the previous chapter. It should be noted that this selection should be a guideline only. A description of the data and an encoding scheme is presented in Appendix C.2.

According to Table 5.1, the amount of data that is sent by the different safety applications is very small—i.e. on the order of a few tens of bytes. However, the additional load added by the different headers must be considered as well. In particular the PHY preamble, the MAC header and the security overhead will add an additional 121 bytes to each message. A detailed description of the PHY preamble and the MAC header is given in Appendix B. The security frame in discussed in Appendix A.

[2]All vehicles send routine messages on a regular basis. RSUs on the other hand are located at some designated places along the road only and are therefore not assumed to generate a lot of communication traffic.

Data	V2V Safety Applications			Routine	Size
	EEBL	V2VHW	AEVW	Message	[bits]
Vehicle Identifier	×	×	×	×	24
Time	×	×	×	×	64
Position	×	×	×	×	96
Heading	×	×	×	×	16
Speed	×	×	×	×	16
Longitudinal Acceleration	×	×	×	×	16
Lateral Acceleration		(×)			16
Vehicle Size and Weight	(×)	(×)		(×)	48
Steering Wheel Angle		(×)			16
Break Status	(×)	(×)			8
Signal and Lights		(×)			8

Table 5.1: Data requirement of the different V2V safety applications.

5.6 Security

There are various potential threads correlated with vehicular safety communication. 'VSC-Task 11' [3] presents a few examples of possible attacks:

> For example, if an attacker could impersonate an emergency vehicle, he or she could mount threats with consequences of different magnitudes. The attacker could use the communication to move quickly through traffic, convincing drivers that an emergency vehicle is approaching and causing them to move to the side of the road. This could result in delays and possibly confusion (when they do not see an emergency vehicle) for the other drivers. That same attacker could alternatively plan to preempt traffic signals in a coordinated fashion to make sure that the vehicle of an elected official, for example, would be properly placed for a terrorist attack. This scenario clearly has more serious consequences. Taking another example, if an attacker can impersonate a vehicle in an emergency braking manoeuvre, he or she could cause traffic jams behind them that could serve several purposes. The disturbance could simply be an amusement for the attacker, or it could be used in collaboration with other attackers to temporarily obstruct traffic arteries around a targeted city. Because the potential consequences of an attack vary so widely, they are not quantified in this report.

According to 'VSC-Task 11', the following security services are required for vehicular safety communication:

- **Authentication**

  The receiver should be able to determine whether a received message is trustworthy or not.

- **Validation**

  It must be ensured that an authenticated message is valid—i.e. does not contain faked data.

- **Privacy**

  It should also not be possible to retrieve private information from a vehicle nor to trace one.

- **Confidentiality**

  Confidentiality is not needed since safety messages do not contain sensitive information.

Security in a vehicular communication environment is addressed by the IEEE 1556 working group. According to a pre-draft version of 'IEEE 1556' [13] and in accordance with 'VSC Task 11', safety messages are required to be signed right before they are sent. A detailed description of the message signing and the security frame can be found in Appendix A.

The matter of security is considered as a black box in this work—i.e. the message is signed right before it is sent. However, the matter of privacy must be addressed as certain functions require some kind of identification—e.g. to keep the context up to date. Unfortunately, identification and privacy are contradictory.

Chapter 6

Implications

This chapter deals with the implications that the application requirements presented in the previous chapter have on the communication architecture.

The first part of this chapter analyses routine messages. It shows how channel congestion can be avoided even though routine messages are sent regularly. It further presents an acknowledgement scheme that allows predicting the occurrence of a packet collision. The second part highlights that a cooperative event distribution is much more efficient than a repeated broadcast from a single station in terms of an effective event distribution. The third part presents the so-called *echo mechanism*. Echoing provides the possibility to forward a message without adding a lot of additional load to the channel. It ensures an almost reliable event distribution and can improve the quality of the context. Furthermore a publish/subscriber scheme is introduced to tackle the issue that safety applications do not transmit messages based on a traditional application-to-application communication.

6.1 Broadcast Routine Messages

Routine messages are sent frequently to update the surrounding traffic about the vehicle's current status. This information provides all the necessary information to predict the vehicles position in the next couple of seconds. During that predictable time window a vehicle will broadcast several routine messages.[1] Hence, some of the routine messages can be missed—due to the channel quality or channel switching—without jeopardising safety.

6.1.1 Congestion Control

In Section 5.2 it was shown that the periodic broadcast of safety messages can cause a channel breakdown. It is of paramount importance that the Control Channel does not congest by any means and is furthermore capable of handling the additional load that is caused by a series of events. How exactly this is achieved is beyond the scope of this work. Nevertheless, a brief discussion about *congestion control* is necessary to understand its functionality. It should be noted that the congestion control is also referred to as *Routine Message Generation Control* in this work.

The channel load caused by routine messages depends mostly on the transmission power and the interval of these messages. The former parameter determines the

[1] Even in a heavy traffic situation it is assumed that vehicles can send one or two routine messages per second without congesting the channel—although with a reduced transmission power.

range in which the routine messages are received whereas the latter determines how often the messages are sent. Current research at the DC RTNA shows that these two parameters combined with the current vehicle density set the message density on the channel. Therefore, it does not make a difference, in terms of channel load, if routine messages are sent in a wide range in combination with a long message interval or in a short range with a short interval.[2]

In order to allow the vehicles to capture the surroundings as accurately as possible, the message density should be maximised without jeopardising the event distribution. It should be noted that the distribution of an event and the channel switching (discussed in Chapter 7) benefit from a certain message density as well.

Routine messages usually have the same importance for all the vehicles within a certain range as they are usually driving with a similar speed in a predictable manner. Therefore, the different vehicles should send their routine messages with a similar parameter setting, i.e. message interval and transmission power, in order to ensure fair channel access. This should not imply that a few exposed vehicles, e.g. speeding ones, are not adapting their parameter setting accordingly. Such a fair channel access must be achieved in a distributed manner since no superior node is available. How exactly this is done is beyond the scope of this work but it can be assumed that the current parameter setting is attached to the safety message.

6.1.2 Push or Pull Routine Messages

The congestion-control unit determines the point in time the next routine message is sent. This can be achieved in two ways:

1. **Push-based transfer**

   Routine messages are generated on a regular basis and are passed—i.e. pushed—to the congestion-control unit that holds the most recent routine message in a buffer. As soon as the channel is capable of handling another routine message, the buffered message is sent. This implies that a message is usually buffered for some time before it is sent, and some of the created messages will be overwritten in the buffer.

2. **Pull-based transfer**

   The congestion-control unit asks for a new message to be generated whenever a new routine message can be sent.

The pull-based approach is assumed to be the better choice, as the messages are always up to date and are not generated unnecessarily.

6.1.3 Message Feedback

The general reception probability is similar to the one depicted in Figure 2.2(a). However, whenever a packet collision occurs—which is usually caused by the hidden terminal effect—the reception quality worsens. A feedback mechanism might enable detection of such a collision and allow rebroadcasting of the packet in order to improve the average reception quality.

[2] This is an approximation only.

Acknowledgement Scheme

The broadcast nature of safety messages does not allow sending an acknowledgement for every received message—the communication channel cannot possibly handle this additional load.

Instead of generating additional packets, the acknowledgement may be attached to the next safety message that is sent. Such an approach is referred to as *piggybacking*. However, this feedback scheme is not ideal for a heavy traffic situation in which dozens of messages have to be acknowledged, which can overly increase the message size.

As a result, the number of piggybacked acknowledgments must be limited and selected carefully to maximise their impact. For instance, failure to receive an acknowledgement from a vehicle 200 metres ahead is not a good indication for a message collision, since the packet is very likely to get lost due to the bad reception quality at that distance.

Current research at the DC RTNA proposes that acknowledgements from vehicles driving $80-110$ metres away from the sending vehicle provide the highest significance for detecting a message collision.[3] Therefore, the acknowledgment shall be limited to vehicles driving in this range. Furthermore, the number of messages should be limited to about five to ten messages so as not to increase the messages size in a disproportional manner.

Based on the received acknowledgement it is possible—according to current research at the DC RTNA—to detect a packet collision due to the hidden terminal effect.

Message Identifier Length

Safety messages need to be identifiable in order to acknowledge them. This can be achieved with a random message identifier that is long enough to ensure a very high probability of uniqueness.

It is important to clarify what uniqueness means in this context: the identifier has to be valid as long as an acknowledgement can be expected on the channel. As mentioned above, the vehicle will acknowledge at most ten messages received from vehicles driving the proper distance away. Therefore, the identifier must be unique for a certain number of messages—referred to as the *message window*. It is assumed that a window of about one hundred messages should be sufficient.

Table 6.1 shows that the expected number of identifier collisions depends on the identifier's size and the message window. The calculation of the blacklisting value is based on the assumption that all vehicles are constantly listening to the Control Channel. However, if some stations cannot maintain a proper blacklist—they might be communicating on a Service Channel—the number of expected identifier collisions is higher. Therefore, the length of the identifier must be based on the assumption that proper blacklisting is not possible.

A one-byte identifier is likely to cause an identifier collision and is probably too short. A two-byte identifier will result in an identifier collision every few thousand messages and even less if blacklisting is done in most vehicles. Considering the purpose of the acknowledge scheme—that is to 'estimate' a message collision—an incorrect feedback every thousandth message is justifiable.

[3] Vehicles very close by are unlikely to experience a message collision and vehicles farther away will miss a lot of messages due to the channel characteristic.

Message Window	Identifier Length [bits]	Number of Identifier Collisions in 1'000 Messages	
		w/o Blacklisting	Blacklisting
50	8	195	5.90
50	16	0.763	$23.0 \cdot 10^{-3}$
50	24	$2.98 \cdot 10^{-3}$	$45 \cdot 10^{-6}$
100	8	391	11.8
100	16	1.53	$46.1 \cdot 10^{-3}$
100	24	$5.96 \cdot 10^{-3}$	$180 \cdot 10^{-6}$
200	8	781	23.6
200	16	3.05	$92.2 \cdot 10^{-3}$
200	24	$11.9 \cdot 10^{-3}$	$360 \cdot 10^{-6}$

Table 6.1: Number of message-identifier collisions based on the length of the identifier and the number of messages that needs to be distinguished. The value for the blacklisting is calculated in a deterministic range based with a fifty percent reception probability.

Event Message Acknowledgement

It should be noted that event messages have their own, more sophisticated acknowledgement scheme. A detailed description about event indications is held in Section 7.6.

6.2 Broadcast Event Messages

6.2.1 Distribution Area

Most event messages are distributed within a range of about 300 metres, but there are certain messages that should cover a distance of up to 1000 metres. This latter range is impossible to be covered by a single station and a multi-hop message distribution is required.

This task is commonly referred to as "geocast flooding". One major concern about flooding is that it is likely to add a lot of additional load to the channel, resulting in a channel breakdown if many vehicles send event messages due to a chain effect. There are different approaches for geocast flooding—see [26] for details. However, it should be noted that all of these protocols generate a lot of additional packets and are therefore not efficient in terms of channel load.

6.2.2 Repeated Broadcast

The vast majority of event messages need to be distributed in a range less than 300 metres. Considering an exemplary reception characteristic, as depicted in Figure 2.2(a), this is at the limit of a DSRC's broadcast range. There is usually a high probability—90 percent or more—for a message reception in a range of about 100 metres. However, farther distances degrade the quality of the signal severely with the distance from the sender.

Broadcasting the message multiple times would increase the reception probability, as indicated in Figure 6.1. Sending the message twice can almost guarantee a successful reception in a range of about 100 metres. Farther away increases the chance that vehicles will not receive the message. To extend the reception range, the message can be repeated several more times. However, even sending the message eight times cannot ensure coverage of more than 200 metres.

Hence a reliable coverage for a wider range than 200 metres is generally not achieved by a single unit and implies the involvement of other stations.

6.2.3 Message Forwarding

Safety messages need to be distributed along the street. Hence the region where a message must be distributed is mostly one-dimensional, i.e. along the road, and not circular.

This makes an efficient message distribution using multiple hops fairly easy—the message has to be forwarded along the road—and no complex routing algorithm is necessary. The effect of such a forwarding is illustrated in Figure 6.2; the message is repeated by a vehicle that is 125 metres and a vehicle that is 250 metres away. These three broadcasts can almost ensure a reliable message distribution in a range of about 300 metres. This is much more efficient than the repeated broadcast that requires up to ten messages to cover a range of 200 metres. It should be noted that the repeated broadcast covers a symmetrical area, while the forwarding covers only one direction, which is generally all a forwarding situation requires.

The event triggering application should provide the communication network the information of the range in which the event should be distributed. Based on this information a station receiving the message decides whether to forwarded the message or not. The event distribution is further discussed in Section 7.5.

6.3 Message Echoing

The basic idea of message echoing is that a safety message is enhanced with the safety content, i.e. without security and other overhead, of a recently received safety message. Subsequently, the impact of the echoing on the message size and the mechanism itself is discussed.

6.3.1 Message Size

The size of the packets should be as small as possible to increase the possibility of a successful broadcast. The message-echoing's impact on the message size must therefore be examined carefully. Table 5.1 presents an overview of the data provided in a safety message. There are two important facts to note:

- The safety information of a V2V safety message is small (~ 40 bytes).
- The largest chunks of the safety information are the position (12 bytes) and time (8 bytes).

The required size for the position and the time can be reduced significantly by sending this information relative to the position and time of the echoing vehicle. The relative position can be encoded with 5 bytes[4] and the relative time with 2 bytes.[5] Hence the $30 - 40$ bytes of the security message can be reduced to less than 30 bytes on average. It should be noted that the message overhead generated in the MAC and the PHY is 121 bytes.

[4]Two bytes each for the relative longitude and latitude and one byte for the altitude.

[5]The delay between the calculation of the relative position/time and the absolute position/time added at the signature (see Appendix Section A for details) is assumed to be at most a few milliseconds. A vehicle is usually not driving more than 4cm/s, hence this minor inaccuracy is assumed to be negligible.

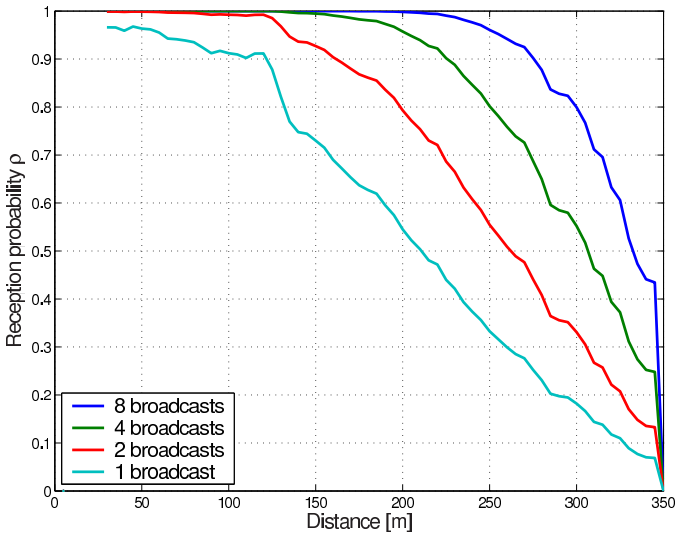

Figure 6.1: Repeated broadcasts of an event message improves the probability of a successful reception.

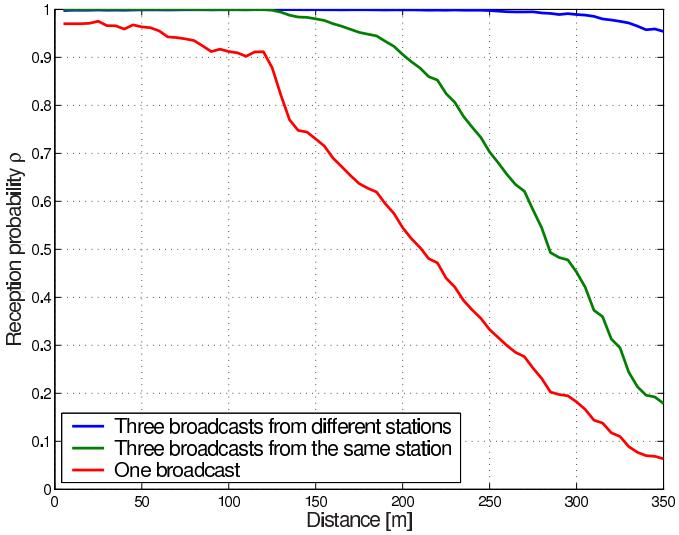

Figure 6.2: There are two approaches to improve the reception probability: The message may be repeated by a single station or repeated by different stations. The latter shows a much better performance, but is unidirectional only.

6.3.2 Echo Mechanism

The echo mechanism can be used to achieve two tasks: ensure an effective event distribution and improve the quality of the context.

Event Echoing

A received event message should contain information about the intended broadcast range and the time window in which the message should be distributed. Therefore, instead of forwarding the message, i.e. generating a new message, the event shall be echoed in the next scheduled routine message. This allows forwarding the event without generating a new message and therefore not loading the channel significantly. The echoing and the event distribution are further discussed in Section 7.5.

Routine Echoing

As there is usually no event on the channel, a routine message should be echoed instead. In that case, on average every routine message is sent twice, resulting in a major improvement in the context quality.

In general, it is of more interest to know what is happening in front of the vehicle and not behind it. Hence, it is assumed that echoing a message received from a vehicle driving about $100 - 150$ metres ahead has the biggest impact on safety improvement. However, the routine-message selection for the echoing is beyond the scope of this work and is not discussed further.

6.4 Safety Application Abstraction Level

Safety applications do not transmit messages using traditional application-to-application communication. As illustrated in Figure 5.3, a safety message can be broken down into so-called *safety information units*—such as position, speed or braking status of the vehicle.

The different safety applications are usually not interested in all the provided information but only a few of these information units. There are two approaches to deal with that issue: Firstly, the received message can be forwarded to all safety applications. Hence all applications must analyse the entire message and filter the relevant information, resulting in an unnecessary processing overhead. Alternatively, a single application parses the incoming message and forwards the relevant information units. This second approach is commonly referred to as the *subscribing mechanism*.

6.4.1 Subscribing

The subscribing mechanism is depicted in Figure 6.3 in the receiving unit and works as follows: Each safety application subscribes to the relevant data. In the example shown in Figure 6.3 the application A_4 is interested in the data D_I and D_{IV} while application A_5 subscribed to the data D_I and D_V.

The incoming message contains the data D_I, D_{II}, D_{III} and D_{IV}; hence application A_4 can be provided with all the requested data. Application A_4 on the other side can only be provided with the datum D_I since the datum D_V is not part of the message.

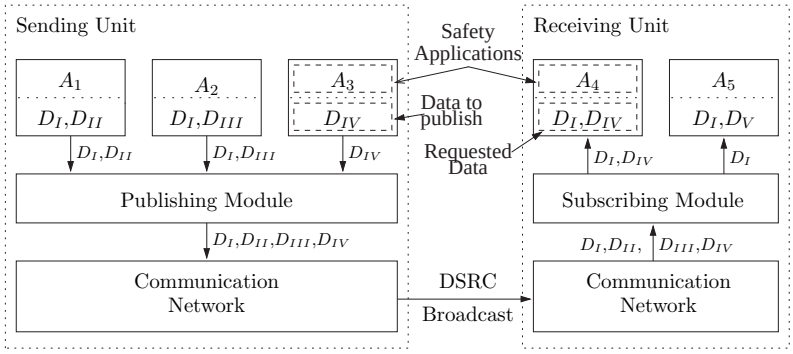

Figure 6.3: The left side of the illustration shows the sending station with the publishing module. The right side shows the receiving station and its subscribing unit.

6.4.2 Publishing

Publishing is the opposite task of the subscription and is depicted in Figure 6.3 in the sending unit. In this example the three applications A_1, A_2 and A_3 need to publish some data. For A_1 this is the data D_I and D_{II}, for A_2 this is the data D_I and D_{III}, and for A_3 this is the data D_{IV}.

The publishing module can combine all the data into one single message containing only one copy of the data D_I, D_{II}, D_{III} and D_{IV}. Hence, there is no redundant data in the message and only one message, instead of three, must be sent.

6.4.3 Compatibility

The publish/subscribe mechanism is based on the assumption that safety messages are built up from small information units that must to be identified. An outmoded DSRC device may not understand all the data sent from a more recent device, but as long as it is guaranteed that the message can be broken down into its safety information units, the outmoded device is able to pick out the data it requires. Detailed information about the safety message format can be obtained in Appendix C.

Chapter 7

Channel Switch

This chapter deals with the issues related with channel switching. In particular, a channel-switch mechanism in combination with an effective event-distribution and indication scheme is presented.

7.1 Objective

Stations exchange safety related messages on the Control Channel, but might demand non-safety communication on a Service Channel. This requires switching the channel, resulting in the temporary cessation of any safety related communication on the Control Channel.[1]

Services are usually provided by RSUs and are therefore available for a very limited time only—e.g. for about $10-15$ seconds while driving on a highway. During that period, the non-safety application is interested in a high throughput and hence in maximising the communication time on the Service Channel.

A channel-switch protocol is therefore required to deal with the following two contradicting goals: ensure safety and maximise the service usage time. In order to design an appropriate protocol, it is necessary to get an idea of the availability and usage behaviour of the provided services:

Even if DSRC is fully deployed, it is not assumed that all streets are covered with service providing RSUs. It is rather assumed that RSUs are located at some designated spots only where they are assumed to have the biggest impact. In particular, a RSU is unlikely to be placed close to a spot that is susceptible to an accident. This implies that the channel-switch scheme should only allocate time for non-safety communication if a service is available.

There are basically two types of non-safety applications: 'short time interactions' and 'drive by information fuelling'. Short time interactions only require a little transaction time, but it is important that they take place—e.g. it must be ensured that all vehicles can register at a toll station. The drive by information fuelling on the other side might require much more time to finish the transaction, so it is interested in a high bandwidth. In addition, it is important to distinguish whether the provided information is public or not. In the former case a lot of vehicles might be accessing the provided information at the same time.

[1] It is assumed that the majority of vehicles will be equipped with one radio only.

7.2 Channel-Switch Scheme – Overview

In contrast to the presented channel-switch schemes in Section 3.3, the one proposed in this work is not based on a synchronisation. This means that the vehicles are deciding independently if the channel can be switched and for how long. In particular, there is no agreement between the vehicles on how they time there channel switches.

The proposed channel-switch scheme—referred to as *independent channel switching* (ICS)—is illustrated in Figure 7.1. The principal idea of the scheme is the following: the time is broken down into safety and non-safety operation intervals in a non-synchronised manner for the different vehicles. This implies that the surrounding traffic is still doing safety communication during the vehicle's non-safety operation interval. After a certain period, the vehicle switches back to the Control Channel to check for important safety messages, in particular for ongoing events. As soon as safety is ensured, the vehicle switches to the Service Channel again to continue the non-safety communication.

In order to maximise the non-safety throughput, the non-safety operation intervals need to be maximised while minimising the safety ones. It is discussed in detail in the remainder of this chapter, how this objective is achieved without jeopardising safety.

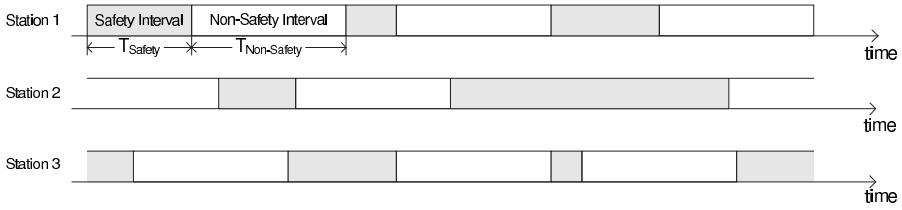

Figure 7.1: Channel-Switch Scheme: The time is broken down into safety and non-safety operation intervals. The operation intervals of the different vehicles are not synchronised.

7.3 Non-Safety Operation Interval

A station communicating on a Service Channel does not receive any safety related messages. Missing some of the routine messages does not jeopardise safety, but missing an event message is critical and must not been delayed for more than a certain amount of time. There are basically two approaches to deal with that problem: The event messages can be forwarded to the Service Channel or the non-safety operation interval is to be limited in combination with a repeated broadcast of the event on the Control Channel.

7.3.1 Message Forwarding

RSU Forwards Messages

A RSU equipped with two radios can listen to the Control Channel and forwards important safety messages to the Service Channel. However, the RSU is not located at the same physical position as the vehicle and does not receive all of the event messages the vehicle would have—e.g. the vehicle has already passed the RSU and a vehicle even farther ahead is braking hard. This forwarding scheme is therefore not reliable.

Virtual Simucast

A vehicle that triggers an event is supposed to broadcast the event message on the Service Channel as well. There are two factors making this approach unfeasible:

- There are six distinctive Service Channels and all of them can provide a service. Hence, the station has to poll all six channels to broadcast the event. This time should rather be spent on the Control Channel in order to ensure an effective event distribution.
- A single broadcast of the event on the Service Channel cannot be assumed to be received by all stations. And there is no time to ensure a certain effectiveness.

Cooperative Message Forwarding

Prior switching to a Service Channel, a vehicle close by can be asked to forward important messages to this distinctive channel. The cooperative message forwarding has the disadvantage that an agreement is necessary that requires sending several messages on the Control Channel. This is neither efficient nor possible without loading the channel too heavily.

7.3.2 Limit Non-Safety Operation Interval

Since reliable message forwarding is not achievable, the vehicle is required to switch back to the Control Channel to check for critical safety messages. This check must be achieved within a time limit that ensures safety.

In the worst case the event is triggered right at the beginning of the non-safety operation interval. Hence, the length of this interval must be less than the maximum allowable latency for an event. This latency is, according to Section 5.3.1, on the order of 100 milliseconds. However, this limit is based on a worst-case scenario and a much higher latency, on the order of $200-300$ milliseconds, would be tolerable in most cases. The vehicle has a detailed view of the surrounding—provided by the context and the vehicle's sensor network. Based on this information the current allowable latency and therefore the appropriate length of the non-safety operation interval can be calculated.

It should be noted that traffic usually does not change in a rapid manner; the length of non-safety operation intervals therefore does not change rapidly either.[2]

The remainder of this section is for completeness only as the presented details are not required for the channel-switch scheme.

It might happen that a prior switch to the Control Channel is necessary. Several reasons could initiate such an early switch:

- **Priority Cancellation from Application**

  Whenever an event is triggered, an immediate switch to the Control Channel is likely to announce the event without any further delay.

- **RSU Suspends Service**

  A RSU that suspends a service should announce this intention.

[2] This assumption might not hold in a case of an emergency. Events are likely to occur in burst, so it might be advisable to decrease the non-safety operation interval after receiving an event.

- **Application Completion**

  All service-using applications finish their transactions.

- **Service not Available**

  Some services imply some kind of interaction between the involved stations. Such a service has usually a limited number of concurrent users. Therefore, if the request for a service is rejected, or not answered at all, the radio should leave the Service Channel and try to access the service later.

- **Service Channel is Idle**

  The Service Channel should be left if the channel is idle for a long time.

7.4 Safety Operation Interval

During the safety operation interval, three tasks must be achieved to ensure safety:

- Check for ongoing events
- Send routine messages
- Update the context

These three tasks are discussed in detail subsequently.

7.4.1 Check for Ongoing Events

In this work it is assumed that a safe switch from an event-check point of view is always possible if the vehicle can ensure that all current events have been received prior to leaving the Control Channel. An exceptional case is if the vehicle itself triggers an event. How non-safety communication can be done in this rare case requires further studies and is beyond the scope of this work. However, it is assumed that a short switch—on the order of fifty milliseconds—is possible to perform high priority non-safety communication, e.g. electronic toll collection.[3]

Several time-critical event messages might have been sent during the non-safety operation interval. The vehicles are therefore required to check the Control Channel frequently for ongoing events. This check has two demands: It must be ensured that all ongoing events are captured in a timely manner and that the communication on the Service Channel should be continued as soon as possible. In particular, it must be ensured efficiently that all events are captured in order to continue the non-safety communication. There are two options to deal with those issues:

1. **Announcement of being away**

   The vehicle sends a routine message announcing that the vehicle was not listening to the Control Channel for a certain time. Vehicles that have sent an event message meanwhile are supposed to repeat the event. Alternatively, a vehicle close by could be asked to provide a safety message comprising all resent events.

   However, these announcements schemes are based on a reliable communication that cannot be assumed with DSRC—i.e. it can neither be assumed that the announcement is received nor the rebroadcasts of the missing events.

[3] Otherwise it would be possible to cheat the system by turning on the hazards lights.

2. **Repeated Broadcast**

   The event messages are broadcasted repeatedly on the Control Channel. In addition, all vehicles are supposed to indicate all received events in their routine messages.

   A vehicle that checks the Control Channel has to screen the routine messages for indicated events that have not yet been received. Whenever such indications are observed, the vehicle is required to wait for these messages to be sent again before continuing the non-safety communication.

Announcement of being away is not feasible. Hence, the repeated broadcast mechanism in combination with the indication scheme will be used. As indicated, the vehicle has to wait for the missed event of being sent again. This implies an increased latency and might jeopardise safety. It is discussed in Section 7.5 how this delay can be minimised and Section 7.6 shows how the indication scheme works in detail.

7.4.2 Send Routine Messages

The channel-switching vehicles are still required to send routine messages with roughly the same interval as with doing safety communication only—the surrounding is still required to get updated regularly. This implies that routine messages must not necessarily be sent in every safety operation interval—e.g. in a heavy traffic situation. It should be noted that the routine message interval could be substantially smaller than the non-safety operation interval. In that case the interval should be increased to a certain maximum that still ensures safety in order to not unnecessarily limit the non-safety throughput.

The routine message should neither be sent at the very beginning nor at the very end of the safety operation interval—sending at the beginning would have an impact on the message echoing while sending at the end would jeopardise the routine message acknowledgment. It is therefore important that the routine message generation control and the channel-switch unit interact with each other.

7.4.3 Update Context

A routine message allows the prediction of the vehicle's position over the next few seconds. This requires receiving such a message from each surrounding vehicle once every few seconds.

Routine messages are assumed to be sent at least once or twice every second. Therefore, one might argue that receiving every third routine message or so ensures that the context is up to date, and one has therefore to spend about thirty percent of the time observing the Control Channel. However, this approach does not work if some of the neighbouring vehicles do non-safety communication as well. For instance, three vehicles are doing non-safety communication, all of them communicating two-thirds of the time on a Service Channel. These three vehicles could be out of sync, resulting in not receiving a single routine message from each other.

The out of sync issue can be tackled doing a complete context update every one or two seconds—i.e. providing a longer non-safety operational interval once in a while. The length of this interval can be assumed to be on the order of the current routine message interval. However, the length of the non-safety operational interval should not be set according to the routine message's frequency, but should depend on the datedness of the context.

Most services are limited to a certain number of concurrent users, implying that only a few vehicles are not exchanging data on the Control Channel at the same time. However, this cannot be assumed for an information-providing service, e.g. map update, that is potentially used by all vehicles in range. In such a situation, a proper context update is not assumed to be achievable in a reasonably short time as all vehicles might spend most of the time doing non-safety communication. It is therefore suggested that information-providing services are suspended once in a while ($\sim 1\text{s} - 2\text{s}$) to ensure that all vehicles switch back and stay on the Control Channel in a coordinated manner.

It should be noted that a longer safety operation interval is not only necessary to keep the context up to date, but also to ensure that V2I safety messages are received.

7.5 Event Message Distribution

An event must be sent on the channel frequently in order to minimise the waiting time. This is achieved based on the echo mechanism as introduced in Section 6.3.

7.5.1 Distribution Parameters

A station that receives an event is supposed to echo it in the next scheduled routine message. It should be noted that all stations echo events, regardless if the event is received directly or via an echo. This results in a flooding of the event. In order to prevent the event of being spread out in an uncontrollable manner, the event message is expected to provide the following distribution parameter:

- **Distribution Range**

  All vehicles in the event's distribution range are supposed to echo the event. For instance, a breaking vehicle might have a distribution range of 250 metres to the back.

- **Time Window**

  The vehicle should echo the message in a given time window only, in order to avoid old events of being echoed on the channel. This time window is not assumed to be larger than a few hundreds milliseconds—the data in the echo becomes obsolete.

- **Event Identification**

  The indication of events requires them to be identifiable. This is achieved by assigning an *event identifier*—a random number large enough to ensure uniqueness with overwhelming probability.

  The event identifier needs to be unique for a certain period only—on the order of a couple of seconds. During this time, a few tens of events are assumed to be triggered at most in the surrounding. In order to minimise the probability of a collision, all current event identifiers should be blacklisted. This blacklisting is assumed to be very effective due to the echo and the indications of the events. It is therefore assumed—this must be verified—that a two byte event identifier provides uniqueness.

- **Sequence Number**

  A vehicle that triggers an event is likely to send more than one message about this very event. These newer messages do still describe the same event, but more recent information about the vehicle's current status is provided. Every

time the station sends a new message about the same event, a so-called *sequence number* should be increased—initialised to zero sending the event the first time. This sequence number should ensure that only the most recent information about the event is echoed.

Due to the echo mechanism, a single event is received frequently but should not be processed more than once in order to avoid an unnecessary processing overhead. It is therefore required that the vehicles maintain a list with the received events and the current sequence number. A received event should only be forwarded to the safety applications if either the event identifier is unknown or the sequence number of a known event indicates an update.

In addition, there should be the possibility that safety applications can indicate that they are no longer interested to get an update of a specific event—e.g. event messages from a braking vehicle behind. This has the effect—as discussed in the next section—that a channel-switching vehicle does not have to wait for a particular event that was indicated with a higher sequence number.

The echoing results in a flooding of the event in the intended range and time without adding a lot of load to the channel and is assumed to be very effective. However, this effectiveness is based on a certain message density on the Control Channel that cannot be assumed in a sparse traffic situation. Such a sparse traffic situation implies a moderately loaded channel and allows generating additional messages without congesting it.[4]

The generation of such new messages should be based on a countdown mechanism and works as follows: A received event initialises a timer based on the vehicles suitability to forward the message. If the timer expires before the event is received again, the message should be echoed immediately. In particular, the event triggering station needs to carefully check the occurrence of its own event on the channel in order to ensure an effective event distribution.

7.5.2 Event Selection

Events do frequently occur in bursts due to chain effects in the traffic. It is therefore likely that a station that is about to send a routine message has different events in a buffer that are eligible to be echoed. In order to restrain the message size, only one event should be echoed, requiring the selection of the most suitable one.

The event selection has the general objective to make the collectivity of the event distributions reliable. This is assumed to be achieved considering the following parameters:

- **Priority**

  Event messages have different priorities. The higher the priority of the event is, the more likely the message should be to be echoed. Hence a high priority event occurs more often on the channel than a low priority one, but all of them must occur on a regular basis.

- **Occurrence**

  The longer an event has not been announced on the channel, the more important the event is to be sent again.

[4] The additional messages should not just forward the event, but rather send a routine message echoing it.

- **Event Indication**

  As discussed in the next section, the vehicles are supposed to indicate received events. An event that is not indicated from vehicles in the event's distribution range should be more likely to be chosen.

- **Position**

  The closer the location of the vehicle to the edge of the distribution range, the less important the echoing of the event is.

- **Time**

  The closer to the end of the time window, the less important the event is.

The exact definition of the selection process requires in depth studies and is beyond the scope of this work.

7.6 Event Indication

An event-checking vehicle does not know if events have been missed while doing non-safety communication. Events shall therefore be indicated in all safety messages, in order to provide the event-checking vehicle with the information whether events have been missed or not.

7.6.1 Event Selection for the Indication

A vehicle should not indicate all received events, but only the ones that are in the event's distribution range. Hence, the vehicle sending the indication is also eligible to echo the event. This ensures that no indications are received referring to events that are out of the event's distribution range.

Events are likely to occur in bursts. To avoid too many events being indicated in one message—resulting in a large overhead—the number of indication should be limited to about ten messages.[5] This is assumed to be a very rare, but delicate case from a safety point of view. So a channel-switching station receiving that many indications is supposed to listen to the Control Channel for a certain minimum of time, to ensure it gets updated.

An event does not last forever. Actually every distributed event contains information about its life span. After this time, the event will not be echoed any longer and must therefore not be indicated anymore.

7.6.2 Indication Scheme

All relevant events—according to the event's distribution range and life span—are indicated in the safety message. A station checking the Control Channel for missed events has to observe these indications for events that have not yet been received.

How many messages are required to be received to ensure a reliable indication of all current events? There are two problems related to this question:

1. **Time Window**

   A vehicle sending a safety message might be doing non-safety communication as well and does not therefore know about ongoing events. In order to tackle that problem, the event indication has to be amended with the information of

[5] A station that has more than ten messages to indicate should choose the ten most recent ones.

how long the indicating station has been listening to the Control Channel. A vehicle checking for events has to verify if this time overlaps with its previous safety operational interval.

2. **Position**

   Two vehicles driving close by are likely to be both either inside or outside of an event's distribution range and can therefore be expected to receive the same event messages. Therefore, event indications from vehicles close by are more relevant than the ones from vehicles driving at a certain distance. On the other hand, if indications are received from vehicles driving in the front and in the back, all necessary indications are received as well—this is due to the overlapping distribution ranges.

   The relevant positions to receive indications are depicted in Figure 7.2. Three ranges are distinguished:

   - All indications received from vehicles farther away than 150 metres shall not be considered. This is due to the fact that vehicles in this distance are unlikely to be in the distribution range of the same events.
   - The event-checking vehicle can assume to have received all relevant indications if a message is received from a vehicle within 50 metres. This is depicted in Figure 7.2(a).
   - The event-checking vehicle can assume to have received all relevant indications if a message is received from a vehicle driving within 150 metres in the front and a message from a vehicle driving within 150 metres in the back. This is depicted in Figure 7.2(b).

   It should be noted that the suggested distances (50m, 150m) are provided for the purpose of a better illustration only. It is actually assumed that these distances depend a lot on the current surrounding.

This second demand implies that a single indication (from a vehicle at the right position that is listening to the Control Channel for the proper period) is sufficient to ensure that all ongoing events are received. It should be noted that this implication is based on the assumed effectiveness of the event distribution due to the flooding effect of the echo mechanism.

As soon as all missed events are received, the vehicle can safely continue the non-safety communication from an event-check point of view. This procedure is illustrated in detail in Figure 7.3. It should be noted that there is a 'Start' but not an 'End' point. This is due to the fact that a station that has received all ongoing events might still have to send a routine message or to update its context. Therefore, the indication process does not end with receiving all events, but with the actual channel switch.

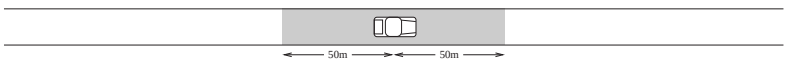

(a) A safety message needs to be received from a vehicle driving close by...

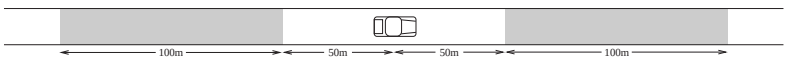

(b) ...or from one vehicle driving in front and one driving in the back in order to ensure that all current events are indicated.

Figure 7.2: Relevant range for event indications.

It might happen that an indicated event is never received. The stations are not expected to wait endlessly for such an event. After a certain period, on the order of 100 milliseconds, an indicated event can safely be assumed to not be of relevance anymore; this is due to the suggested event distribution scheme that ensures a frequent occurrence of the event on the channel.

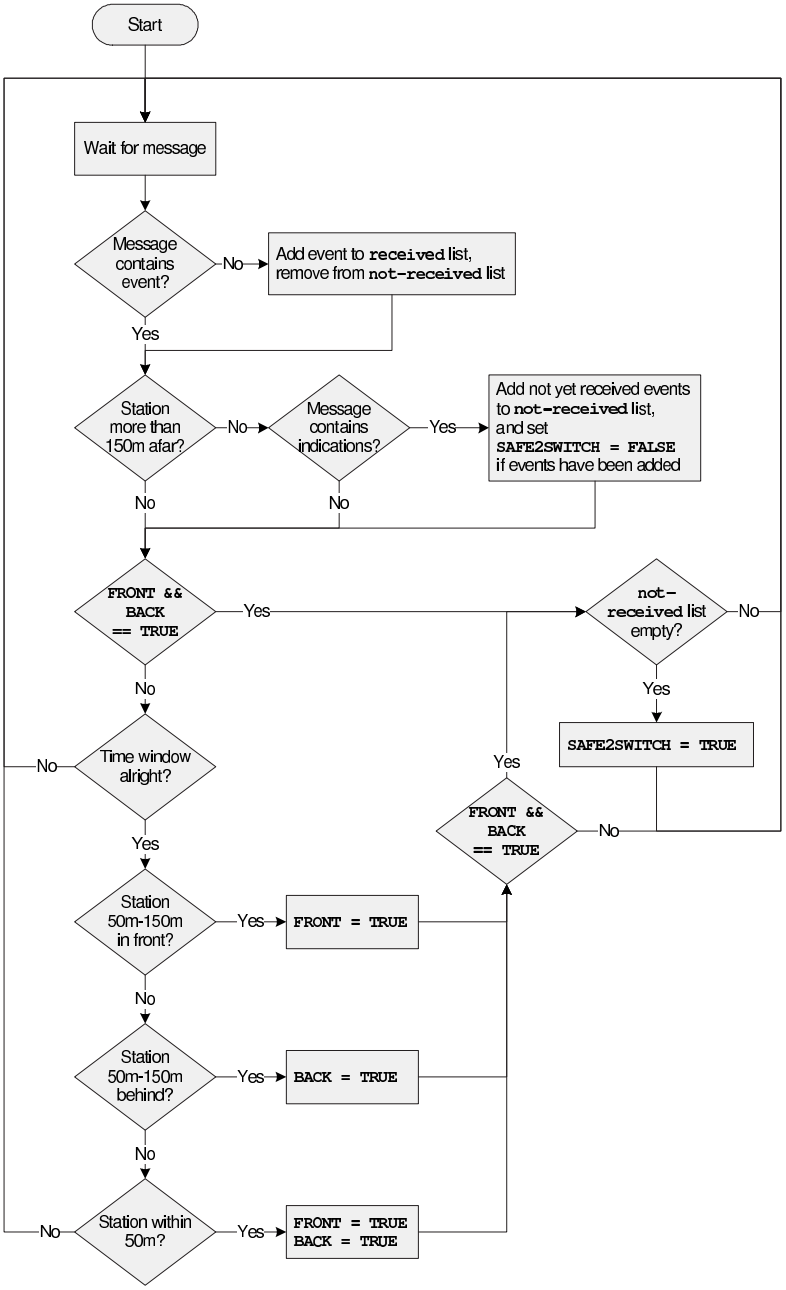

Figure 7.3: Channel switch indication scheme.

7.7 Service Availability

Safety messages are exchanged on the Control Channel, but services are provided on a Service Channel. So how would a station know if a service is provided on one of the Service Channels and if an application is interested in using it?

There are two approaches to deal with those issues: either polling the Service Channels regularly or announcing services on the Control Channel.

Polling the Service Channels

Whenever one or more applications are interested in using a service, the station polls the different Service Channels periodically to figure out if one of the required services is provided. There are six different Service Channels and all of them need to be screened for at least a couple of milliseconds. Hence, a lot of time is wasted checking for services that are most likely hardly ever offered. This time should rather be spent doing safety communication on the Control Channel. Polling the Service Channels is therefore not a suitable approach.

Announce Service on the Control Channel

The second approach is to announce services on the Control Channel sending so-called *channel-switch announcements*. These announcements should provide all necessary information about the offered service, in particular the channel the service is offered in.

As was seen in Chapter 2, the non-safety traffic on the Control Channel is strictly limited in terms of transmission duration ($750\mu s \sim 500$ bytes) and interval (100ms). Therefore, the channel-switch announcements are assumed to have only a minor impact on the channel load.

It should be noted that channel-switch announcements should have the same demands on security as safety messages. Otherwise an attacker could send faked channel-switch announcements provoking the vehicles to leave the Control Channel even though no service is provided or could offer a non-existent high priority service on another Service Channel in order to compromise a low priority service.

7.8 Comparison

A channel-switch scheme has two major goals: ensure safety and maximise the non-safety throughput. The ICS and the synchronised switching schemes, as presented in Section 3.3, are compared subsequently according to these two goals.

7.8.1 Safety

Safety is most important and must be ensured. This is achieved if the channel-switch scheme ensures an up-to-date context and a timely delivery of event information.

It is assumed—based on the previous discussion in this chapter—that the ICS ensures safety. The safety aspect is therefore discussed for the synchronised switching schemes only.

Event Distribution

Both synchronised channel-switch schemes do not make any assumptions on how an event is distributed, but offer a time slot to achieve it. Event messages have high priority and are therefore likely to be sent at the beginning of the safety time

slot. Hence, an effective event distribution can be achieved by the echo mechanism if synchronisation is ensured.

The global synchronisation scheme ensures a synchronised beginning of the safety time. As soon as the channel is idle for a certain time, the non-safety time slot starts. So a vehicle can get out of sync for a part of one safety time slot only. This is not assumed to be critical. Especially since the event messages are sent at the beginning of the time slot.

This is different with the distributed synchronisation scheme of the *i*-Channel that does not provide any fixed synchronisation points. However, it is beyond the scope of this work to analyse whether a proper synchronisation is possible or not.

Overall Safety Time

The synchronised switching schemes attempt to maximise the available time for the non-safety time slots permanently in order to maximise the throughput for non-safety communication whenever a service is available. Services are assumed to be unavailable for a substantial part of the driving time. Hence, a substantial part of the time that could be used for safety communication is lost for nothing.

One might argue that this should not be a concern as long as safety is ensured. However, this 'reliability' is always based on a probability and can never be assumed to be one hundred percent. It is therefore advisable to only provide non-safety time slots if a service is available.

7.8.2 Throughput

The non-safety throughput should be distinguished between the single station and the overall system—the later regards the amount of data the RSU can exchange.

Single Station

A station using a service requires spending as much time as possible on the service channel. This time is given by the ratio of the safety and non-safety intervals.

The non-safety time slots in the global synchronisation scheme can be fragmented or have zero length. On the other side, both the *i*-Channel and the ICS do not suffer from this fragmentation and additionally ensure a certain throughput.[6]

The ICS is based on the assumption that a certain amount of safety communication can be missed without jeopardising safety. In particular, the ICS provides a very high non-safety throughput if only a few events occur at once—in such a case an efficient event-check is assumed, resulting in short safety operation intervals. It is therefore assumed that the average throughput of a single station is higher with the ICS than it would be with the synchronised switching schemes. However, simulations are necessary to analyse the throughput of the different schemes appropriately.

Overall Throughput

The system bandwidth of the synchronised schemes is given by the ratio of the safety and the non-safety time slots. The synchronised scheme has therefore a limited bandwidth, but provides synchronised non-safety time slots, allowing communication that requires a predictable availability of the stations—such as routing of non-safety data. In addition, a RSU can provide safety and non-safety data in accordance with the current operation interval.

[6] It can be assumed that sooner or later all ongoing events will be received with the ICS.

The ICS does not synchronise the different stations and therefore offers to use the whole bandwidth of the Service Channel. However, routing of non-safety data is assumed to be rather difficult with the ICS due to the unpredictable availability of the stations.

7.8.3 Overview

The following lists provide an overview of the advantages and disadvantages of the different channel-switch schemes.

Independent Channel Switching

+ Most of the driving time is spent for safety communication,
  i.e. non-safety operation intervals are only provided if they are required
+ High throughput—even in a heavy traffic situation
+ The non-safety throughput of the system is maximised
+ Short time non-safety communication is always possible
− Non-safety communication is not synchronised

Global Synchronisation

+ Safety and non-safety operation intervals are synchronised
+ RSU can easily provide safety and non-safety information
− Limited overall non-safety throughput
− Most of the non-safety operation intervals are not used
− Non-safety communication cannot be guaranteed
− Non-safety operation intervals can be fragmented

***i*-Channel**

+ Safety and non-safety operation intervals are synchronised
+ Minimum non-safety throughput of about $0.25 \cdot B_{\mathrm{Max}}$
+ RSU can easily provide safety and non-safety information
− Out of sync possible and might jeopardise safety
− Most of the non-safety operation intervals are not used
− Limited overall non-safety throughput
− Limited non-safety throughput in a heavy traffic situation

The ICS ensures safety and provides a high non-safety throughput for a single station and the overall system. The global synchronisation is not suited for a high non-safety throughput and the i-Channel could get out of sync resulting in jeopardising safety. The ICS is therefore assumed to be the most suitable approach doing channel switching.

Chapter 8

Communication Stack

This chapter proposes a communication stack for DSRC providing the functionality as discussed in the previous chapters. The communication stack is depicted in Figure 8.1 and shows three distinctive stacks. Two of them, namely TCP/IP and SDT (single hop data transfer), are meant to be used by non-safety applications. The third stack, referred to as *Safety Stack*, is used for safety communication exclusively. In addition, there is a vertical layer, the so-called *Information Connector*, accessible by all layers allowing data exchange between them.

The following discussion addresses the different layers from the bottom to the top focussing on the data processing, i.e. process the packet in order to hand it over to the next layer, but describes management functions as well. However, unless otherwise stated the described tasks are meant for data processing.

In the discussion of the different layers, the following terminology for data units is used: A *safety information unit* is the core unit of the safety data, such as acceleration or heading of the vehicle. A *safety message* is either a routine message or an event message and comprises several safety information units. A *safety packet* comprises up to two safety messages. A *packet* can comprise a safety or a non-safety packet. And the term *frame* is used for a packet in the MAC and PHY.

8.1 Physical Layer

The physical layer shall be implemented in accordance with 'IEEE 802.11p' [9]. This layer deals with all issues involved in the transmission of raw bits on one of the channels approved to be used with DSRC.

The data is encoded using OFDM in combination with either BPSK, QPSK, 16-QAM or 64-QAM as modulation scheme. The first 128 bits of the message are modulated with BPSK and contain information about the modulation scheme of the rest of the message. After a successful decode of the PHY preamble, the device switches to a so-called *reception mode* to receive the rest of the message in the appropriate modulation scheme.

A mechanism referred to as "improved capture effect" is recommended to be used to increase the overall performance of the reception.[1] According to this mechanism, the radio should not switch to reception mode if the signal strength of a received header is weak. The signal strength of the remaining message, encoded with a higher order modulation scheme, is even lower and a successful reception of the message not likely. The following situation should be considered to get an idea about the

[1] This is not standardised in 'IEEE 802.11p' yet.

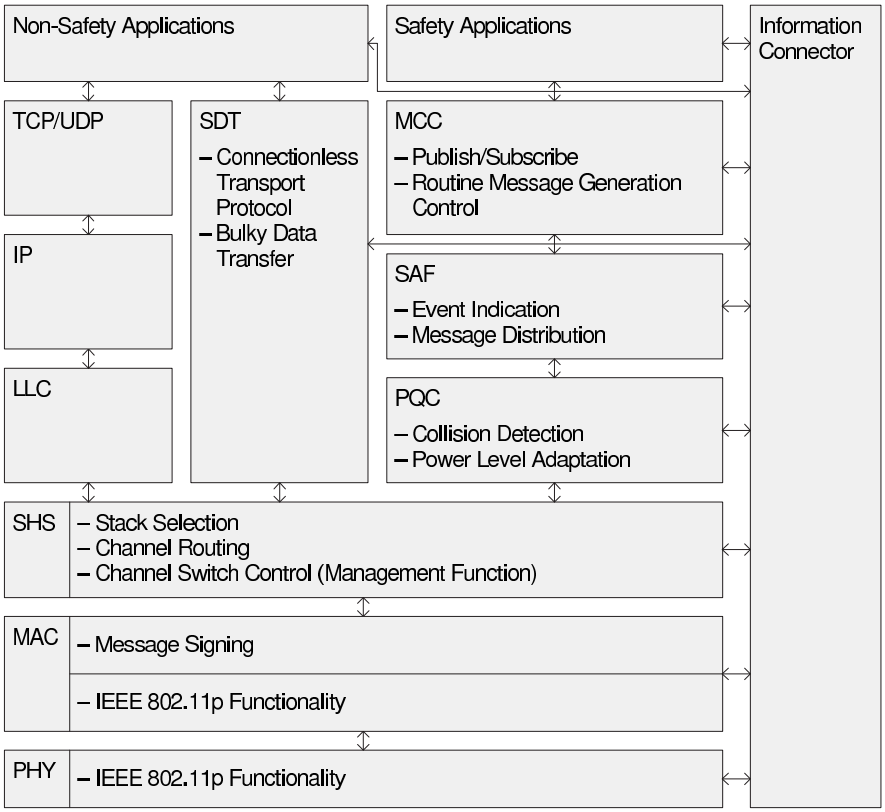

Figure 8.1: The proposed communication stack for DSRC provides three distinctive stacks. Two of them, TCP/IP and SDT, are meant to be used for non-safety applications. The third one, referred to as Safety Stack, is exclusively used for safety communication. All layers, but the TCP/IP stack, have access to the Information Connector.

advantage of this approach: A message with a weak signal is overlapped by a message with a strong one. The former one will be corrupted by the later one, but not necessarily vice versa. Therefore, if the radio does not switch to reception mode to receive the first message, the second one might be received successfully.

8.2 Medium Access Control Sublayer

The medium access control (MAC) sublayer coordinates the channel access in a distributed manner and deals with security concerns.

8.2.1 Enhanced Distributed Coordination Function

The channel access protocol of the MAC sublayer determines the station that has access to the channel in a distributed manner. It minimises the probability of a message collision and provides a fast access to the channel for high priority messages.

This is achieved by a mechanism referred to "enhanced distributed coordination function" (EDCF)—as to be standardised in 'IEEE 802.11e'. The EDCF is an extension of the DCF mechanism and works roughly in the following way. The frame's AIFS (arbitration inter-frame space) length and contention window size are adjusted according to the frame's priority level, of which a total of four are available. The higher the priority of a frame, the smaller its AIFS length and contention window size are chosen. Hence a high priority frame is more likely to get channel access over the lower prioritised ones. It should be noted that this scheme is not yet standardised in 'IEEE 802.11p'.

8.2.2 Security

According to 'VSC Task 11' [3] and 'IEEE 1556' [13] all safety messages and all channel-switch announcements are signed to check the frame's integrity and to authenticate its origin. An incoming frame on the other hand has to be validated before the message is handed to the upper layer. A comprehensive discussion about security is held in Appendix A.

Privacy is another important security issue. The MAC frame contains the sender's MAC address, providing the possibility to allocate a frame to a specific vehicle and has therefore an impact on privacy. One might argue that a vehicle identifier is not needed at all due to the broadcast nature of the safety communication and therefore should not be added to the MAC header of a safety packet. However, it should not be assumed that future safety applications do not need to address a specific station.

In order to grant the necessary privacy, the MAC address is suggested to be dynamic and random. The random MAC address provides anonymity as nobody can influence or predict a randomly chosen address. The demand on the MAC address to be dynamic, i.e. to change frequently, prevents the communication of a vehicle to be linked over a long period. It should be noted that RSUs have a static MAC address assigned.

8.3 Single Hop Switch Layer

The task of the single hop switch layer (SHS) is to switch packets to the appropriate channel on one side and to the appropriate stack on the other side.

8.3.1 Stack Selection

An incoming packet has to be routed to the appropriate stack. This is a common multiplex/demulitplex task and is achieved adding a stack identifier to the SHS header.

8.3.2 Channel Routing

The channel routing protocol ensures that an outgoing packet is sent with the corresponding transmission profile—i.e. on the appropriate channel with the proper priority, power and modulation. This task is dependent on the stack the packet is provided from:

- **Safety Stack**

  Safety messages provide the transmission profile on a per packet basis.

- **SDT Stack**

  A channel-switch announcement for a service using the SDT stack requires providing a transaction identifier and the transmission profile. An application requests to use the service by registering the transaction identifier and providing it with every sent packet.

  This scheme does not claim to be the smoothest solution. It is added here to show that the channel routing is feasible.

- **TCP/IP Stack**

  In 'IEEE P1609.4' it is defined how the channel routing should be done with DSRC for TCP/IP data transfer. According to 'IEEE P1609.4', the channel routing is based on the packet's destination MAC-address for unicast and based on a registered profile for multicast and broadcast. A detailed description of this process is given in 'IEEE P1609.4' [12].

Based on the transmission profile, in particular the priority and the channel information, the packet can be routed to the appropriate queue in order to be sent as soon as possible. The queue's implementation is not specified. However, it can be assumed that there is a queue—maybe just logically—for every combination of channel number and priority.

8.3.3 Channel Switch

The channel-switch control is a management function and can either be done in the MAC or in the SHS. A higher layer is out of question due to the different stacks.

In contrast with other wireless standards, such as 'IEEE 802.11 a/b/g', DSRC provides concurrent communications in independent channels. However, the radio can only transfer data on one channel at the time, implying that channel switching is required to provide 'concurrent' communication.

The channel switching should not be done on a per packet basis in the MAC, but in SHS having the overview of the packet flow, in particular on the incoming traffic. In addition, it is an advantage to keep the MAC as simple as possible in order to implement the MAC and the PHY on a single chip.

8.4 Packet Quality and Control Layer

The packet quality and control layer (PQC) provides two functions: first, it provides a collision and quality detection, and second, it translates the intended *broadcast range* of the packet to the appropriate power level.

8.4.1 Collision and Quality Detection

Based on a piggyback acknowledgement scheme, the PQC estimates the average quality of the channel and detects hidden terminal collisions with a certain probability. It should be pointed out that the PQC does not differ between different message types, i.e. event or routine messages. The PQC provides two services:

- **Acknowledged Packet Transfer**

  The acknowledged packet transfer requests a piggyback acknowledgment from stations within a certain range and time window. If the received packets do not allow the ruling out of the possibility of a collision, the packet will be sent again using the unacknowledged packet transfer.

  In addition to the collision detection, the received acknowledgments provide the possibility to estimate the current channel quality analysing the overall feedback of the last few messages.

- **Unacknowledged Packet Transfer**

  The unacknowledged packet transfer does not request an acknowledgment. It should be noted that requested acknowledgements are piggybacked all the same.

The acknowledged transfer is assumed to be used by routine messages while the event messages are transferred using the unacknowledged service. This is not a contradiction per se since event messages have a much faster and more reliable acknowledgement scheme in the safety layer.

8.4.2 Power Level Adaptation

Safety messages have to be distributed within a certain range. The reception range of a message depends a lot on the rapidly changing channel quality, requiring to adapt the transmission power accordingly. This rather challenging task is suggested to be done in the PQC in order to allow the upper layers in the Safety Stack to deal with the abstract parameter of the broadcast range. It is therefore the task of the power level adaption to translate the requested broadcast range to the according transmission power. It should be noted that this translation is provided on a best-offer basis only—this is due to the non-deterministic channel characteristic.

This task requires to estimate the current channel quality. In order to do that, it is assumed that the PQC adds the intended transmission power to the safety packet header. This should allow the estimation of the current signal attenuation by comparing this value with the received power level. However, it is beyond the scope of this work to address the question of the channel quality estimation.

8.5 Safety Layer (SAF)

The safety layer (SAF) resides on top of the PQC and provides the functionality to distribute safety messages. In particular, the SAF provides an effective event distribution.

8.5.1 Outgoing Safety Message

An outgoing safety message is processed as follows:

1. **Event Identifier and Sequence Number**

   An outgoing event message is provided with an *internal event identifier*. If this internal event identifier is unknown to the SAF, a new event identifier has to be assigned to it and the sequence number has to be initialised to zero. A known event on the other hand is sent with the previously assigned event identifier, but with an incremented sequence number.

2. **Message Echoing**

   An outgoing routine message, referred to as *primary message*, is enhanced with a prior received safety message. This is an event message whenever a suited one is available or a routine message otherwise. However, no message will be echoed if no suitable message is available—e.g. no surrounding traffic.

   The information of the position, time and priority of the primary message is attached in the security frame of the MAC and the SHS respectively. This information will not be available for the echoed message and needs to be attached to the SAF header. As discussed previously, the relative position and time is attached to decrease the overall packet size.

3. **Event Message Indications**

   All events that are eligible to be echoed, except the one that is echoed, are indicated in the SAF header with their event identifiers and sequence numbers. In addition, the Control Channel time is required to be attached to the header as well.

8.5.2 Incoming Packet

An incoming safety packet is processed as follows:

1. **Missing Event Detection**

   This is a management function. The provided event indication is compared with a list of the received events in order to detect missing ones.

   As discussed previously, a safe switch to the Service Channel is only possible if all ongoing events have been received. It is therefore required to indicate to the channel-switch module whether it is safe to switch or not from an “missing event detection” point of view.

2. **Safety Packet Split**

   An incoming safety packet usually contains two safety messages—one from the sending station, the primary message, and an echoed one. As discussed previously, the lower layers provide the position, time and priority of the primary message. The echoed message on the other hand has this information provided in the SAF header.

   Subsequently the two messages are treated as independent safety messages according to their priority.

3. **Message Distribution**

   The routine message that is suited the most for being echoed and all event messages eligible for echoing are held in a buffer and processed according to the event distribution protocol.

4. **Handover to Upper Layer**

   Routine messages are always forwarded to the upper layer. Event messages on the other hand are forwarded only if the event has not been received before.

8.5.3 Event Distribution Protocol

This module has the objective to distribute all current events in a distributed manner as discussed in Section 7.5.

- **Ensure Reliable Distribution of One's Own Events**

  It is of paramount importance to keep track of the distribution of one's own event by checking if the event is echoed or indicated by vehicles at the proper position.

  If this is not accomplished, the event has to be sent again. Instead of rebroadcast the old event message, an updated version of the message is requested using the internal event identifier.

- **Echo in Routine Message**

  Event messages are echoed in the regular scheduled routine messages. If several events are suited to be echoed, the appropriate one is selected in accordance with the goal of making the collectivity of the event distributions reliable.

  It should be noted that whenever no event is eligible to be echoed, the most suited routine message is echoed instead.

- **Forward Event Message**

  The echoing mechanism is not assumed to be effective in a sparse traffic situation. However, such a situation implies that the channel is not under a heavy load and allows the generation of additional messages to improve the effectiveness of the event distribution.

  The forwarding of the event message is based on a countdown mechanism. A received event is forwarded, i.e. piggybacked to a requested routine message, if the timer expires while not receiving the event again. The value of the timer depends highly on the current environment.

- **Set Broadcast Range**

  Set the broadcast range of an event message in accordance with the distribution range of the event, the surrounding, and the channel load.

  It should be noted that the broadcast range of routine messages is determined by the 'routine message generation control unit' in the layer above.

The presented event distribution protocol implies that it is necessary to request the generation of event and routine messages respectively.

8.6 Safety Message Creation and Control Layer (MCC)

The safety message creation and control layer (MCC) deals with the creation and the reception of the actual safety data.

8.6.1 Routine Message Generation Control

The routine message generation control adjusts the communication parameters of the routine messages, namely interval and broadcast range, in order to meet the following demands:

- The Control Channel must be available at any time to handle an additional peak load due to events, requiring to strictly regulate the channel load caused by routine messages.

  As discussed previously, the routine messages absorb the major load of the event distribution. It is therefore assumed that the routine messages' channel load can be a substantial part of the channel's capacity.

- All stations should send their routine messages with a similar set of parameters in order to ensure fair access to the channel. This requires most likely to attach the current parameter set to the header.

It should be noted that the routine message generation control is assigned to the same layer as the publishing module. This ensures that the information is available if an unscheduled outgoing message—requested by a safety application or the SHS—contains the routine information and allows adapting the schedule for the next routine message accordingly.

However, the MCC is separated form the physical medium by several layers and has therefore no information about the current load on the Control Channel. This information will be provided by the Information Connector as discussed in Section 8.9.

8.6.2 Subscription

The subscribing module parses the incoming safety message and provides the safety information units to the applications according to their subscription.

8.6.3 Publish

The publish module provides three different services:

- **Publish Routine Information**

  A basic set of routine information—such as speed, acceleration and heading—should always be provided in the routine messages.[2] In addition to this "basic information set", a safety application can request to publish additional data in the routine messages.

- **Publish Event Information**

  An event triggering application can request to publish event related information within a certain range and time window. The priority level of the event is set according to the event's importance.

  The routine information is assumed to be always included in the event messages. In the rare case of concurrent events, the publish module combines them into one single event message.

- **Publish Application Based Message**

  A safety application can request to send an application-based message that is distributed the same way as a regular event, but must not be combined with other events. This service is particularly useful for applications such as the approaching emergency vehicle warning.[3]

[2] The position and time are attached to the signature and do not belong to the routine message's basic information set.

[3] An emergency vehicle requires announcing its approaching. However, the information that an ambulance brakes should not be distributed far ahead.

As discussed previously, the publishing module assigns an internal event identifier to the event and application based messages in order to allow the underlying safety layer to request an updated version of a specific message.

8.6.4 Vehicle Identifier

It is required that a safety message can be assigned to a specific vehicle in order to keep the context up to date. Such an identifier is provided in the MAC header. However, the vehicle's MAC address cannot be assumed to be fully controllable by the Safety Stack due to the concurrent non-safety communication. It is therefore suggested to provide a separate vehicle identifier for the purpose of safety only. It is suggested that the vehicle's identifier is a three byte long, random and dynamic number.[4] Randomness guarantees anonymity while the frequent change of the identifier prevents the vehicle of being traceable over a long period.

It should be noted that unlinkability is not necessarily a contradiction to the necessary of linking consecutive safety message in order to keep track of the surrounding traffic. The vehicle's identifier is not meant to be changed for every single message; it rather changes on the order of once every few minutes. So whenever the identifier changes, the corresponding vehicle will virtually vanish and reemerge as a new vehicle in the context. However, this happens only once in a while and has therefore only a minor impact on the context's accuracy.

8.6.5 Context Management

The context management combines all the received status information of the surrounding traffic to maintain the context. This is a management functionality and is not directly involved in the data processing.

The context management has to ensure that the context is always up to date even though non-safety communication is done a substantial part of the time. This is achieved by indicating to the channel-switch module whether it is safe to do non-safety communication or not from a "context up-to-date" point of view.

8.7 Single Hop Data Transfer Layer (SDT)

The single hop data transfer (SDT) layer provides two services: the bulky data transfer and a connectionless data transport. In order to understand the functionality of the former, a brief introduction on LT-codes is held first.

8.7.1 LT-code

Information-providing applications, such as map update, require the broadcasting of a large amount of data to several interested vehicles. Instead of sending the same data to each vehicle individually, the data should be distributed to all interested stations at once in order to increase the overall throughput.

The amount of provided data can be large, but the packet size for a DSRC message is limited to about four kilobytes. It is therefore likely that the whole data requires to be split up into several packets—e.g. a one-megabyte digital map is split up into about 250 packets.

[4] A safety message usually contains an echo. Hence, the overall size of the packet is the same if the MAC address or a second but smaller identifier is used.

These packets are sent to a basically unlimited number of vehicles. This implies that a feedback scheme to acknowledge packets does not scale. In particular, a station cannot request a missed packet to be sent again.

The different packets can be sent repeatedly in a loop. This permits to wait for a missed packet to be sent again. However, if that specific packet is missed again the station has to wait for another round to get the next chance to receive the missing packet. So the station might be listening to the Service Channel for a long time before the whole data chunk is received. This results in two major problems: First, the radio needs to listen to the Service Channel for a long time not receiving any useful data most of the time. This time would rather be spent doing safety communication on the Control Channel. And second, the vehicle's time to finish the transaction is very limited as the vehicle is in the service-providing RSU's reception range for a short period only. The approach of sending the packets in a loop is therefore a reasonable approach only if a small amount of data is to be distributed.

A more suitable approach for a bulky data transfer is to use a forward error correction technology such as LT-code [29, 30, 31]. The fundamental idea of the LT-code is the following: In a first step, the data is split up into so-called "original blocks" of equal size. In a second step, "new blocks" are created based on a linear combination of subsets of the original ones. It should be noted that each possible subset of original blocks creates a different new block. This implies that a basically endless stream of unequal new blocks can be broadcasted, all of them containing partial information of the original data. A station interested in the data has to collect an arbitrary selection of the encoded blocks. As soon as enough blocks are received—there is usually an overhead of about $5-10$ percent—the original data can be restored.

The small overhead, the fact that these stations can gather the packets whenever they have time to do so, and the fact that no feedback is required to distribute that data to an unlimited number of interested stations, makes this code ideal to be used in combination with a unreliable wireless environment in which the stations have a sporadic, unpredictable, and not synchronised access to the channel.

8.7.2 Bulky Data Transfer Protocol (BDTP)

The Bulky Data Transfer Protocol (BDTP) provides two functions: On the one hand it encodes and publishes packets and on the other it gathers them to decode the original data.

Data Distribution

The data is encoded in a stream of packets using a well-defined LT-code. This stream can be endless; hence, the BDTP requires to be instructed how this stream shall be limited. This limit can be based on the packet rate, the time, the number of packets in terms of overhead, or a combination of them. In addition, it should be possible to stop or pause a current transfer.

Besides the encoded packet stream, a packet describing the encoded content is required to be sent once in a while.[5]

[5]The channel-switch announcement might contain information about the provided data. However, this information is likely to be insufficient due to the limited size of the channel-switch announcement.

Data Aggregation

A received packet describing the encoded data is passed to the application that requested the service. This application responds to the BDTP indicating whether the message should be gathered or not. If the data is requested, the BDTP gathers as many packets as necessary to decode the data and forward it to the appropriate application. The BDTP does not wait for the application's request—i.e. the reception of a data-describing packet—to start collecting the encoded packets. It starts collecting the packets right away and discards them if the application is not interested in the data. It should be noted that the bulky data aggregation scheme is depicted in Figure 8.2.

8.7.3 Connectionless Transport Protocol

The SDT provides a connectionless transport protocol on a single-hop basis. It is similar to UDP but does not provide a checksum. In particular, there is no reliability function provided.

This protocol is meant to be used for applications doing short time interactions, such as the electronic toll collection. These applications are assumed to have a fairly simple message exchange and do not require a sophisticated transport protocol with its reliability overhead. It should be noted that this assumption is in accordance with 'IEEE P1609.1', suggesting using UDP for short time interactions.

8.8 TCP/IP Stack

The TCP/IP stack provides the well-known Internet protocol suite.

8.8.1 Logical Link Control (LLC)

DSRC provides a best-effort datagram delivery service only and neither error control nor flow control. Hence, the LLC shall provide a connectionless unacknowledged message delivery, referred to as "Type 1", as standardised in 'IEEE 802.2' [5].

8.8.2 IPv6 versus IPv4

The Internet Protocol version 6 (IPv6), as defined in 'RFC 2460' [35], shall be supported. IPv6 is preferred to IPv4, as defined in 'RFC 791' [36], for three reasons:

1. In order to provide privacy, the vehicles must not have a static IP address assigned to them. Hence the vehicles have to choose a random IP address that has a very high probability to be unique within a certain range. IPv4 has a very limited address space and is therefore not suitable.

2. It is very tricky, if not impossible, to route packets based on completely random IP addresses. IPv6 allows choosing a random address based on the current position to tackle that problem.

3. IPv6 is supposed to replace IPv4 in the next twenty years. So it would be absurd to design a new communication architecture that is based on a protocol that is to be replaced in the near future.

Nevertheless, IPv4 should be supported to be compliant with 'IEEE 802.11a' access points, but is not meant to be used with DSRC.

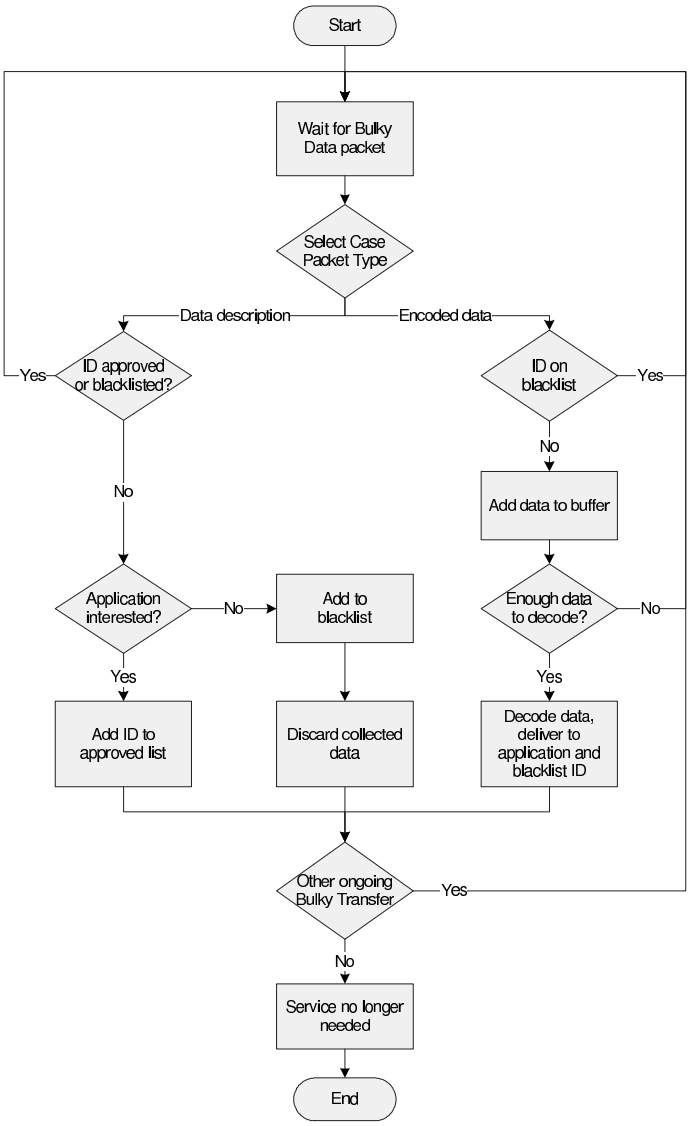

Figure 8.2: Data aggregation scheme for the bulky data transfer.

8.8.3 User Datagram Protocol (UDP)

The user datagram protocol (UDP)—a message oriented connectionless transport protocol—shall be supported as defined in 'RFC 768' [37]. It provides multiplexing and a data checksum on top of an IP datagram.

8.8.4 Transmission Control Protocol (TCP)

The transmission control protocol (TCP)—a connection oriented, reliable byte stream transport communication protocol—shall be supported as defined in 'RFC 793' [38].

8.9 Information Connector

The Information Connector is a vertical crossbar that all layers have access to. It provides a data pool to exchange protocol information that is of interest in other layers. In particular, the Information Connector provides a common interface to the vehicle's sensor network. It should be noted that the idea of this vertical layer originates from the staircase approach proposed by Füssler et al. [16].

The Information Connector is neither meant to interact with a protocol nor to allow such interactions between protocols. Published data in the Information Connector shall therefore not trigger direct actions, but protocols looking up the data can react accordingly.

Such an information crossbar does not fit into the traditional layered architecture, in which all layers work independently from each other. However, such a data exchange is necessary as the following list of provided data indicates:

- **Vehicles Sensor Data**

  The Information Connector provides a common interface to access the vehicle's sensor data. This data is usually provided by the vehicle's sensor network, e.g. CAN bus, and should be buffered to be accessible at all times. It should be highlighted that some sensor data, such as the GPS position, is not updated frequently and is therefore meant to be extrapolated in order to be more accurate.

- **Position and Time**

  The position and the time are provided by the security module in the MAC.

- **Channel Load**

  The average and the current load of the different channels are published by the PHY or the MAC. This information is of particular interest for the routine message generation control, the event message distribution and the bulky data transfer.

- **Current Channel Quality**

  The current channel quality, estimated in the PQC, is of particular interest for the event distribution protocol.

- **Current Context**

  The context is available in the MCC and is of major interest for several protocols in the Safety Stack.

- **Control Channel Time**

  The event indication protocol requires including the Control Channel's time span.

The published data in the Information Connector depends on the actual design of the protocols. The provided list therefore neither claims to be complete nor to be specific about the presentation of the data.

8.10 Management Functionality

The previous discussion of the communication stack was presented from a data processing point of view and the management functionality has not been addressed as a whole. Especially the channel-switch functions are spread among several layers. The following list should therefore provide an overview of the functionality involved in the channel switching.

- **Event Detection**

  The missing event detection—assigned to the SAF—indicates to the channel-switch entity whether it is safe to switch the channel from a "missing event detection" point of view.

- **Context Update**

  The context management—assigned to the MCC—indicates to the channel-switch entity whether it is safe to switch the channel from a "context up-to-date" point of view.

- **Application Interrupt**

  Safety applications indicate whether it is safe to switch the channel from their point of view.

- **Channel Switch Announcement**

  A service providing application can request to send a channel-switch announcement on the Control Channel.

  A received channel-switch announcement is indicated to the applications. An application interested in the service can request to use it.

- **Channel Switch**

  A channel switch from the Control Channel to a Service Channel is performed whenever it is safe to do so and a requested service is announced.

  A channel switch to the Control Channel is performed if:

  - Safety requires it
  - A service should be announced
  - The service is no longer required
  - The service is not available

- **Service Suspend Indication**

  The Safety Stack and the SDT should be indicated about a channel switch in order to react appropriately. For instance, the Bulky Data Transfer Protocol should not generate new packets while the station is not transferring data on the appropriate channel.

This list provides an overview of the necessary functionality only and does not claim to be very detailed about it.

It should be noted that various other management functionality is necessary that does not belong to the actual data processing—such as the initialisation of the protocol stack or the management of the different layers. However, a thorough analysis of the management related issues requires a detailed definition of the protocols and is beyond the scope of this work.

Chapter 9

Discussion

This chapter discusses the presented communication architecture. A first part shows that the proposed communication stack provides the necessary functionality. The second part points out that the communication stack allows an effective event distribution in a timely manner. The last part discusses the communication architectures presented in Section 3.2.

9.1 Communication Stack

The proposed architecture provides three distinctive stacks—namely TCP/IP, SDT, and the Safety Stack. This section points out that the TCP/IP stack itself cannot provide the necessary safety communication, shows that two stacks are required to do non-safety communication and closes by highlighting the Safety Stack.

9.1.1 TCP/IP

The TCP/IP protocol stack has been proven to be suitable in combination with many different communication technologies, in particular with wireless ad-hoc networks such as a wireless home networks. So why not use TCP/IP with DSRC?

The main reason for TCP/IP's success is its capability to do internetworking, i.e. packets can be routed between different networks based on radically different technologies. This is achieved masquerading TCP/IP from the underlying network. This implies that these protocols neither can nor do have to deal with the characteristics of the underlying network.

A vehicular ad-hoc network that provides safety has unique demands on the communication, such as distribute a message to all nodes within a certain range in a reliable manner. In particular, there is no requirement to route a safety message within the network to a specific node. However, an effective distribution cannot be achieved if the underlying communication technology is masqueraded. For instance, the channel must not congest for any reasons. TCP tackles congestions by reducing the packet rate as soon as the protocol senses a congested network. Hence, the network congests first before the appropriate action is taken.

It is therefore necessary that the transport layer—this would be the MCC in the presented communication stack—knows about the characteristic of the network technology and receives information from the underlying layers. This implies that TCP/IP is not suited to do safety communication.

9.1.2 TCP/IP Versus SDT

Non-safety communication does not require to understand the underlying network technology. This implies that TCP/IP can be used to exchange non-safety data. Nevertheless, it is suggested to primarily use the SDT to do non-safety communication. The reasoning for this suggestion is discussed subsequently.

Most of the non-safety applications are based on a direct link between the RSU and the vehicle—i.e. the vehicle exchanges data with the RSU while driving in reception range. As discussed previously, it is not meant that packets are routed along the road to the RSU.

The majority of non-safety applications perform 'short time interactions' or 'information fuelling'. Short time interactions are assumed to have a fairly simple message exchange and information fuelling should use the Bulky Data Transfer Protocol. Hence, the Internet protocol suite is not required for these two types of applications, but provides an unnecessary complexity (routing and reliability) to the communication and the assigned IP address raises privacy concerns.

The only but important application that benefits from the TCP/IP stack is the Internet access, for which the RSU provides a gateway to route the packets to its destination. However, not all applications are known yet and future applications may require a reliable multi-hop transport protocol.

9.1.3 Safety Stack

The Safety Stack is exclusively used by safety applications and provides the functionality to do safety communication. Its structure is discussed in the remainder of this section.

Safety communication contains a great deal of tasks. Combining them into a single layer would therefore result in one complex composition. This is avoided spreading out the tasks to several layers, each of them processing distinctive data units. A safety packet contains three elemental data units:

- **Safety Information Units**

  The safety applications send and receive distinctive safety information units.

- **Safety Message**

  A safety message contains the collectivity of safety information units. Two message types can be distinguished: routine messages and event messages.

- **Safety Packet**

  Up to two safety messages are combined into one safety packet.

The Safety Stack is divided into three layers according to these data units. The assignment of the services to these layers bases upon two fundamental requirements: First, all necessary information to process the data is available in the layer. Second, the information flow between the layers is consistent. The proposed design satisfy these criteria as illustrated in Figure 9.1 and described in the following list:

- **Publish/Subscribe**

  The publish/subscribe mechanism deals with safety information units.

- **Routine Message Generation Control**

  This service is required to know if the outgoing message comprises the routine information. This knowledge is only available in the MCC—the layer that creates the message. In addition, the routine message creation control and the

message generation should belong to the same layer in order to avoid synchronisation problems.

- **Event Indication**

  The reception of an event message must be indicated. This task is based on the reception of a safety message and should be assigned to the SAF—the layer that deals with safety messages.

- **Message Distribution**

  The message distribution ensures an effective event distribution. This task deals with safety messages and belongs to the SAF.

- **Collision Detection**

  The collision detection is achieved piggybacking received safety packets and does therefore belong to the PQC.

- **Power Level Adaptation**

  The MCC and the SAF should not deal with transmission specific parameters, such as transmission power, but rather deals with the abstract parameter of the broadcast range instead. This requires a power level adaptation in the PQC.

It should be noted, that the three safety layers have very similar tasks than the ones in the TCP/IP stack. The LLC and the PQC layer do both masquerade the underlying technology and provide different operational modes—in particular, an acknowledged and unacknowledged connectionless mode. The IP and the SAF layer do both ensure that the messages are routed, i.e. distributed, appropriately. And the TCP and the MCC layer do both have congestion-control functionality. However, there are differences as well. For instance, the reliability is mainly achieved on the network (SAF) and not the transport level (TCP).

9.2 Event Message Distribution

The distribution of time-critical events is the crux issue of safety communication. This is discussed below analysing the reliability and the possible delays.

9.2.1 Reliability

It is utterly impossible to guarantee a certain quality of service for a broadcast based communication over a non-deterministic channel. Nevertheless, a very reliable message distribution is assumed by the proposed distribution scheme in combination with the congestion control.

The distribution scheme is based on two cooperating mechanisms: the event echoing and the event indicating. The echoing results in an observable flooding effect and does already ensure a very high reliability. The event indication scheme provides the information that a particular station did not receive an event. This triggers a rebroadcast if the event is not echoed regularly.

The congestion control ensures that the channel is never under too much load by restraining the number of packet collisions. This is achieved by the routine message generation control (limiting the average channel load) and the echoing (levelling the peak load imposed by events).

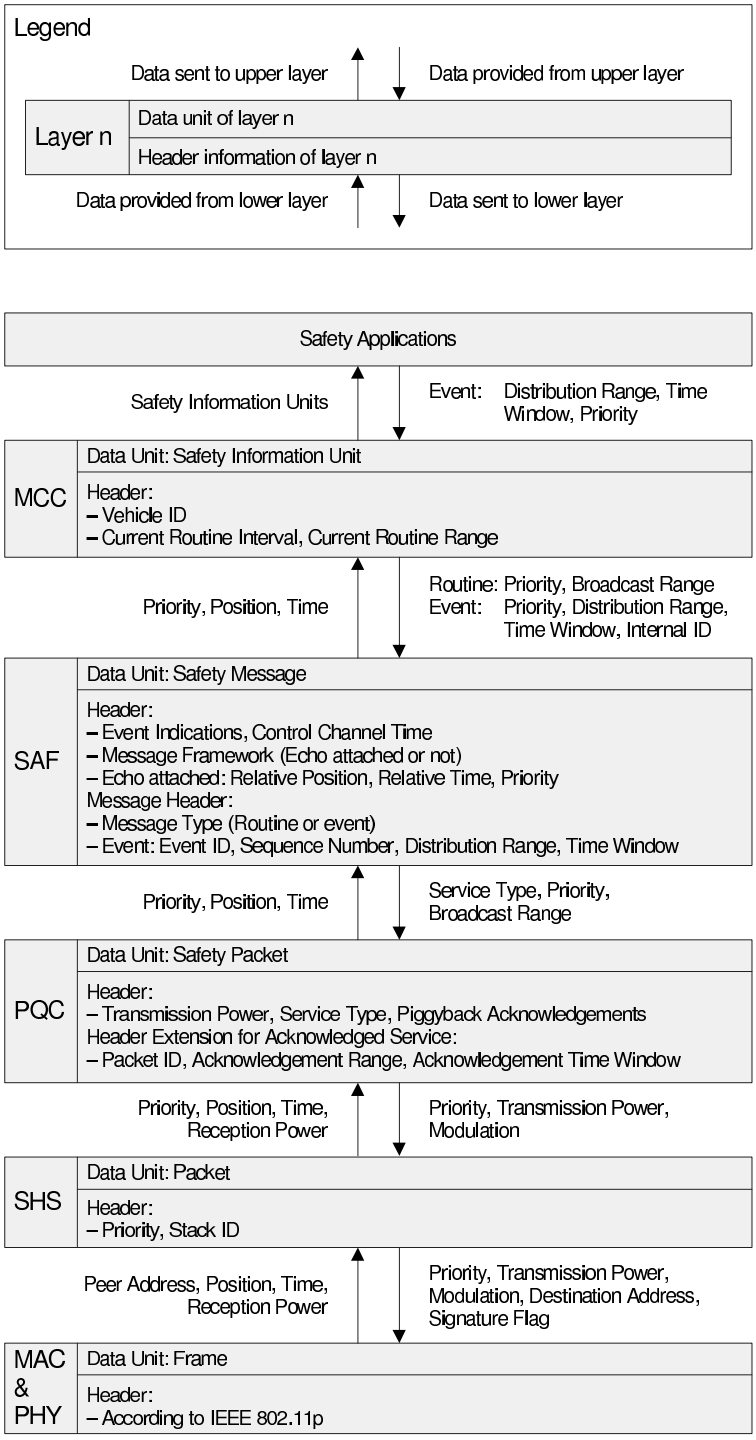

Figure 9.1: Information flow of the communication stack for safety applications. It presents the header information of the different layers and shows the information flow between them. Some of the information is message type specific and is denoted with 'Event:' and 'Routine:' respectively.

9.2.2 Delay

Three possible delays are involved in the event distribution: processing delay, distribution delay and channel-switch delay.

Processing Delay

The procession delay refers to the time that is required to process the message in the communication stack.

First of all, the publish module collects the data to create the safety message. This can be done very efficiently as the Information Connector provides all the vehicle's sensor information. In particular, there is no delay waiting for data on the vehicle's sensor network.

No extensive processing of the event message has to be done in the SAF and the PQC—events do not echo messages and use an unacknowledged packet transfer. In addition, the event message has a high priority assigned to it and is processed accordingly. It is therefore reasonable to assume only a minor delay in these two layers.

The SHS is a potential bottleneck since all the packets from the different stacks are queued to be routed to the according channel. The queue is priority based and does not affect a high priority message, but the channel routing adds a delay if a channel switch is required. Therefore, it should be considered whether to announce the creation of an event message to trigger the channel switch prior to the arrival of the packet in the queue.

The MAC layer has to sign the safety message. This task demands a certain amount of processing power and time. It is therefore recommended to sign the message with hardware support to minimise this delay. The channel access time on the other side is assumed to be very short: First, the routine message generation control ensures that the channel is not under too much load. Second, the priority based channel access scheme decreases the average waiting time of high priority packets.

The reception unit should process the incoming message according to its priority. There is no potential bottleneck and the processing—except the signature verification—is fairly simple. It is therefore reasonable to assume a minor reception delay only. It should be noted that the processing delay of the outgoing message is a threshold for all receiving stations. This implies that an outgoing packet, having the same priority as an incoming one, should always be processed first.

Distribution Delay

The distribution delay occurs if the original message is missed due to the non-deterministic channel or a packet collision. The station has therefore to wait for the event of being either echoed or rebroadcast.

The echoing results in a flooding effect, given a certain routine message density. A low message density on the other side implies a moderately loaded channel only and allows broadcasting additional packets to distribute the event. In either case, the probability that all vehicles receive the event in a timely manner is assumed to be overwhelming.

Channel Switch Delay

Switching channels in order to do safety communication implies a certain delay threshold. The proposed channel-switch scheme adapts the non-safety operational interval—and therefore the maximal delay threshold—based on the current environment. Hence, a timely event delivery is ensured if this task is done appropriately.

9.3 Comparison With Existing Architectures

9.3.1 Layered Approach

The different layers are supposed to work independently. In particular, it is meant that the different protocols do not exchange information with each other. However, the analysis of the communication requirements points out that some of the functions cannot be assigned to a single layer. Some sort of data exchange between layers is therefore necessary making a strictly layered approach for a vehicular safety architecture not feasible.

9.3.2 IEEE P1609.3

A communication architecture, to be used with 5.9GHz DSRC, is about to be standardised by the IEEE working group P1609 and 1556. The proposed design, available in a draft version, focuses on non-safety applications and considers safety as a black box.

The IEEE P1609.3 communication stack, depicted in Figure 3.1(b), is similar to the one proposed in this work. Except that the whole Safety Stack and the Information Connector are missing. Hence, according to the results of this work, the architecture is not suited to do safety communication at all.

Apart from safety, the proposed architecture seems to be reasonable. However, it is highly recommended to provide the Bulky Data Transfer Protocol in the WAVE short message layer. In addition, it should be considered to shift the lower layers according to the structure proposed in this work for two reasons: The WAVE short message layer does not require the functionality provided by the LLC and the channel switching should rather be done in a layer above the MAC.

9.3.3 Staircase Approach

The staircase approach is meant to provide thoughts of a vehicular ad-hoc network and very important functions, such as event distribution and channel switching, are either not addressed in detail or not considered at all. Two of the staircase approach's key features are discussed in detail:

- **Information Connector**

  The idea of a vertical information crossbar was adopted in this work, but its functionality differs in one important point. The staircase approach suggests providing the data in a subscriber pattern. Hence, the subscriber gets an indication if the data is changed. This work proposes that the Information Connector provides look up functionality only.

- **Staircase Approach**

  The staircase approach proposes that the applications can bypass a service offered by a layer. For instance, the application can directly access the single-hop layer to send a message. The architecture proposed in this work provides several access points also, but differs in its realisation: the layers are not bypassed, but several stacks are provided.

There are no major disagreements between the thoughts of the staircase approach and the architecture proposed in this work. However, a detailed analysis is not possible due to the lack of details in the staircase approach.

Chapter 10

Conclusion

The objective of this Master's Thesis was to design a vehicular safety communication architecture to be used with 5.9GHz DSRC.

Safety communication based on DSRC has to deal with two principal problems. First, a deterministic channel model cannot be assumed in the fast changing topology of a vehicular ad-hoc network with its rapidly moving nodes located quite far apart. And second, hidden terminal collisions cannot be avoided efficiently due to the unbounded system and the broadcast nature of the safety communication.

The occurrence of most traffic accidents is based on the simple fact that two or more vehicles are at the same place at the same time. Usually this happens based on an event, such as a hard breaking vehicle, but collisions occur as well if all vehicles are driving in a normal and predictable manner. It is therefore necessary to send so-called routine messages on a regular basis whose information allows the prediction of the vehicle's position for the next few seconds. Whenever this prediction is jeopardised—the vehicle might be accelerating or changing its trajectory vigorously—a so-called event message has to be distributed in a reliable and timely manner based on an unreliable channel. It is through this distribution of messages that will hopefully reduce the likelihood of accidents. It should be noted that events are generally long lasting since human beings do not operate the vehicle in a very fast manner.

The event distribution is based on the so-called echo mechanism and works as follows. All vehicles that receive an event message shall attach the event information to their next scheduled routine message. This echoing does not increase the packet size substantially due to the enormous frame overhead the packets have compared to the safety information itself. Therefore, the event distribution does not add a lot of load to the channel but is very effective due to the flooding effect of the echoing.

Routine messages on the other side need to be received only once in while to keep the so-called context up to date. Missing a routine message due to the non-deterministic channel does not therefore jeopardise safety. However, packet collisions—usually caused by the hidden terminal effect—do decrease the overall performance of the routine message greatly and should be detected. This is achieved based on a so-called piggyback acknowledgement scheme.

All vehicles send routine messages on a regular basis. This adds a lot of load to the channel and can result in a channel breakdown. This must not happen for any reason. Therefore, a congestion control is necessary that manages the routine message generation in a distributed manner.

Besides the safety communication, non-safety data transfer should be offered as well. The safety communication is done in a distinctive channel, referred to as the Control Channel, and must not be used for non-safety data exchange. It is therefore

required to switch to another channel for non-safety communication resulting in temporarily not receiving any safety related messages, in particular event messages. In order to ensure safety, the Control Channel has to be checked frequently for ongoing events.

Prior to continuing the non-safety communication, the station has to make sure that all current events have been received. This is achieved by an event indication scheme that requests all vehicles to report all ongoing events in their messages. In addition to the event detection, it is necessary that channel-switching vehicles ensure that their context is updated once in a while and routine messages are still sent on a regular basis.

The communication stack shows two distinctive stacks for the non-safety communication and a third one exclusively to be used for safety communication. The latter comprises the functionality to provide an effective safety communication. In addition to these three stacks, a vertical information crossbar provides the functionality to deal with cross layer issues.

10.1 Contributions

This Master's Thesis provided the following contributions:

- **Messaging Scheme**

  Proposing a messaging scheme, based on a comprehensive analysis of how safety can be improved in combination with DSRC, featuring event and routine messages.

- **Message Distribution**

  Propose the echo mechanism in order to ensure an effective event distribution.

- **Channel-Switch Scheme**

  Introducing a channel-switch scheme that provides safety while maximising the non-safety throughput not relying on a deterministic channel characteristic.

- **Communication Stack**

  Design of a communication stack that provides safety communication while enabling the possibility of non-safety data transfer.

Among these major contributions various details and ideas are provided in this work.

10.2 Further Work

This work presented the overall structure for a vehicular safety communication architecture. A lot of functions had been emerged to be necessary and protocol ideas for them were suggested in this work. In a next step it would be necessary to analyse and simulate the behaviour of these protocols in order to prove their feasibility and to tweak their parameter settings. The following protocols are most important to be analysed in detail:

- **Event Distribution Protocol**

  Events are likely to occur in bursts. This makes the distribution of the collectivity of the events a challenging task and needs to be addressed.

The echo mechanism requires a certain traffic density to show its full potential. It needs to be analysed what the appropriate scheme for additional message creation is.

- **Channel Switch Protocol**

  The proposed channel-switch scheme needs to be simulated to maximise the throughput while safety is ensured.

- **Routine Message Generation Control**

  The channel must not congest for any reasons. A protocol for the routine message generation needs to be developed that grants appropriate channel access for all stations in a distributed manner.

- **Power Adaptation**

  The indented communication range of a safety message is translated to the according transmission power. This requires a good estimation of the current channel quality.

- **Context Enhancement**

  A routine message should be echoed if no event is distributed. A selection scheme for the appropriate routine message has to be found.

- **Collision Detection**

  This task is currently analysed at the DC RTNA and is of major interest.

The requirements of the various safety applications have not been thoroughly analysed yet. Actually most of them exist only as a very first draft. While analysing the requirements of the safety applications a lot of assumptions were made, trying not to exclude any possible demands the applications might have. It is of paramount importance to analyse the requirements of the various safety applications and to review the assumptions made in this work.

Appendix A

Security

Some of the text in this appendix originates from 'IEEE 1556 / D1' [13].

A.1 Security Frame Format

The V2V safety communication is assumed to be the major part of the overall communication on the Control Channel. It is important that the size of these messages is small; hence the security frame should not contain any unnecessary fields. Instead of having a very flexible security frame, as proposed in 'IEEE 1556' [13], it is recommended to provide a well-adapted frame for V2V safety messages as depicted in Table A.1:

- **Protocol Version**

  Format of the signed message. The dedicated protocol number '0' indicates that the packet is not signed—i.e. the security footer is not provided.

- **Transmission Time**

  A 64-bit integer, encoded in big-endian format, giving the number of microseconds since the Unix epoch (00:00:00 GMT, 1 January, 1970).

- **Transmission Location**

  The lat/long/altitude coordinates are used to represent position. The latitude and the longitude are encoded with a 6-bit resolution field and a 34-bit fixed-point value. The altitude contains an altitude position indicator in metres.

- **Lifetime**

  Validity period of the signed message in seconds.

- **Signer**

  This field contains the information to determine the keying material and hash algorithm used to sign the message—though not necessarily the identity of the signer.

- **Signature**

  The signature is assumed to be 165 bit long [13].

Frame	Field	Size
Security Header	Protocol Version	1 byte
Data		
Security Footer	Message Identifier	8 bytes
	Transmission Time	8 bytes
	Lifetime	1 byte
	Transmission Location - Latitude - Longitude - Altitude	12 bytes - 5 bytes - 5 bytes - 2 bytes
	Signer	20 bytes
	Signature	165 bits
Total		**71 bytes**

Table A.1: Security Frame: A V2V safety message requires an additional 71 bytes to sign it.

A.2 Overview of Signed Message Processing

A.2.1 Transmission Processing

Creating a signed message for transmission involves the following steps:

1. Generate a random message identifier.
2. Get the current position and time.
3. Encode the unsigned message.
4. Digitally sign the unsigned message.
5. Create and encode the signed message.

A.2.2 Reception Processing

When a unit receives a signed message it must use the following process, or its equivalent, in processing it.

1. Decode the message.
2. Check that the *Transmission Time* is within the acceptable time window. If not, discard the message.
3. Check that the *Transmission Position* is within the acceptable position window. If not, discard the message.
4. Look up the message in the cache of recently received messages. If the message has already been received, discard it as a replay.
5. If the sender's certificate contains a scope restriction, verify that the *Message Position* is within the geographic scope of every certificate in the sender's certification path. If not, discard the message.
6. Verify that the application field in the message is consistent with the scope restriction in the certificate.

7. Verify that the sender's certificate has not been revoked. If the sender's certificate has been revoked, discard the message.

8. Verify the sender's certificate. If the sender's certificate does not verify, discard the message.

9. Verify the signature on the message. If the signature does not verify, discard the message.

10. If all the previous tests verify, cache any previously unseen certificates and pass the message up to the application layer.

It should be noted that performing the steps in this order will minimise the number of public key operations. Thus, checks 1-7 all can be performed with minimal computational overhead. Check 8 can be cached if the same certificate is seen multiple times.

A.3 Policy Requirements

The OBU signing keys for safety messages must be embedded in a tamper-resistant Hardware Security Module (HSM). This HSM must be compliant with FIPS 140-2 level 3. The HSM must be designed not to release these signing keys from the module. In addition, it must not be usable for signing arbitrary messages.

All messages signed by the HSM must be wrapped in a *Signed Message* structure. The HSM must populate the *message id*, *Transmission Time*, and *Transmission Location* fields. The *Transmission Time* and *Transmission Location* fields must be populated with data received respectively from a clock and GPS unit which are housed within a FIPS 140-2 level 3 module. We recommend that the clock and GPS unit be housed within the same module as the signing module. However, if they are housed within a separate unit, then communications between the two modules must be authenticated with an algorithm which provides at least 100 bits of security and measures must be taken to ensure timing synchronisation between the two modules.

The clock must be periodically updated from the GPS unit in order to correct for clock slip. However, because the GPS unit gets its input from radio signals outside the tamper boundary, mechanisms must be used to isolate the system from GPS spoofing. The clock must be calibrated for maximum slip values and must not allow for corrections beyond those values. In addition, "backward" corrections must be performed by slow-ticking rather than by rolling time backward. In addition, the system should enforce physical plausibility rules, such as rejecting speeds in excess of the maximum speed of the vehicle, impossible altitudes, etc.

Appendix B

PHY Preamble and MAC Header

B.1 PHY Preamble

The PHY preamble for DSRC is the same as the one in 'IEEE 802.11a' and is depicted in Table B.1. The physical layer adds an additional 16 bytes of data to each packet.

	Data	Size
PLCP Preamble	Sync	80 bits
	Start Frame Delimiter	16 bits
PLCP Header	Signal	8 bits
	Service	8 bits
	Length	16 bits
Total		**16 bytes**

Table B.1: PHY Frame Format: The physical layer adds an additional 16 bytes of data to the packet.

B.2 MAC Header

Current DSRC prototypes are using the data frame format as standardised in 'IEEE 802.11' [6] and shown in Table B.2. The MAC frame adds an additional 34 bytes to each packet, but provides three fields, namely—address 3, sequence control and address 4—that are not required for safety communication. It is therefore suggested that a new MAC frame format is standardised in 'IEEE 802.11' as follows:

The MAC header does support different frame formats according to the two bit long subfield *Type* and the four bit long subfield *Subtype* in the *Frame Control* field. All packets containing a '10' in the *Type* subfield contain data that is meant to be delivered to the upper layer—they are neither management nor control frames. The frame format of the data packets are distinguished based on their subtype. Only 8 of the 16 available subtypes are defined yet, so a new subtype shall be defined providing a header that contains the essential fields only. This new header is depicted in Table B.3 and adds an additional 20 bytes to each packet only.

Data	Size
Frame Control	2 bytes
Duration	2 bytes
Address 1	6 bytes
Address 2	6 bytes
Address 3	6 bytes
Sequence Control	2 bytes
Address 4	6 bytes
Data	$0 - 2312$
FCS	4 bytes
Total	**34 byte**

Table B.2: MAC frame format: The MAC header adds an additional 34 bytes of data to the message.

Data	Size
Frame Control	2 bytes
Duration	2 bytes
Address 1	6 bytes
Address 2	6 bytes
Data	$0 - 2326$
FCS	4 bytes
Total	**20 byte**

Table B.3: New MAC frame format: The new MAC header adds only an additional 20 bytes of data to the message.

Appendix C

Safety Message Data

C.1 Safety Message Coding Scheme

According to the discussion in Section 5.4, the safety-message coding scheme has three demands:

1. The application can pick out the required safety information units.

2. The safety message can contain safety information units that not all applications understand.

3. The overall message size should be as small as possible.

In order to meet the first requirement, the different safety information units are suggested to be encoded individually providing a well-defined identifier for every possible safety information unit. The second demand requires framing the combination of data and identifier by indicating the frame size. Hence, the frame of a single safety unit contains the frame size, an identifier and the safety information unit, and is depicted in Table C.1.

Frame Format		
Length	Identifier	Safety Information Unit

Table C.1: Frame format of a single safety information unit.

In order to meet the third requirement, it is suggested that the most common combinations of safety information units are combined into a single frame—e.g. the combination velocity, acceleration and heading indicated by a well-defined identifier. This is shown in Table C.2. In order to meet the first demand, such combinations should be defined before the first generation of DSRC equipment gets integrated into vehicles.

Frame Format				
Length	Identifier	Safety Information Unit 1	Safety Information Unit 2	...

Table C.2: Frame format of combined safety information units.

C.2 Safety Message Content

Subsequently a discussion about the most important safety information units for V2V communication is held.

C.2.1 Time

It is important to know the exact time the message was generated. This allows determining how much time has passed meanwhile in order to extrapolate the new situation if necessary. Nowadays, a precise time adjustment, in order of microseconds, can be achieved using GPS [32]. It should be noted that the time is added in the security header.

C.2.2 Position

The vehicle receiving a message has to know its relative position to the message-sending vehicle—e.g. the vehicle that sent the message is 200 metres behind. It should be noted that information about the altitude is necessary too. Otherwise a vehicle driving on a overpass are likely to crash into the ones beneath—from the vehicles point of view—resulting in a series of unnecessary alerts.

It is important to know how accurate this position is. A state of the art stand-alone GPS can yield an accuracy of about seven metres. This is possible since the government-imposed degradation of the GPS signal—referred to as "Selective Availability"—was turned off in May 2001. Differential GPS (DGPS) uses code based differential corrections to achieve a positioning accuracy of about one metre. This can further be improved with a carrier based RTK GPS resulting in an accuracy on the order of centimetres. A summarisation of these values is given in Table C.3. Please refer to [34] for more detailed discussion about that topic. It should be noted that additional information—such as speed sensor, rate gyros or map data—can substantially improve the GPS accuracy and the confidence interval.

Positioning system	Achievable accuracy (m, 1σ)	Confidence interval (%, 3σ)	Maturity date	Volume cost at maturity date [$]
GPS	7.0	90	1998	$50 - 60$
DGPS	1.0	85	2004	$30 - 40$
RTK GPS	0.02	80	2012	$70 - 80$

Table C.3: Overview of the GPS accuracy. The data derives from [33].

In Section 4.3 it is shown that some applications require to know the lane the vehicles are driving in, in order to show their full potential. This requires a position accuracy of about $0.5 - 1$ metres and additional map data. According to Table C.3 this is not likely to be available in the first DSRC generation. It should be noted that the position is added in the security header.

C.2.3 Heading

Receiving the sending vehicle's heading is required to extrapolate its position. In addition, it allows to determine whether the sending vehicle is in the same direction and most likely on the same street. The encoding of the heading is shown in Table C.4 and C.5.

Field	Heading	Precision
Coding	LSBit = 0.05°, 0° = North, signed, positive: clockwise	Table C.5
Size [bits]	13	3

Table C.4: Safety message data: Heading

Precision	000	001	010	011	100	101	110	111
Value [degrees]	N/A	90	45	10	3	1	0.2	0.05

Table C.5: Heading precision scheme.

C.2.4 Vehicle Speed and Acceleration

Receiving the sending vehicle's speed allows to extrapolate the current situation. This extrapolation can be improved if the acceleration is known. The encoding of the vehicle speed is shown in Table C.6 and C.7, the one of the acceleration in Table C.8 and Table C.9.

Field	Velocity	Precision
Coding	LSBit = $0.05m/s$, unsigned	Table C.7
Size [bits]	13	3

Table C.6: Safety message data: Velocity

Precision	000	001	010	011	100	101	110	111
Value $[m/s]$	N/A	5	2	1	0.5	0.2	0.1	0.05

Table C.7: Velocity precision scheme.

Field	Longitudinal		Lateral	
	Acceleration	Precision	Acceleration	Precision
Coding	LSBit = $0.01m/s^2$, signed	Table C.9	LSBit = $0.01m/s^2$, signed	Table C.9
Size [bits]	13	3	13	3

Table C.8: Safety Message Data: Acceleration

Precision	000	001	010	011	100	101	110	111
Value $[m/s^2]$	N/A	100	20	5	1	0.2	0.05	0.01

Table C.9: Acceleration precision scheme.

C.2.5 Break Applied Status and Pressure

There are different possibilities that cause a vehicle to decelerate: It might be going upwards, the driver is not pushing the accelerating pedal anymore or the driver is braking. Hence the information if the braking pedal is applied allows a better judgement of the current situation. The encoding is shown in Table C.10.

Field	Break Applied Status	Break Pressure	Anti-lock braking
Code	0 On 1 Off	00000 N/A 00001 Minimum pressure 00010 ... 11111 Maximum pressure	00 N/A 01 Off 10 On 11 Engaged
Size [bits]	1	5	2

Table C.10: Safety message data: Brake.

C.2.6 Vehicle Type

Different types of vehicles can have a radically different behaviour in certain situations—e.g. a truck has a much longer braking distance than a passenger car or a bike. The information about the vehicle type should therefore be included in the message. The most general approach is to send the vehicle's weight and dimension, but the dimension could be encoded in predefined ranges to reduce the size of this information unit. A possible encoding is shown in Table C.11.

Field	Width	Height	Length	Weight
Coding	LSBit = $0.01m$, unsigned, '00...0' = N/A	LSBit = $0.01m$, unsigned, '00...0' = N/A	LSBit = $0.01m$, unsigned, '00...0' = N/A	LSBit = 5kg, unsigned, '00...0' = N/A
Size [bits]	10	10	14	14

Table C.11: Safety message data: Dimension and weight.

C.2.7 Signal and Lights

This information indicates that the vehicle has its hazard lights on, or indicates a lange change/turn. The encoding is shown in Table C.12.

Field	Turn Signal / Hazard Lights	Headlights	Availability	Filler
Code	00 = Off 01 = Left Turn 10 = Right Turn 11 = Hazard	00 = Off 01 = Parking Light 10 = Low Beam 11 = High Beam	XX1 = Turn Signal N/A X1X = Hazard Light N/A 1XX = Headlight N/A	
Size [bits]	2	2	3	1

Table C.12: Safety Message Data: Signal and lights.

C.2.8 Steering Wheel Angle

The steering wheel angle might be of interest if the vehicle is out of control. The encoding is shown in Table C.13 and C.14.

Field	Steering Wheel Angle	Precision
Coding	LSBit = 0.2°, 0° = Straight Forward, signed, positive: clockwise	Table C.14
Size [bits]	13	3

Table C.13: Safety message data: Steering wheel angle.

Precision	000	001	010	011	100	101	110	111
Value [degrees]	N/A	20	10	5	2	1	0.5	0.2

Table C.14: Steering wheel angle precision scheme.

Bibliography

[1] Vehicle Safety Communications Project. Task 3 Final Report. Identify Intelligent Vehicle Safety Applications Enabled by DSRC. *Vehicle Safety Communications Consortium – CAMP.* 2004.

[2] Vehicle Safety Communications Project. Task 6D: WAVE Radio Module – Final Task Report. *Vehicle Safety Communications Consortium – CAMP.* 2004.

[3] Vehicle Safety Communications Project. Task 11: WAVE/DSRC Security Extension – Final Report. *Vehicle Safety Communications Consortium – CAMP.* October 2004.

[4] Vehicle Safety Communications Project. Task 12: Intersection and Intervehicle Safety Messaging Effectiveness Evaluation. Final Task Report. *Vehicle Safety Communications Consortium – CAMP.* October 2004.

[5] IEEE Std 802.2-2001: IEEE Standard for Information technology – Telecommunications and information exchange between systems – Local and metropolitan area networks – Specific requirements. Part 2: Logical Link Control.

[6] IEEE Std 802.11-1999(R2003): IEEE Standard for Information technology – Telecommunications and information exchange between systems – Local and metropolitan area networks – Specific requirements. Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) specifications.

[7] IEEE Std 802.11a-1999(R2003): Supplement to IEEE Standard for Information technology – Telecommunications and information exchange between systems – Local and metropolitan area networks. Specific requirements. Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) specifications High-speed Physical Layer in the 5 GHz Band.

[8] IEEE Std 802.11i-2004: IEEE Standard for Information technology – Telecommunications and information exchange between systems – Local and metropolitan area networks – Specific requirements. Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) specifications. Amendment 6: Medium Access Control (MAC) Security Enhancements.

[9] IEEE 802.11p Task Group.
http://grouper.ieee.org/groups/scc32/dsrc/index.html

[10] IEEE Std P1609.1-2005 – Draft.

[11] IEEE Std P1609.3-2005 – Draft.

[12] IEEE Std P1609.4-2005 – Draft.

[13] IEEE Std 1556 (Draft Version 1) – Draft Standard for 5.9 GHz Intelligent Transportation System (ITS) Radio Service Security.

[14] Federal Communications Commission. *FCC 99-305.* FCC Report and Order, October 1999.

[15] Federal Communications Commission. *FCC 03-024.* FCC Report and Order, Feburary 2004.

[16] H. Füssler, M. Torrent-Moreno, M. Transier, A. Festag, H. Hartenstein. Thoughts on a Protocol Architecture for Vehicular Ad-Hoc Networks. *Accepted for 2nd International Workshop in Intelligent Transportation (WIT 2005).* March 2005.

[17] V. Taliwal, D. Jiang, H. Mangold, C. Chen, R. Sengupta. Empirical determination of channel characteristics for DSRC vehicle-to-vehicle communication. *VANET '04: Proceedings of the first ACM workshop on Vehicular ad-hoc networks.* Philadelphia, PA, 2004.

[18] M. Torrent-Moreno, D. Jiang, H. Hartenstein. Broadcast reception rates and effects of priority access in 802.11-based vehicular ad-hoc networks. *in Proc. of ACM VANET 2004.* Philadelphia, PA, October 2004, pp. 1018.

[19] Q. Xu, D. Jiang. Design and Analysis of Highway Safety Communication Protocol in 5.9GHz Dedicated Short Range Communication Spectrum. *Proc. of the IEEE Conference of Vehicular Technology.* Korea, Spring 2003.

[20] X. Yang, J. Liu, F. Zhao, N. Vaidya. A Vehicle-to-Vehicle Communication Protocol for Cooperative Collision Warning. *Technical Report. University of Illinois at Urbana-Champaign.* December 2003.

[21] Q. Xu, T. Mak, J. Ko, R. Sengupta. Vehicle-to-Vehicle Messaging in DSRC. *First ACM Workshop on Vehicular Ad-Hoc Networks.* 2004.

[22] H. Füssler, H. Hartenstein, J. Widmer, M. Mauve, W. Effelsberg. Contention-Based Forwarding for Street Scenarios. *1st International Workshop in Intelligent Transportation (WIT 2004).* pp. 155-159, Hamburg, Germany, March 2004.

[23] Vehicle Infrastructure Integration – Architecture and Functional Requirements. *FHWA.* January 2005.

[24] R. Braden, T. Faber, M. Handley. From Protocol Stack to Protocol Heap – Role-Based Architecture. *First Workshop on Hot Topics in Networks (HotNets-I).* October, 2002.

[25] C. Tschudin. Flexibile Protocol Stacks. *in Proc. of ACM SIGCOMM 1991.* pp. 197-208, Zurich, Switzerland, September 1991.

[26] C. Maihöfer. A Survey of Geocast Routing Protocols. *IEEE Communications Surveys & Tutorials.* Vol. 6, no. 2, pp. 32-42, Q2 2004.

[27] M. Green. 'How Long Does It Take to Stop?' Methodological Analysis of Driver Perception-Brake Times. *Transportation Human Factors.* Vol. 2, No. 3 pp. 195216, 2000.

[28] A. S. Tanenbaum. Computer Networks. *Prentice Hall.* 4th edition, 2003.

[29] D. J. C. MacKay. Information Theory, Interference, and Learning Algorithms. Chapter 50 – Digital Fountain Codes.
http://www.inference.phy.cam.ac.uk/mackay/DFountain.html

[30] Digital Fountain – Core Technology Overview – *www.digitalfountain.com*

[31] M. Luby. LT Codes. *Whitepapaer – www.digitalfountain.com*

[32] D. W. Allan, N. Ashby, C. C. Hodge. The Science of timekeeping. *Hewlett Packard Application Note 1289.* 1997.

[33] J. Wang. An Examination of Vehicle Positioning and Application Requirements. *DaimlerChrysler Telematics Research – Internal Report.* April, 2004.

[34] GPS-Infos
http://www.kowoma.de/gps/index.htm

[35] IETF Request for Comments: RFC 2460, Internet Protocol, Version 6 (IPv6)
http://www.ietf.org/rfc/rfc2460.txt

[36] IETF Request for Comments: RFC 791, Internet Protocol
http://www.ietf.org/rfc/rfc791.txt

[37] IETF Request for Comments: RFC 768, User Datagram Protocol
http://www.ietf.org/rfc/rfc768.txt

[38] IETF Request for Comments: RFC 793, Transmission Control Protocol
http://www.ietf.org/rfc/rfc768.txt

[39] The DSRC Project
http://www.leearmstrong.com/DSRC/DSRCHomeset.htm

[40] Network on Wheels (NOW)
http://www.network-on-wheels.de/

[41] FleetNet
http://www.et2.tu-harburg.de/fleetnet/

[42] Inegrated Project PReVENT
http://www.prevent-ip.org/

[43] Car-to-Car Communication Consortium
http://www.car-2-car.org/

[44] Intelligent Transportation Society (ITS)
http://www.itsa.org/

[45] PATH
http://path.berkeley.edu/dsrc/index.html

[46] Federal Communications Commission (FCC)
http://wireless.fcc.gov/services/its/dsrc/

[47] Atheros Communications
http://www.atheros.com/

[48] DENSO International America, Incorporated
http://www.densocorp-na.com/