Advanced Transport Protocols for Wireless and Mobile Ad Hoc Networks

Inauguraldissertation
zur Erlangung des akademischen Grades
eines Doktors der Naturwissenschaften
der Universität Mannheim

vorgelegt von

Dipl.-Ing. (FH) Ralf Schmitz
aus Heidelberg

Mannheim, 2006

Dekan: Professor Dr. Matthias Krause, Universität Mannheim
Referent: Professor Dr. Wolfgang Effelsberg, Universität Mannheim
Korreferent: Professor Dr. Hannes Hartenstein, Universität Karlsruhe (TH)

Tag der mündlichen Prüfung: 26. Februar 2007

Zusammenfassung

Die Themen dieser Arbeit umfassen Transportprotokolle in den folgenden Forschungsgebieten:
Fast Handover ermöglichen mobilen IP Endgeräten den unterbrechungsfreien Übergang zwischen Zugangsroutern bei direktem, drahtlosem Zugang in ein Infrastrukturnetz (z.B. Internet). Der *Fast Handover* Algorithmus wurde optimiert und die Leistung der Transportprotokolle UDP und TCP während des Zellübergangs durch Messung quantitativ ermittelt und bewertet.
Im Folgenden beschäftigt sich die Arbeit mit fahrzeugbasierten Ad Hoc Netzwerken. Für diese Netze wird ein Punkt-zu-Punkt Transportprotokoll und ein Algorithmus zur zuverlässigen und effizienten Informationsverteilung in einem geografischen Zielgebiet entwickelt und durch Simulationen bewertet.
Abschließend wird der Einfluss von Schwankungen der Signalstärke auf die Leistung eines Ad Hoc Netzwerkes untersucht. Messungen ermitteln die Höhe und Verteilung dieser Schwankungen. Aus diesen Ergebnissen wird ein einfaches, aber realistisches Funkmodell entwickelt, welches den Einfluss auf die Leistung eines Ad Hoc Netzwerks bewertet. Daraus ergeben sich Vorschläge und Richtlinien für zukünftiges Protokolldesign.

Advanced Transport Protocols for Wireless and Mobile Ad Hoc Networks

by

Ralf Schmitz

Dissertation in co-operation between

Faculty of Computer Science IV, University of Mannheim, Germany
and
NEC Europe Ltd., Network Laboratories Heidelberg, Germany

2006

Advisor and Co-Advisor:

Prof. Dr. Wolfgang Effelsberg, University of Mannheim, Germany

Prof. Dr. Hannes Hartenstein, University of Karlsruhe, Germany

Abstract

This thesis comprises transport protocols in the following different areas of research:
Fast Handover allows mobile IP end-devices to roam between wireless access routers without interruptions while communicating to devices in an infrastructure (e.g., in the Internet). This work optimizes the *Fast Handover* algorithm and evaluates the performance of the transport protocols UDP and TCP during fast handovers via measurements.
The following part of the thesis focuses on vehicular ad hoc networks. The thesis designs and evaluates through simulations a point-to-point transport protocol for vehicular ad hoc networks and an algorithm to facilitate the reliable and efficient distribution of information in a geographically scoped target area.
Finally, the thesis evaluates the impact of wireless radio fluctuations on the performance of an Ad Hoc Network. Measurements quantify the wireless radio fluctuations. Based on these results, the thesis develops a simple but realistic radio model that evaluates by means of simulations the impact on the performance of an ad hoc network. As a result, the work provides guidelines for future ad hoc protocol design.

Acknowledgements

First of all I thank my supervisors Prof. Dr. Effelsberg and Prof. Dr. Hartenstein for all their efforts.

I like to thank the NEC Network Laboratories, headed by Dr. Heinrich Stüttgen, and particularly the Mobile Internet Group headed by Amardeo Sarma, for the specific contractual framework that made this thesis possible.

Last but not least a big thanks goes to Silke Lampson, Christian Kücherer, Sabine Schneider, Martim Ferrer and my girlfriend Nicole Prinz for all their support, reading, good comments, motivation and friendship during this thesis.

Contents

List of Figures

List of Tables

Glossary

3G	Third Generation Mobile Communication Network
AAA	Authorization, Authentication and Accounting
ACK	Acknowledgment
ACTP	Application Controlled Transport Protocol
ADSN	Ack Duplication Sequence Number
AODV	Ad Hoc On Demand Distance Vector Routing
AR	Access Router
ATCP	Ad Hoc TCP
ATP	Ad Hoc Transport Protocol
BACK	Binding Acknowledgment
BCMP	Brain Candidate Mobility Protocol
BER	Bit Error Rate
BMB+F	Bundenministerium Bildung und Forschung - German Ministry for Education and Research
BU	Binding Update
CBF	Contention-Based Forwarding
CBR	Constant Bit Rate Traffic
CDF	Cumulative Distribution Function
CN	Correspondent Node
CoA	Care-of Address
CSMA/CA	Carrier Sense Multiple Access with Collision Avoidance
CTS	Clear To Send
DCF	Distributed Coordination Function
DSDV	Destination-Sequenced Distance-Vector Routing
DSL	Digital Subscriber Line
DSR	Dynamic Source Routing

EBSN	Explicit Bad State Notification
ERDN	Explicit Route Disconnection Notification
ERSN	Explicit Route Successful Notification
FHO	Fast Handover
FIFO	First In - First Out
FMIP	Fast Handovers for Mobile IPv6
GAMER	GeoCast Adaptive Mesh Environment Routing
GeoCast	Geographically scoped Broadcast
GPS	Global Positioning System
GPSR	Greedy Perimeter Stateless Forwarding
HA	Home Agent
HMIP	Hierarchical Mobile IPv6 Mobility Management
HWGUI	Highway Graphical User Interface
ICMP	Internet Control Message Protocol
IEEE	Institute of Eletrical and Electronics Engineers
IETF	Internet Engineering Task Force
IP	Internet Protocol - version 4 or version 6
LACK	Local Acknowledgment
LAN	Local Area Network
LBM	Location-Based Multicast
MAC	Medium Access Control
MANET	Mobile Ad Hoc Network
MGRP	Mesh-Based GeoCast Routing Protocol
MIPv6	Mobile IPv6
MN	Mobile Node
MPR	Multipoint Relay
MTNM	Mobile Terminal Network Manager
MTU	Maximum Transfer Unit
NACK	Negative Acknowledgment
NAT	Network Address Translation
NoW	Network on Wheels Project
NS-2	Network Simulator version 2
Obs	Observer

OFSGP/OFMGP	Obstacle-Free Single / Multi-Destination Geo-casting Protocol
OLSR	Optimized Link State Routing
PBR	Position-Based Routing
PN	Pivoting Node
QoS	Quality of Service
RFN	Route Failure Notification
RMFTP	Reliable Multicast File Transfer Protocol
RMTP	Reliable Multicast Transport Protocol
RRN	Route Re-Establishment Notification
RTO	Retransmission Timeout
RTS	Ready To Send
RTT	Round Trip Time
RWP	Random Way-Point
SACK	Selective Acknowledgment
SBA	Scalable Broadcast Algorithm
SRF	Simple Reliable Flooding
SRGF	Simple Reliable Geographical Flooding
SRM	Scalable Reliable Multicast
TA	Target Area
TCP	Transmission Control Protocol
TCP DOOR	TCP Detection of Out-of-Order and Response
TCP-BuS	TCP Buffering Capability and Sequence Info
TCP-F	TCP-Feedback
TCR	Topological Change Rate
TERGF	Time-Extended Reliable Geographical Flooding
TORA	Temporary Ordered Routing Algorithm
TPSN	TCP packet Sequence Number
UDP	User Datagram Protocol
UMTS	Universal Mobile Telecommunication System
VANET	Vehicular Ad Hoc Network
VTP	Vehicular Transport Protocol
WLAN	Wireless Local Area Network
ZRP	Zone Routing Protocol

Chapter 1

Introduction

Communication and computer technology experienced a tremendous advance and growth in the recent years. There is the success attained by the Internet on the one hand and the boom of mobile communication via cell phones on the other hand, just to mention the most prominent examples. Broadband Internet access provided to households (e.g., through Digital Subscriber Line (DSL), television, power cable or even satellites), as well as mobile phones have become part of the daily life of the majority of the population in industrialized countries.

The vision of ubiquitous computing with many invisible computers per person surrounding and supporting the user any time and any place with a broad variety of applications to use comes closer to reality than ever. The availability of portable, powerful communication and computing devices paves the way for technological expansion in the future evolution, like merging Internet access and mobility as a next step. Expanding the capabilities and protocols of mobile devices enhances the services that can be offered to mobile users in a manifold fashion. As a simple example, one could consider a voyager in a foreign city. The wireless access to the infrastructure or even the connection to further travelers along his route may supply him with information according to his needs, such as traffic conditions, parking availability, hotel or sightseeing information and many more. The demand for communication in self-organized, mobile networks has arisen.

The technical implementation of ad hoc networks requires communication protocols, such as network and transport protocols. Established protocols, like Internet protocols, cannot be adopted to mobile ad hoc networks, as these protocols are designed for hard-wired infrastructure networks. As a result, they perform poorly. The performance of routing and transport protocols in ad hoc networks mainly depends on the ability to adapt to changes in the connectivity or network topology. In addition to the routing protocol development in the current research, this thesis focuses on the design and performance of transport protocols for wireless and mobile ad hoc networks in order to bring ad hoc networks closer to reality.

1.1 Contribution of this thesis

This thesis covers several independent topics of actual research in the field of wireless, mobile ad hoc networks. The work comprises the design and performance evaluation of transport layer algorithms in wireless and mobile ad hoc networks.

The Internet provides two levels of transport algorithms. The *connectionless* user datagram protocol (UDP) is a lightweight protocol that offers minimal transport services by simply injecting packets into a packet-switched network. In contrast, the *connection-oriented* transmission control protocol (TCP) provides reliable, efficient and in-order packet transmissions. Reliable data transmission includes retransmission of lost packets. TCP uses acknowledgments to indicate successful transmission or detect the loss of packets. Flow control mechanisms pace the transmission rate to the receiver's buffer capacity, and congestion control mechanisms avoid and resolve network congestion. Finally, TCP orders packets by sequence numbers to provide in-order data delivery to the applications. This thesis uses the Internet transport protocols, particularly TCP, as a reference for performance evaluations in ad hoc networks.

The first part of the thesis evaluates UDP and TCP performance in presence of Mobile IPv6 handovers. A handover represents the process when a roaming mobile user changes its single-hop access point to the Internet. In this thesis, the *Fast Handover in IPv6* algorithm, as discussed and defined in the Internet Engineering Task Force (IETF) [61] , is enhanced and implemented to suit in an All-IP network. An All-IP network deploys IP technology up to the wireless end device. This network comprises IPv6-based mobility management and includes quality of service (QoS), Authentication, Authorization, Accounting and Charging (AAAC). The fast handover algorithm and its implementation is optimized for this environment. In order to measures the UDP and TCP performance in a real All-IP test network upon handover occurrence, this thesis presents the results of comparing the IPv6 standard and the fast handover approach.

The second part of the thesis advances to transport issues in wireless and mobile ad hoc networks, such as vehicular ad hoc networks (VANETs). VANETs are a promising candidate to deploy ad hoc networks in practical and valuable applications, by enabling multi-hop vehicle-to-vehicle and vehicle-to-roadside communication. Thus, existing applications can be integrated into vehicles to provide a broad range of new exciting applications. These applications fall into two categories: Unicast applications, such as media transmission or email, and broadcast applications, such as active road traffic safety or forecast services. Both areas pose specific requirements on the transport layer that demand for novel transport algorithms, as developed in the second part of this thesis.

Finally, the third part of the thesis evaluates the impact of wireless signal strength fluctuations since field measurements in the ad hoc protocol development for VANETs have shown significant effect of radio characteristics on the network performance. The performance of an ad hoc network mainly depends on the ability of its network protocols to adapt to topology changes. In a mobile ad hoc network, the main reason for topology changes is the continuous node movement. However, signal strength fluctuations also contribute significantly to the topological change rate, e.g., because radio propagation characteristics change the radio transmission range over time. Thus, the final part of the thesis quantifies signal strength fluctuations in field measurements, derives a simple but realistic radio model out of the measurement results and provides guidelines for ad hoc network simulations and protocol design.

Summarizing, this thesis covers different aspects of research in the area of ad hoc networks. The results contribute to the deployment of future wireless and mobile ad hoc networks by increasing robustness and performance of these networks. The definition of an advanced transport protocol for mobile ad hoc networks of VANETs provides an important achievement. Furthermore, the evaluation of radio characteristics on the ad hoc network performance allows giving guidelines for future protocol design and simulative evaluations.

1.2 Outline and Structure of the thesis

This thesis is structured in five chapters. It begins with single-hop wireless connectivity of roaming mobile devices that connect to an infrastructure. Subsequently, it advances to pure mobile ad hoc networks for point-to-point and point-to-multipoint vehicular communication. Finally, it evaluates the impact of radio characteristics on the ad hoc network performance.

Chapter 2 evaluates UDP and TCP transport layer performance in the presence of handovers in a Mobile IPv6 environment, including quality of service (QoS), Authorization, Authentication, Accounting and Charging (AAAC). Uninterrupted services are crucial for roaming users since interruption of data streams are not tolerable for certain applications, like voice calls or video streams. Particularly when the mobile device changes its access points to the fixed network in a so-called *handover*, the user should not experience a noticeable interruption during a real-time communication or performance degradation of a download. This work, as developed in the framework of the IST project Moby Dick, integrates and evolves the IETF *Fast Handovers in Mobile IPv6* approach towards the IPv6-based mobility enabled architecture that comprises QoS and AAAC. It measures handover interruption and transport protocol performance in an All-IP test network and compares transport protocol performance upon standard Mobile IPv6 and fast handovers for IPv6 handover.

Chapter 3 presents the design of a unicast vehicular transport protocol (VTP), which is tailored to the unique characteristics of VANETs and evaluates simulatively the protocol performance. Unicast applications in VANETs require reliable and in-order delivery of data in combination with flow and congestion control, similar to the service provided by TCP in the Internet. However, TCP performs poorly in wireless networks, particularly, in mobile ad hoc networks. The unique characteristics of VANETS, such as frequent topology changes, round trip time (RTT) jitter or reordering, necessitate the development of a new transport protocol that is specifically tailored to these conditions. Prior to the design of a vehicular transport protocol (VTP), the chapter analyzes the path characteristics communication/disruption duration, packet loss, packet reordering, RTT and RTT jitter for typical German highway scenarios. Based on these results, the chapter describes the design and evaluation of VTP that consider the unique characteristics of the vehicular environment. A simulative evaluation compares the VTP performance against standard TCP.

Chapter 4 designs and evaluates an efficient, time-extended reliable geographical flooding (TERGF) algorithm. This algorithm provides efficient and reliable distribution of information in a geographical area over time, which is important to inform vehicles in a target area, e.g., for safety applications. Particularly, TERGF informs vehicles that enter the target area after the initial distribution of the message. TERGF combines GeoCast, self-pruned flooding (i.e., explicitly addressing all single-hop neighbors in the header of the broadcast message) and acknowledgment-based reliability via passive acknowledgments in order to provide time-extended reliability in a geographical area. A simulative evaluation compares TERGF and the state-of-the art GeoCast algorithm.

Chapter 5 evaluates the impact of radio signal strength fluctuations on the ad hoc network performance. Field measurements in the framework of protocol evaluation have shown a significant impact of these fluctuations. The thesis measures signal strength fluctuations in field trial experiments, considering a *best case* stationary scenario without obstacles around and with sender and receiver in line-of-sight. Based on the measurement results, the evaluation derives a simple, but realistic signal strength fluctuation model. A simulative study uses this model to quantify the impact of signal strength fluctuations on the metrics topological change rate and link stability, which directly relate to the ad hoc network performance. The simulation results compare the impact of signal strength fluctuations with the impact of mobility. The results provide guidelines for ad hoc network simulations and protocol design.

Finally, Chapter 6 concludes the work by summarizing and interconnecting the results of the separate chapters. It points out possible areas of future research to carry on the work and bring mobile ad hoc networks into life.

Chapter 2

Transport Protocol Evaluation in Presence of Fast Handovers

2.1 Introduction

The success attained by the fixed-line Internet and the success of mobile cell phone networks facilitate the vision of an *All-IP* next generation network. These All-IP networks integrate the different philosophies of both environments. Mobile Internet technology is moving towards a packet-based (i.e., IPv6- based) network, empowered by the availability of portable, powerful computing and communication devices. This vision creates the demand for a mobility-enabled and security-aware architecture, including Quality of Service (QoS), which is independent of the access technology.

The EU IST project Moby Dick [94] has taken on the challenge of providing a solution that integrates IP-based mobility, QoS and AAAC (Authorization, Authentication, Accounting and Charging). The project integrates so far separated approaches in a heterogeneous access environment. The Moby Dick design is independent of the deployed access technology, and the implementation employs IEEE 802.11b Wireless LAN [63], TD-CDMA of the Universal Mobile Telecommunication System (UMTS) [130] and Ethernet [64] as exemplary access technologies. The Moby Dick architecture focuses purely on the next generation Internet protocol IPv6 [115], *(i)* to account for the rapidly growing number of mobile devices and *(ii)* because the current Internet protocol IPv4 was not designed taking terminal mobility into account.

The design of IPv6 already considers portability, i.e., devices auto-configure their network settings at boot time in order to allow network connectivity at different locations. However, a global mobility management scheme is required in order to support global reachability and transparent mobility, as provided by Mobile IPv6 (MIPv6) [33]. MIPv6 provides the continuation of ongoing connections when the mobile device changes its access points to the fixed network, i.e., when the routing responsibility for a mobile node changes. This process is termed *handover*. In this

context, the provisioning of Fast Handovers (FHO) with small handover latency / interruption time and small or zero packet loss, in next generation mobile IP networks is essential in order to provide

(i) uninterrupted real-time services (e.g., real-time audio/video streaming)
(ii) acceptable download performance for connection-oriented data streams.

Roaming mobile users demand for IP services with quality comparable to traditional networks and uninterrupted real-time services as experienced in today's cellular mobile phone system. Therefore, the detailed analysis of transport-layer protocol performance, in presence of Fast Handovers in Mobile IPv6, is essential for the deployment of future packet switched networks. The minimization of handover latency and packet loss aims at avoiding noticeable communication disruption in real-time UDP data streams and preventing performance degradation due to mobility in TCP connections.

This chapter explains the integrated Moby Dick architecture and evaluates the network performance in presence of Fast Handovers.

2.2 Background

This section describes the basic protocols employed in the evaluation of this chapter, including the *Internet protocol version 6 (IPv6)*, the Internet transport protocols *user data protocol (UDP)* and *transmission control protocol (TCP)*, the mobility support for IPv6 *Mobile IPv6* and the *fast handovers in Mobile IPv6* extension, which aims at reducing interruptions and packet loss due to handover.

2.2.1 Internet Protocol Version 6 (IPv6)

The Internet Protocol version 6 (IPv6) [115], as designed and specified by the Internet Engineering Task Force (IETF) [61], defines a successor of the current IPv4. The design of IPv6 overcomes constraints of IPv4, such as the limited number of available addresses (i.e., 2^{32}), ineffective allocation of addresses by a class hierarchy or unmanageably large routing tables.

Particularly, the limited number of addresses is problematic when considering the rapidly growing number of Internet hosts. Approaches like Network Address Translation (NAT) [38] hide the shortage of IP addresses by mapping internal addresses to a small set of global, external addresses. However, such approaches increase network complexity and raise scalability problems. Thus, an adequate and manageable global address space in the design of IPv6 is essential for future networking.

Beyond that, the trend towards mobile networking with small, powerful communication devices demands for mobility support in IPv6.

The following list summarizes the key features in the design of IPv6.

Expanded address space and capabilities. The 128 bit address size of IPv6 represents a significant increase compared to 32 bit addresses in IPv4. This change supports a much greater number of directly addressable nodes, more levels of addressing hierarchy and simpler auto-configuration mechanisms. The auto-configuration uses the IPv6 address concept that includes the medium access control (MAC) protocol address in the IPv6 address because a MAC address is typically a unique identifier. Beyond, scalability of multicast routing is improved by adding a *scope* field to multicast addresses. Finally, a new address type termed *anycast address* is introduced. An anycast packet is destined for *any* node of a specific group of nodes.

Header format simplification and improved support for extensions. The IPv6 header defines optional fields that are only used when needed. This reduces the average header processing costs and average required bandwidth when compared to IPv4 where all header fields are always present.

The concept of *header chaining* increases the flexibility and processing while avoiding unnecessary headers. Optional information or protocol headers are encoded as separate headers that may be placed between the IPv6 and upper layer headers, as shown in Figure 2.1.

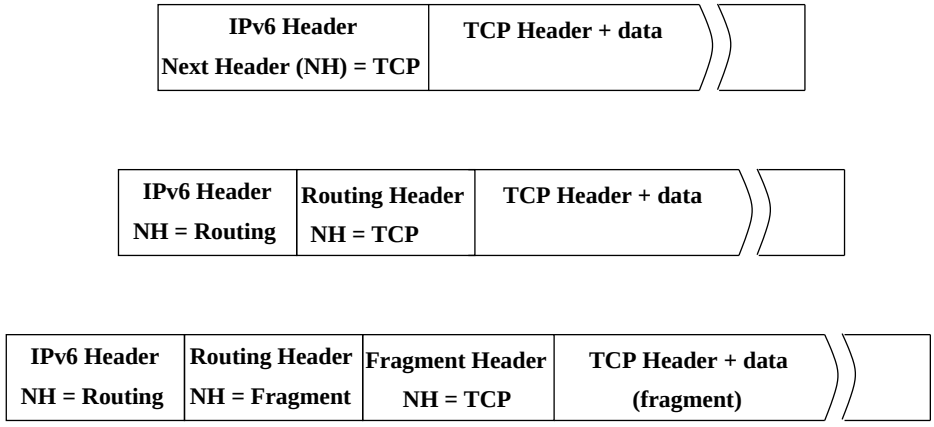

Figure 2.1: Examples of IPv6 Header Chaining.

The header processing is improved by the elimination of the header checksum. Intermediate nodes, such as routers, need not calculate the checksum, which increases the forwarding and routing performance. IPv6 assumes that higher layer protocols provide their own checksum if required.

Flow label capability. IPv6 provides the capability to label packets of specific flows. This allows routers to treat packets of different flows differently. In this way, Quality of Service (e.g., real-time services), as requested by the senders, is supported in the inner nodes of the network.

Reasonable security. IPv6 defines extensions for authentication support, data integrity, and data confidentiality (i.e., the latter is optional).

Reduced administrative overhead by auto-configuration. A major goal of IPv6 is the simplification of administration work. Protocols such as Stateless Address Autoconfiguration [120] and Neighbor Discovery [126] enable auto-configuration of IPv6 nodes to acquire their network settings (e.g., IPv6 address or default gateway) without human support or interaction. Beyond, auto-configuration and the support for address renumbering already supports *nomadic* computing: Nodes are no longer bound to a static position. A node can connect to the network at different access points. However, IPv6 does not provide ongoing connections. The *nomadic networking* of IPv6 requires a reboot or network re-start for each change of access points. Thus, the auto-configuration capabilities of IPv6 provide a basic support for the emerging mobility of users.

2.2.2 The User Datagram Protocol (UDP)

The user datagram protocol (UDP) [65] is a connectionless transport protocol. It is commonly employed on top of packet switched IP networks. UDP offers a minimal transport service, which allows applications to directly access the datagram service of the IP layer. UDP does not provide reliability or error recovery. The only services provided by UDP are checksum calculation and multiplexing by port number.

Typically, applications with specific requirements use UDP, such as real-time applications like IP telephony or video conferencing. These applications do not require reliability or congestion control, but rather aggressively use the network according to their bandwidth and delay requirements. UDP can quickly transmit data because it introduces only minimal overhead. For real-time applications, packet loss up to a certain limit is tolerable because human perception is not sensible to small interruptions (i.e., depending on the codec, assuming that the codec is able to cope with packet loss). The retransmission of a lost packet would be useless anyway due to the delay boundaries in real-time communications. The retransmission of a lost packet would just waste bandwidth because the receiver drops the packet when it is too late, i.e., the data stream has already been presented to the user. Furthermore, real-time applications require a constant transmission rate, which conflicts with transport layer services, such as congestion control. A throughput reduction due to congestion control could result in a quality decrease or even connection abort. Applications that use UDP optimistically rely on best effort networking service.

2.2.3 The Transmission Control Protocol (TCP)

The Transmission Control Protocol (TCP) [66] is a connection-oriented protocol, which provides reliable and in-order byte stream delivery to applications. TCP

uses a *sliding window* mechanism in combination with timers (e.g., retransmission timer RTO) in order to adapt to network conditions and retransmit lost packets to provide reliability. The window size determines the number of bytes of data that can be sent before an acknowledgment from the receiver must arrive. TCP establishes a full-duplex virtual connection between two endpoints where the IP address and the port number define each endpoint. The byte stream is transferred in segments. Typically, applications that require guaranteed delivery of data use TCP as their transport layer.

The TCP algorithm as a whole is quite complex and there are many different versions and extension proposals available. Therefore, this section provides an overview of the main TCP algorithms and components, referring to the different TCP versions. For more detailed information, the reader is referred to the respective references.

TCP congestion control: Slow start and congestion avoidance phase. TCP provides window-based congestion control in order to avoid network overload and resolve network congestion. TCP assumes network congestion upon the detection of packet loss. The different phases of TCP's congestion control are illustrated in Figure 2.2.

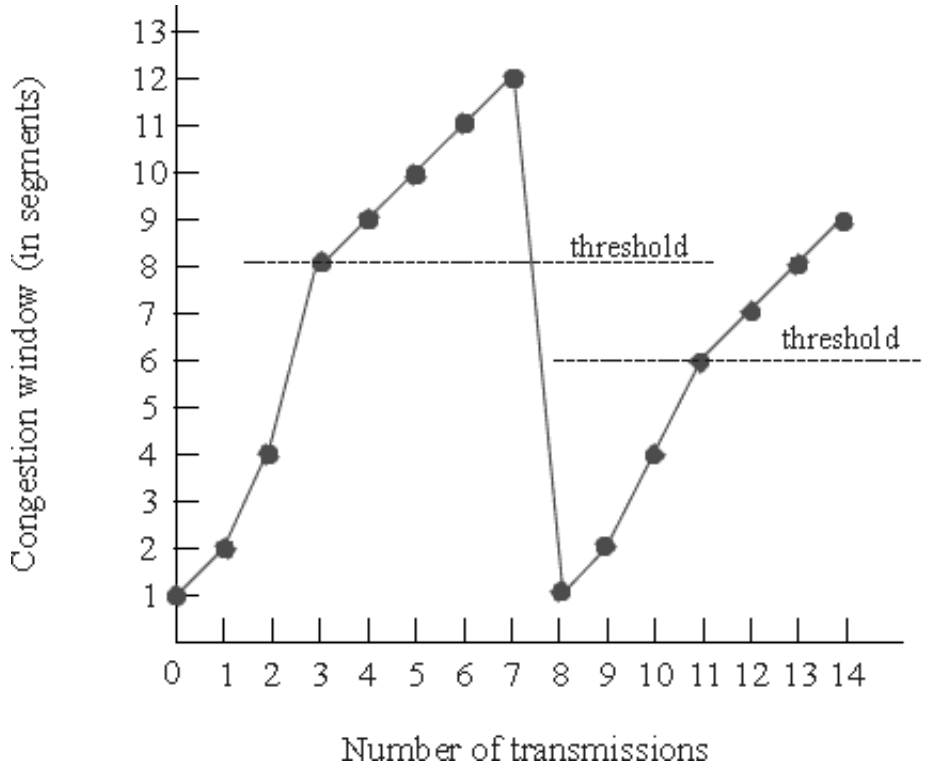

Figure 2.2: TCP congestion window illustration of [121].

When a new TCP connection is established between two end points, the Slow-Start (SS) mechanism takes place in order to probe the capacity of the network: Starting from a congestion window (cwnd) equal to one, the sender increases the cwnd by the current segment size upon reception of a new (i.e., non-duplicated) acknowledgment (ACK). In this way, the window size increases exponentially up to an estimated capacity, termed SS-threshold. Figure 2.2 illustrates the slow-start algorithm for the first four packets (i.e., between packet number zero and three). When the SS-threshold is reached, TCP enters the Congestion Avoidance (CA) phase, as shown between packet number four and seven. In CA, the cwnd increases linearly up to the receiver's maximum advertised window or until packet loss is detected. Regular TCP (Tahoe) [66] assumes a packet loss when the retransmission timer (RTO) expires before the respective segment is acknowledged. In this case, SS is unavoidable. In Figure 2.2, the sender detects a packet loss upon the transmission of packet number seven. Consequently, it enters SS: When the sender transmits packet number eight, the cwnd is reset to one and the SS-threshold drops by half of the current cwnd.

TCP fast recovery algorithm. Packet loss in a TCP stream can have other reasons than congestion. In order to improve performance in case of non-congestion packet loss, TCP Reno [134] introduces the *Fast Recovery* algorithm. Fast Recovery uses duplicated acknowledgments: When the third duplicate ACK is received, the TCP Reno sender enters the Fast Recovery state: The sender reduces the SS-threshold by half of the current cwnd and retransmits the missing segment. After that, the sender sets the cwnd to SS-threshold plus three segments (i.e., one segment per duplicate ACK). The sender increases the cwnd by one segment upon reception of each further duplicate ACK that arrives after the fast retransmission. Thus, further data can be send even in the Fast Recovery phase. When an ACK arrives that confirms all outstanding data, Fast Recovery is terminated by setting the cwnd to SS-threshold, and the sender enters the CA phase again.

TCP NewReno [116] extends the Fast Recovery algorithm of TCP Reno, in case more than one packet is lost in the same window. TCP NewReno introduces a Fast Retransmission interval, which allows the sender to retransmit several lost segments in the Fast Retransmission phase, whereas the Fast Retransmission in TCP Reno restricts the retransmission to single packet loss.

TCP selective acknowledgments. TCP may experience poor performance when multiple segments are lost from one window of data. The cumulative acknowledgments of TCP provide only limited information about packet loss. With the cumulative acknowledgment scheme, a TCP sender can only learn about a single lost packet per round trip time. An aggressive sender may choose to retransmit packets early, but retransmitted segments beyond the cumulative acknowledgment number may have already been successfully received.

The selective acknowledgment (SACK) option [90] overcomes this limitation by reporting blocks of successfully received segments beyond the cumulative acknowledgment. Implicitly, this reporting scheme includes the packets not received. Consequently, the sender can retransmit only the missing packets.

The SACK extension uses two TCP options. The first option (i.e., termed `SACK-permitted`) may be sent in a `SYN` segment to indicate that the sender is capable to use the SACK extension. The second option is the SACK option itself. In case SACK is permitted, the receiver appends the SACK option to acknowledgments when it received non-contiguous segments. The SACK option specifies the left and right edge of received segment blocks. Since the TCP options are restricted to 40 bytes, the receiver may at most append four SACK blocks to a single acknowledgment. In case the receiver detects more non-contiguous segments, it appends the first segments in the flow.

2.2.4 Mobile IPv6

Mobile IPv6 [33] provides global mobility management for portable IPv6 devices, without any modifications to *non-mobile* hosts and routers in the Internet. The protocol intends to enable nodes to conveniently roam between different IP subnetworks, independent of the access technology. Mobile IPv6 introduces the following entities:

- Mobile Node (MN): Any non-stationary host in the network (e.g., notebook, PDA or mobile phone).
- Home Agent (HA): A router/proxy in the MN's *home network*, which keeps track of the locations of MNs that belong to this home network. In absence of a MN, the HA intercepts packets destined for the MN and tunnels the data to the actual MN's location.
- Correspondent Node (CN): Any host in the Internet, which communicates with a MN.
- Access Router (AR): A router that offers network connectivity and forwarding services to a MN in a *foreign network*.

As indicated above, Mobile IPv6 distinguishes between the *home network* and *foreign networks* of a MN. The *home network* represents the network where CNs expect the MN to be according to its permanent IP address. When a MN roams, it visits *foreign networks*. According to this distinction, to provide an unique, permanent identifier for reachability (e.g., TCP uses the IP address as part of the connection identifier) and to provide a temporal identifier for connectivity in a foreign network, each MN has two IP addresses:

Home address The home address is the MN's IP address in its home network, which is used as unique identifier. This address remains unchanged when the MN is roaming into foreign networks.

Care-of address The care-of address (CoA) is the topologically correct address in a visited foreign network, which provides connectivity via the AR. The CoA is assigned in addition to the home address.

According to its current location, the MN can be considered at home or attached to a foreign network. When the MN is at home, it is connected to the HA's sub-network and thus, it is reachable via its home address (i.e., like any stationary host). When a MN connects to a foreign network, it creates a CoA according to the prefix of the Router Advertisement, which announces the presence of the foreign network. In order to inform its HA about the current point of attachment, the MN sends a Binding Update (BU) message to its HA. This BU contains the current CoA of the MN.

When only the HA is aware of the current location of a MN, the HA must route/tunnel all data packets for the MN. In order to optimize the routing and reduce the load of the HA as a bottleneck, Mobile IPv6 provides a *route optimization* option. When route optimization is enabled, the MN sends a BU also to active CNs (i.e., CNs in the binding cache of the MN's ongoing communications). In this case, the respective CNs must include the Mobile IPv6 stack, whereas standard Mobile IPv6 without route optimization is transparent to fixed Internet nodes. CNs can address packets directly to the CoA of the MN and avoid triangle routing via the HA. In case a CN intends to contact a MN, only the first packet is routed via the HA. Upon reception of a tunneled packet, the MN sends a BU update to the respective CN, allowing direct communication via its current home address.

The process when a CN contacts a MN located in a foreign network is schematically illustrated in Figure 2.3.

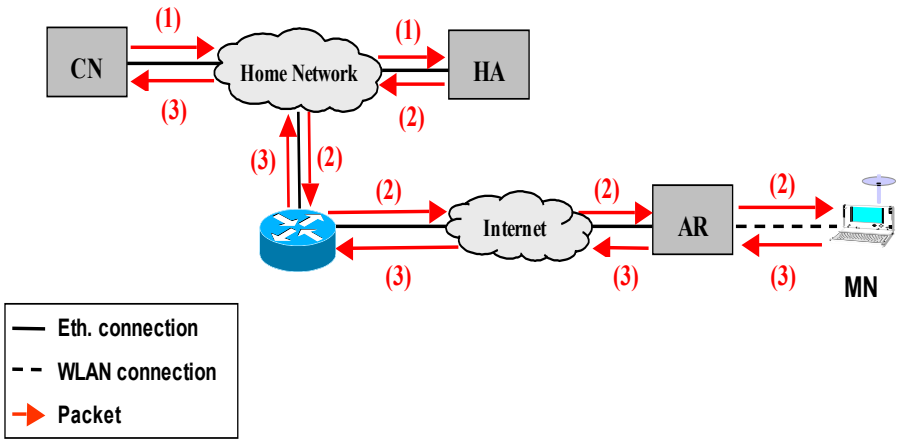

Figure 2.3: Mobile IPv6: Communication initialization to a roaming MN.

Since the CN is not aware of the actual location of the MN, it addresses the packet to the MN's home address (1). The HA intercepts the packet for the MN as proxy, encapsulates it and tunnels the packet to the current location of the MN

(2). The MN replies directly to the CN (3). The MN may include a BU in packet (2) if the MN has the route optimization option enabled. When the CN implements a Mobile IPv6 stack and also has the route optimization option enabled, it may address further packets directly to the CoA of the MN. When the MN changes its point of attachment (i.e., the AR), it updates its new CoA to its HA and the CN in separate BUs.

The process when a MN changes its point of attachment is termed *handover*. This means that the routing responsibility for the MN changes from one AR to another. Though Mobile IPv6 provides support for ongoing connection during a handover (i.e., home address as static identifier), there is a significant interruption in a Mobile IPv6 handover. The provisioning of Fast Handovers is beyond the scope of Mobile IPv6. The respective IETF working group tried to accelerate the standardization by focusing on ongoing connections, leaving Fast Handovers for future standardization in a separate, dedicated working group. In Mobile IPv6, the MN first terminates the old connection before establishing the connection to the new access point. This procedure is commonly termed *break-before-make* philosophy. Figure 2.4 illustrates the Mobile IPv6 handover signaling flow.

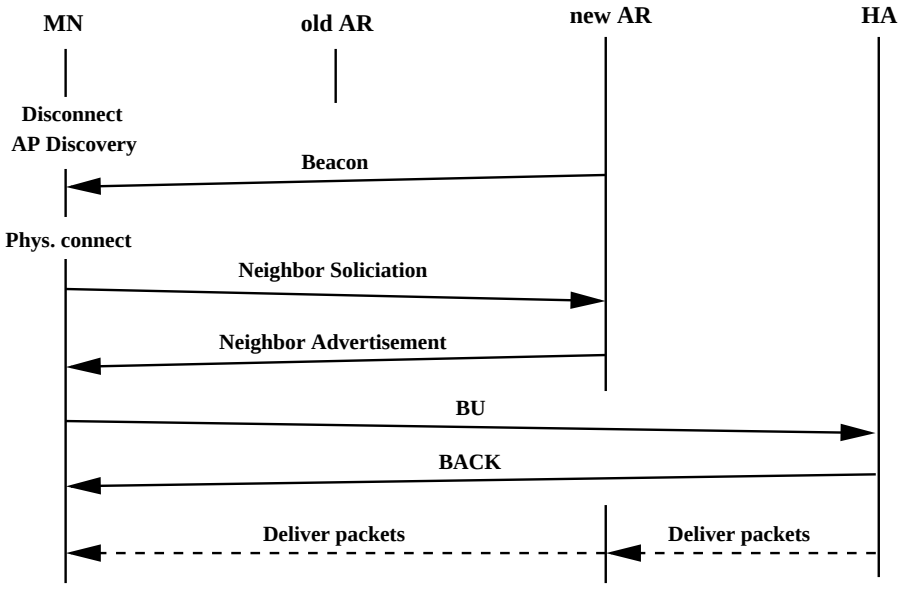

Figure 2.4: Mobile IPv6 handover signaling flow.

First, the MN disconnects from the old AR. Afterwards, the MN searches for new ARs, which announce their presence via beacons. Upon reception of a beacon, the MN physically attaches to the new AP. However, the IPv6 connection cannot resume yet. The MN requests the IPv6 configuration information via a Neighbor Solicitation message. The new AR replies to the solicitation with a Neighbor Advertisement, which contains all relevant IPv6 configuration information. With this information, the MN informs its Home Agent with a Binding Update (BU).

The HA confirms the BU with a Binding Acknowledgment (BACK). Now, the HA can deliver IPv6 packets to the new location of the MN, and the IPv6 connection resumes.

The handover latency, as shown in Figure 2.4, leads to unacceptable interruptions during a handover for applications on top of Mobile IPv6. The provisioning of seamless services to roaming users is not possible. Respective enhancement proposals are subject to current research, as presented in the following section.

2.2.5 Fast Handovers in Mobile IPv6

The main motivation of Fast Handovers in Mobile IPv6 is to provide uninterrupted real-time communication and avoid performance degradation in case of handovers. Therefore, the main goal of the different handover enhancement proposals is to reduce the interruption time, termed *handover latency*, during a handover.

This section surveys the main handover enhancements, as discussed within the IETF, such as Hierarchical Mobile IPv6 Mobility Management (HMIP) [50] or Fast Handovers for Mobile IPv6 (FMIP) [113]. Both approaches localize the signaling of the handover to specific nodes in the local sub-network. Beyond, both approaches follow the *make-before-break* philosophy: The MN prepares the new connection while still being connected to the old AR.

HMIP introduces a hierarchy of anchor routers that localize the handover procedure: When a MN initiates a handover, it first localizes and contacts the anchor router within the sub-network that connects both the old and the new ARs. This anchor router takes over the routing responsibility for the MN, particularly during the handover: The anchor router forwards all data destined for the MN to its current AR. Beyond, the anchor router may *bicast* the data during the handover, i.e., it may duplicate and transmit all packets to both ARs simultaneously for a restricted time interval. When the MN physically attaches to the new AR, its data packets already arrive. After the handover, the anchor router forwards the data to the new AR only. The handover is transparent to nodes outside the sub-network, such as HA and CNs.

In FMIP, the MN uses the ARs to prepare the handover and the new connection: When a MN initiates a handover, it informs the old AR about the handover and provides the address of the new AR. The old AR contacts the new AR via the fixed network and prepares the connection. During the handover, the old AR bicasts packets destined to the MN to its access network and to the new AR. Figure 2.5 illustrates the Fast Handover signaling flow.

FMIP can be described in three phases:

(i) Handover initiation phase Each AR transmits periodic beacons to announce its presence. The MN detects new, potential ARs by listening to these beacons. The MN may decide to handover to a new AR, e.g., as a result of beacon signal strength measurements or due to QoS offers contained in the beacon. When the MN decides for a handover, it initiates the handover by

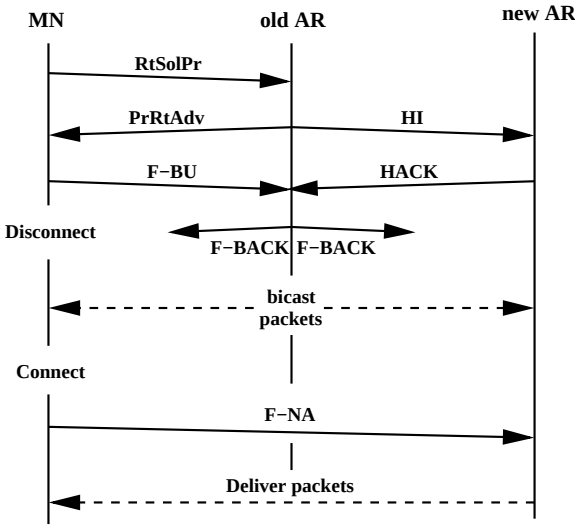

Figure 2.5: Fast handover signaling flow.

the transmission of a Router Solicitation for Proxy (RtSolPr) to the old AR. This message contains the link layer address of the new AR, as acquired in the beacons. In response, the old AR sends a Proxy Router Advertisement (PrRtAdv) including the network prefix of the new AR for address auto-configuration (assuming the new AR is known). In parallel, the old AR informs the new AR about the upcoming handover via the Handover Initiate (HI) message. The old AR replies a Handover Acknowledge (HACK) message. When the MN has all necessary information, it sends a Fast Binding Update (F-BU) to the old AR, which is actually the last message sent before executing the handover, and initiates the second phase of the Fast Handover.

(ii) Tunnel establishment phase When the old AR receives a F-BU message, it establishes a bi-directional tunnel to the new AR. The Fast Binding Acknowledgment (in response to the F-BU) is already duplicated and sent to both potential access points. The tunnel establishment assumes successful exchange of HI and HACK message between the ARs. When the tunnel is established, packets are duplicated in the following bicasting phase.

(iii) Bicasting phase During this phase, all packets destined for the MN are simultaneously transmitted to both ARs. Bicasting is restricted in time to the duration of the handover. When the MN performs the actual physical handover and attaches to the new AR, packets of its ongoing communications already arrive at the new point of attachment. The MN sends a Fast Neighbor Advertisement (F-NA) in order to announce its presence and initiate the forwarding of bicasted packets on the link. Afterwards, the MN can update the location information with its HA (and CN - in case the route optimization option is active).

In the simulative studies [108, 131], the handover latency in the HMIP and FMIP approach is evaluated and compared, considering also the costs in terms of additional network complexity of both approaches. Based on these results, the FMIP approach has been selected in the Moby Dick project and for the following evaluation in this thesis because: *(i)* The handover latency in both approaches is similar. *(ii)* FMIP does not require changes to nodes other than the MN and the involved ARs. Thus, the additional network complexity in FMIP is less than the hierarchical network structure of HMIP. *(iii)* The inter-access router communication provides the framework for the integration of QoS and AAAC. The required QoS and AAAC messages between the ARs can be attached to FHO signaling messages.

The following sections describe the adaptation of FMIP in an overall mobility architecture, comprising AAAC and QoS. In this framework, the evaluation in this thesis provides handover latency measurements and studies the network performance for UDP and TCP data streams.

2.3 Related Work

In [39] the authors categorize the related work in three main categories: Link layer, end-to-end and split-connection protocols.

The *link layer protocols* aim at hiding the high bit error rate (BER) of the wireless medium from the transport-layer. An enhanced link layer autonomously recovers packet losses by retransmissions without affecting the upper layers. One example is the Snoop Protocol [8] which installs a *snoop agent* on every access point. This agent keeps track of the TCP packets sent from the stationary host that have not been acknowledged by the MN. Whenever a packet loss is detected (e.g., via duplicate acknowledgments), the agent checks its cache and retransmits the packet autonomously. Though the TCP performance is increased by the artificial improvement of the link quality, the sniffing of packets raises additional security concerns and requires large, efficient caches on the access points.

The *end-to-end protocol* approaches provide modifications or extensions to TCP which handle losses in a way that improves the performance compared to regular TCP. Therefore, respective proposals maintain the end-to-end semantics of TCP. These approaches comprise e.g., the Reno, NewReno, SACK and Fast Retransmission options of TCP as well as the Explicit Bad State Notification (EBSN) mechanism of the network. In [22] the authors combine TCP and the handover algorithm by introducing artificial acknowledgments after handover completion in order to avoid unused connected times after re-connection due to TCP back-off timers.

Split connections separate the wired and the wireless connection in order to isolate wireless and mobility related problems from the fixed network. The access router or basestation splits a connection between a CN in the fixed network and a MN in a wireless network. Examples include the I-TCP [5] protocol, which intro-

duces an agent on the AR in order to maintain both connections, or the M-TCP [20] protocol, which organizes the network in a hierarchical architecture. However, split connections violate the end-to-end semantics of TCP. When the sender receives acknowledgments from an intermediate entity, it believes the packets to be successfully delivered, whereas the packets may not have reached the final destination via the wireless part of the network. Actually, positively acknowledged packets may never reach the receiver, e.g., when the wireless connectivity is permanently disrupted.

The impact on handover latency on the performance of TCP has been evaluated in previous work, e.g., [41]. However, this case study is based on handover latency times of three to four seconds and therefore, the slow-start algorithm is invoked, which negatively impacts the TCP performance.

Other research projects use different IP mobility management schemes. The IST project WINE GLASS [21, 119] uses Mobile IPv6 to handle IP mobility but does not implement any solution to optimize local mobility at the IP layer. The efficiency in local mobility purely depends on MAC layer technologies and thus, mobility is restricted within the same IP subnet. The IST projects BRAIN and MIND [137] proposed their own local mobility management solution: BCMP (Brain Candidate Mobility Protocol). This protocol combines properties of the IETF hierarchical solutions (like HMIP) and the IETF Fast Handover solution. For this purpose, BCMP introduces special components, such as Anchor Points and Access Routers. Anchor points are special routers that provide addresses to visiting MNs in a set of IP subnets and tunnel packets to the MNs. Access Routers provide the physical access and terminate the tunnel from the Anchor Point.

2.4 Mobility Architecture and Implementation Details

The main objective of the Mobility Architecture, as developed in the Moby Dick project, is to evolve 3^{rd} Generation mobile and wireless infrastructure towards the Internet in order to provide uninterrupted, interactive and distributed multimedia services to roaming mobile users. The overall approach is independent of the deployed access technology. Therefore, the Moby Dick test network comprises TD-CDMA, IEEE 802.11b Wireless LAN and Ethernet as example access technologies for verification, validation and demonstration. The Fast Handover design and implementation mainly follows the IETF Draft *Fast Handovers for Mobile IPv6* [47], as explained in section 2.2.5. Note that in the mean time the respective work within the IETF advanced to RFC-status in [113]. However, the complexity of the overall system implemented by different project partners requires the determination of versions at an early stage. Therefore, the implementation is based on the draft version, as cited above. However, the results, as presented in the following Section 2.5, are generally valid since the actual interruption time is not affected when adapting to more recent Fast Handover for Mobile IPv6 versions.

The Moby Dick approach extends the basic IETF Fast Handover signaling flow, mainly by the adoption of the AAAC and QoS messages, as shown in Figure 2.6. On the reception of a Router Advertisement (1) from a new AR, the MN evaluates the signal strength of this message. In case the signal strength exceeds the signal strength of the current connection by a threshold, the MN initiates the Fast Handover (FHO) process. The MN issues a Router Solicitation for Proxy (RtSolPr) message (2) to the old AR. The old AR sends the Handover Initiation (HI) message to the new AR (3) and waits for the respective Handover Initiate Acknowledgment in message (4). This FHO inter-AR communication is used to transfer AAAC context in an intra-domain handover (i.e., assuming an existing security association between the ARs). The AAAC attendant on the old AR piggybacks the AAAC context on top of the ICMP FHO messages (3) and (4) in order to relay all necessary AAAC information to the AAAC attendant on the new AR.

Simultaneous to the HI message, the old AR informs the old QoS broker about the upcoming Fast Handover via message (A). The old QoS broker sends message (B) to the new QoS broker, which requests the previously used QoS context from the new QoS broker. The new QoS broker reports the available QoS parameters directly to the new AR in message (C).

Upon the arrival of a positive Handover Initiate ACK message, the old AR informs the MN about the successful preparation via a Proxy Router Advertisement (PrRtAdv) in message (5). In reply, the MN sends the Fast Handover Execute (FHE) message (6) to the old AR, which initiates the establishment of the bicasting in (7). When the bicasting is established and confirmed via the Fast Binding ACK in message (8), the MN executes the physical handover (9) and announces its presence to the new AR with the Neighbor Advertisement message (12). Upon the expiration of the bicasting timer in (10), the old AR sends the respective accounting data of the MN to the AAAC server. Likewise, the new AR reports the reception of the Neighbor Advertisement message to the AAAC server for accounting purpose in message (X). Finally, the MN updates its binding with the Mobile IPv6 Home Agent via the Binding Update / Acknowledgment in messages (13) and (14).

The potential delay, which could be introduced in the communication with the QoS broker (i.e., messages (A),(B) and (C)), does not impact the handover latency since the message exchange takes place during the handover preparation phase. However, this parameter affects the radio cell planning (i.e., size of overlapping radio coverage areas). The overall handover time linearly depends on the QoS answer time. The AAAC context is acquired locally on the old AR, transferred in the Handover Initiate message and locally relayed to the AAAC attendant on the new AR. Since there is no impact of AAAC and QoS delay or jitter on the handover latency, the following experiments assume *idealized* AAAC and QoS modules with minimal processing time in the order of 1 ms.

The Fast Handover algorithm, as described above, is implemented as a kernel module in the Linux kernel 2.4.16. The module extends the basic MIPL Mobile

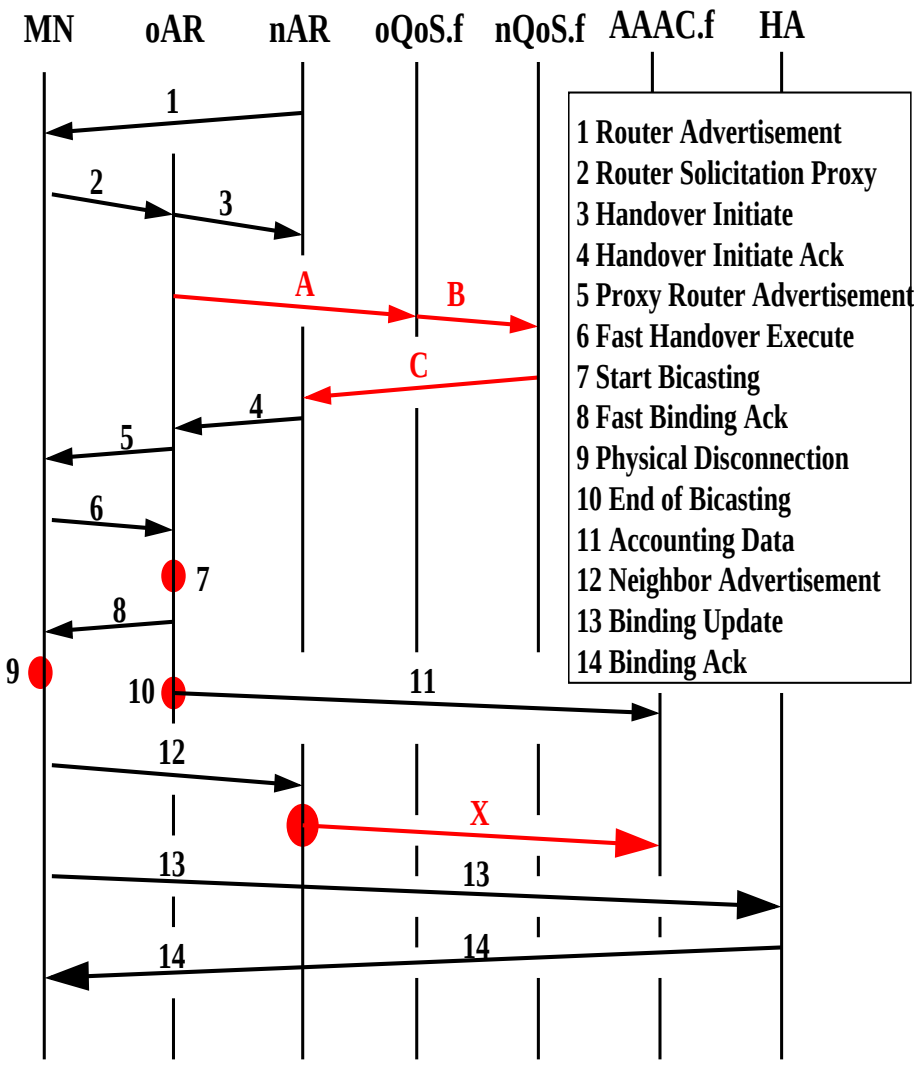

Figure 2.6: Fast handover signaling flow including QoS and AAAC.

IPv6 implementation of the Helsinki University of Technology (HUT) [54]. The FHO signaling is implemented via ICMPv6 messages.

Beyond the FHO functionality, the module manages the kernel-userspace interfaces (i.e., Linux *character device* pointer management to user- and kernel-space functions) to the AAAC and QoS attendants on the ARs, as well as the interface to the Mobile Terminal Network Manager (MTNM) on the MN. The latter component is responsible for the movement detection. The movement detection scheme is still based on Router Advertisements, but it is enhanced to be *network aware*: Router Advertisements from all surrounding Access Routers are captured and relayed to the MTNM. The MTNM evaluates the signal strength of the Router Advertisements, and only when the signal strength of a potential candidate exceeds

a pre-defined threshold for a certain duration of time, the Router Advertisement triggers the Fast Handover procedure. This signal strength measurement supports real movement of a mobile device. However, the presented measurement results are based on manually determined signal strength modifications in order to provide automated, precise, uniform and comparable results.

The majority of Moby Dick partners, particularly network operators, assume that future (i.e., beyond the current third generation - 3G) communication networks will deploy Wireless LAN in infrastructure mode. In contrast to this belief, the Moby Dick mobility architecture employs the 802.11b ad-hoc mode because of hardware and software restrictions in the current Wireless LAN systems: The physical and MAC layer handover in infrastructure mode comprises a mandatory frequency scanning and a physical attachment to all available cells. The additional interruption due to the Wireless LAN infrastructure mode increases the handover latency to at least 150 ms. This handover latency is not acceptable for real-time services. Therefore, the project decided for the ad-hoc mode, being aware that filtering of traffic to different access points is required. The implementation comprises an additional filtering mechanism in the Wireless LAN driver which extracts router advertisements for signal quality evaluation and movement detection. Otherwise, only (data) traffic from the current point of attachment is allowed to pass the filter.

One of our contributions is the coupling of FHO, AAAC and QoS. This includes for example the combination of FHO and AAAC messages, and the scheduling of QoS message exchange in the FHO sequence.

In order to simulate individual traffic and to evaluate the system under varying conditions, a fully IPv6 capable traffic generator has been developed and implemented. This traffic generator allows transmitting data via UDP or TCP transport protocols. It enables the configuration of data flow parameters, such as packet size, inter-packet delay or total amount of data to be transmitted. This traffic generator facilitates reproducible scenarios for the evaluation of network performance in the following FHO analysis.

2.5 Performance Evaluation - Experimental Results

This section presents the experimental evaluation results of the Fast Handover latency measurements and of the TCP and UDP performance in presence of Fast Handover.

2.5.1 Studied Scenarios and Measurement Setup

The experimental FHO performance evaluation is based on measurements in test networks, comprising Linux implementations of all modules, such as FHO kernel module, AR enhancements, movement detection via signal strength measurements,

AAAC and QoS components. Each Moby Dick partner maintains a test network, e.g., for the development, implementation and testing of a specific component. The frequent integration and exchange of modules allows operating each network autonomously. The different trial sites can also be interconnected via IPv6 over IPv4 tunnels in the Internet in order to test interoperability or to evaluate the system over the unpredictable long distance links of the Internet.

Figure 2.7 illustrates the test network used in this thesis and at NEC. The majority of experiments are conducted in this network. A set of handover latency measurements is conducted at the Moby Dick trial site at the University Carlos III Madrid (UC3M). The UC3M network set up is similar and, therefore, it is not illustrated separately. The respective sections indicate results of the measurements at UC3M. The combination of measurements and results of different partners shows the strong interaction and cooperation in the Moby Dick consortium.

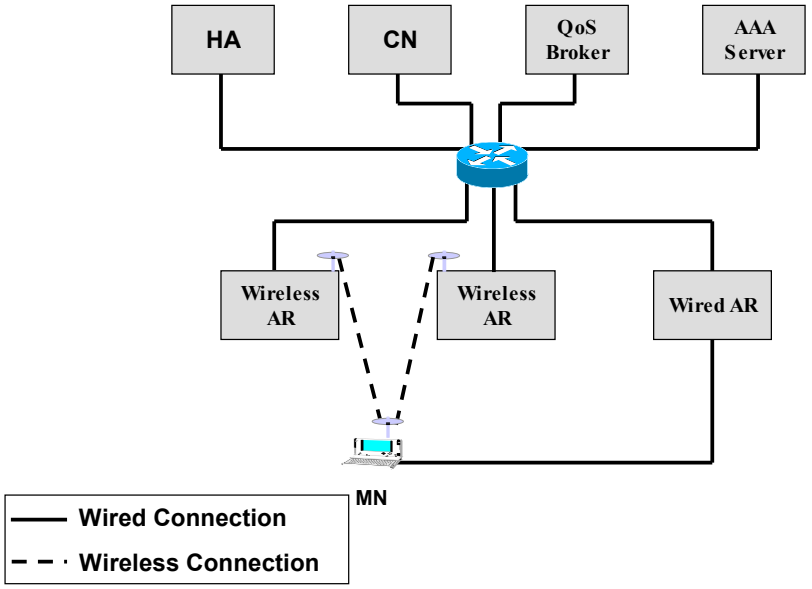

Figure 2.7: Mobility enabled IPv6 Testbed including AAAC and QoS.

All nodes in Figure 2.7 use the Linux operating system with kernel 2.4.16 and the MIPL Mobile IPv6 stack in the kernel. The test network consists of the following components.

The HA is the mobility proxy of the MN in the home network. The CN represents an arbitrary node in the Internet. For the evaluation, the CN is one of the communication end-points. AAAC server and QoS broker implement the required AAAC and QoS databases and functionalities, respectively. The core router separates different IP sub-networks, such as the home network and foreign networks. Each foreign network consists of a separate IP sub-network. The core router provides the main routing functionality, e.g., it routes the data stream for the mea-

surements to the respective foreign sub-network. The network includes three ARs: Two IEEE 802.11b wireless ARs and one wired Ethernet AR. All ARs include the FHO module, AAAC- and QoS attendants. The MN is a laptop, which is equipped with an on-board Ethernet and an IEEE 802.11b wireless LAN PCMCIA card. The MN roams between the foreign networks, using its FHO module, while communicating with the CN. The MN performs intra-technology handover, staying in the same technology and using the same wireless LAN interface, and inter-technology handover which uses the different interfaces wireless LAN and Ethernet.

In this test network, the evaluation measures the handover latency, UDP and TCP performance in presence of standard Mobile IPv6 and Fast Handovers. All of the following scenarios evaluate intra-domain handovers, i.e., handovers within the same administrative domain, which assume a security association between the ARs. Furthermore, the scenarios use ideal QoS and AAAC entities since *(i)* QoS signaling takes place prior and after the actual handover, and AAAC signaling is coupled and combined with the FHO signaling, as explained before. Thus, this assumption does not impact on the performance evaluation. The evaluation focuses on WLAN-WLAN intra-technology, as well as WLAN-Ethernet inter-technology handovers. The remainder of this section describes the evaluation scenarios in detail.

Handover Latency Measurement Scenario This scenario measures the handover latency of *standard* Mobile IPv6 handovers and Fast Handovers in the test network, as described above. The handover latency is measured in two different ways: *(i)* The measurement of packet loss in a data stream with predefined inter-packet delay allows the calculation of the handover latency. However, the granularity of this method is restricted by the inter-packet delay. The measurements use the ping6 tool with an inter-packet delay of 10 ms to 20 ms, which determines the measurement granularity. *(ii)* The handover latency is measured by time-stamps in the source code of the FHO module. This method increases the accuracy of the measurement.

The packet loss measurements are conducted in the UC3M test network and the time-stamp measurements is carried out in the NEC test network. Beyond, the focus of the measurements differs, as follows:

(i) The packet loss measurements focus on the impact of different delays between MN and CN or MN and HA on Mobile IPv6 and FHO handovers. The different delays emulate different locations of the respective nodes, with varying number of hops between the nodes. In the test network, the measurements use the NISTNET [100] tool to emulate different delays. NISTNET allows a single Linux PC, set up as a router, to emulate a wide variety of network conditions, such as e.g., delay, jitter or packet loss. Since NISTNET

supports IPv4 only, the measurements use an IPv6-in-IPv4 tunnel between the ARs and the CN. This tunnel does not influence the measurements since the encapsulation time is negligibly small and constant for all packets.
Further measurements evaluate the impact of the Router Advertisement interval on the handover latency. As explained before, the FHO movement detection scheme utilizes Router Advertisements, similar to standard Mobile IP handovers. In contrast to standard Mobile IP handovers, FHO evaluates the signal strength of the respective Router Advertisement during the handover preparation phase. The measurements use different Router Advertisement intervals of 0.5 - 1.5 s and 2.0 - 4.0 s. The results compare the impact of these different intervals on standard Mobile IPv6 and FHO handovers.

(ii) The focus of the second measurement campaign in the NEC test network is on preciseness. Therefore, time-stamps are added to the FHO source code on the MN. These time-stamps measure the interruption time of the IPv6 connection. The measurement starts when the the MN leaves the old AR and ends when the MN receives the first data packet via the new AR. This method represents a very accurate granularity because the precision of the operation relies strictly on the operating system time (i.e., CPU TSC-Timestamp Counter Register), and the measurement follows immediately the respective FHO primitives.

UDP Performance Measurement Scenario The UDP performance evaluation includes qualitative and quantitative aspects. The qualitative performance measurements evaluate the user's satisfaction when watching a video on a mobile device, while performing frequent handovers. The user compares the perceived quality under standard Mobile IPv6 and Fast Handovers. However, qualitative perception of the aural and visual human abilities is different for every individual person. In contrast, the quantitative evaluation measures the packet loss in case of handovers. These quantitative measurements provide comparable and reproducible results for a meaningful evaluation. The remainder of this paragraph describes both, the qualitative and the quantitative UDP performance measurement scenarios.

In the qualitative UDP performance evaluation, a video trailer is shown to a group of users while the MN frequently executes handovers. The CN transmits the video and audio stream to the MN, using the Linux tool VideoLAN [133]. The buffering capability of the tool is disabled. The CN transmits a video trailer of about 2 min and 20 s length at a data rate of approximately 468 kBit/s. The tool is configured to use uncompressed, full frames since the codec is beyond the scope of this evaluation. The MN executes handovers every 40 s. The users evaluate the performance, i.e., the perceived noticeable interruptions, of Mobile IPv6 and Fast handovers in a questionnaire.

The quantitative evaluation uses the self-implemented traffic generator, as described in Section 2.4. The CN transmits continuously UDP data in packets of 1000 bytes and an inter-packet delay of 25 ms to the MN. Again, the MN executes handovers every 40 s. Tcpdump [129] captures the traffic on the sender and the receiver side. Furthermore, Tcptrace [117], which is recommended by the IETF [118], has been adapted to process basic IPv6 functionalities. The tool analyzes quantitatively the traffic and retrieves the information for the generation of the graphs, as presented in section 2.5.3.

TCP Performance Measurement Scenario The TCP performance evaluation measures quantitatively the TCP throughput in case of handovers. This includes the evaluation of the TCP state when the TCP sender detects packet loss due to handover. Typically, TCP assumes network congestion in case of packet loss and invokes its congestion control mechanisms, such as Fast Retransmit or Slow Start. The reaction of TCP, which depends on the number of lost segments, determines the throughput.

In this scenario, the CN establishes a TCP connection to the MN, using once more the self-implemented traffic generator tool. The evaluation uses TCP NewReno, which includes the Fast Retransmit and Selective Acknowledgment options. In presence of the established TCP connection, the MN executes handovers. Lost segments invoke the TCP congestion control. The evaluation traces the TCP state and measures the resulting TCP performance.

2.5.2 Handover Latency Measurement Results

This section presents the results of the handover latency measurements according to the scenarios, as specified in the previous section. Section 2.5.2.1 measures the number of lost packets in a pre-defined data stream which leads to the derived results of the handover latency by multiplying the number of lost packets with inter-packet delay. These measurements include the evaluation of the impact of round trip time variations and different router advertisement intervals for Mobile IPv6 and Fast Handover. Section 2.5.2.2 presents the handover latency measurement results for time-stamps, added to the FHO module.

2.5.2.1 Packet Loss Measurement for Varying Network Conditions

In this set of measurements, the handover latency derives from the measured packet loss in a pre-defined data stream. The sender transmits ping6 packets every 15 ms. This inter-packet delay represents an upper bound for the accuracy: In case a packet is lost, the actual handover latency is between zero and two times the inter-packet delay (i.e., 30 ms). The interruption might be short and connectivity might resume immediately after the potential arrival of the lost packet, but is realized only with the reception of the subsequent packet.

Figure 2.8 shows the handover latency of WLAN-WLAN intra-technology handover for both standard Mobile IPv6 and the Fast Handover implementation over the network delay between the CN and MN. Note that NISTNET emulates the network delay in both directions. The two lines in the graph represent the minimum and maximum border of the latency, according to the inter-packet delay.

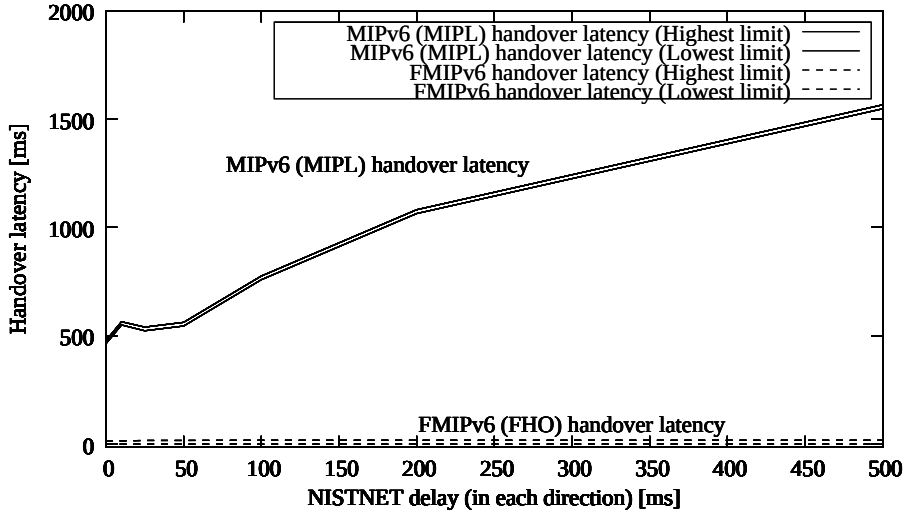

Figure 2.8: Mobile IPv6 and fast handover latency via packet loss measurement versus network delay.

Figure 2.8 shows that Mobile IPv6 handover delay linearly increases with increasing network delay. The Mobile IP handover latency is directly proportional to the round-trip time required for a Binding Update to reach the CN.

In contrast, the Fast Handover latency is independent of the network delay. The Fast Handover latency remains constant because the old AR forwards (i.e., bicasts) the data to the new AR until the MN updates its new location to its HA and CNs. Therefore, the handover latency is completely independent of location and network delays to other nodes.

Figure 2.9 shows the handover latency measurement results for different Router Advertisement intervals in scenarios with and without emulated network delay increase. Two different intervals are chosen: *(i)* The minimum permitted interval in the Mobile IPv6 Draft, i.e., MinRtrAdvInterval: 0.5 s and MaxRtrAdvInterval: 1.5 s. *(ii)* The Router Advertisement interval is increased to MinRtrAdvInterval: 2.0 s and MaxRtrAdvInterval: 4.0 s. Note that these values are lower than the recommended ones in the *Neighbor Discovery for IPv6* RFC 2461 [126], according to the modifications proposed in the Mobile IPv6 Draft version 15. For each of these values two experiments are conducted, one with and one without adding an emulated network delay of 500 ms. A network delay of 500 ms emulates a situation of a very long distance between the CN and MN.

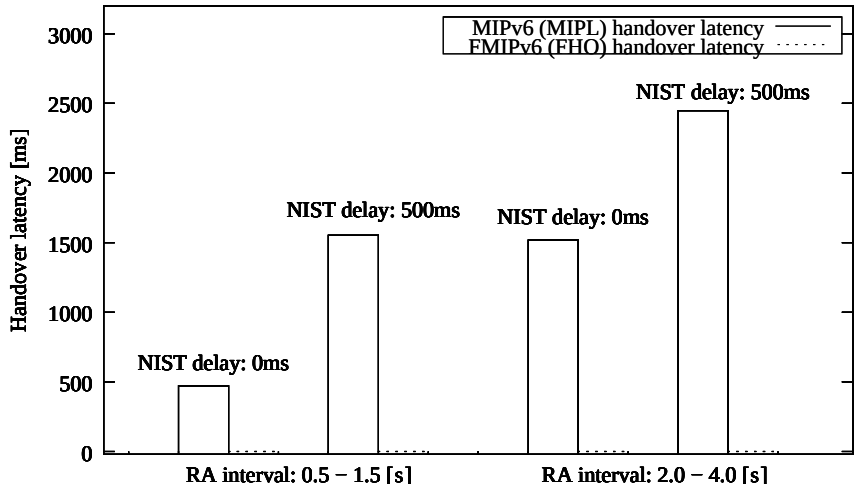

Figure 2.9: Mobile IPv6 and fast handover latency via packet loss measurement versus router advertisement interval.

The Router Advertisement interval has a significant impact on the latency of a standard Mobile IPv6 handover because the detection of the new AR via Router Advertisements takes place when the MN is already disconnected from the old AR. While sending more frequent Router Advertisements would reduce the Mobile IPv6 handover latency, the bandwidth consumption of unsolicited Router Advertisements increases. Particularly in access technologies with scarce bandwidth, such as Wireless LANs, short Router Advertisement intervals are not an option.

The Fast Handover solution is independent of the interval between Router Advertisements. The MN scans and discovers new, potential ARs, while still being attached to the old AR.

2.5.2.2 Time-Stamp Measurements

This section presents the handover latency measurement results via time-stamps in the Fast Handover module. Thus, the evaluation focuses on FHO only. Intentionally, the comparison to standard Mobile IPv6 is omitted because the previous results already show that standard Mobile IP handover latency is not suitable for multi media traffic. Actually, the provisioning of small handover latency is beyond the scope of the Mobile IPv6 definition.

Figure 2.10 depicts the statistical distribution of WLAN-WLAN intra-technology Fast Handover latency for 100 experiments. The mean handover latency in this scenario is 0.23 ms. However, the graph shows a deviation of this mean value, similar to a normal distribution.

In case of WLAN-Ethernet inter-technology Fast Handover, there is no interruption. The availability of two different interfaces allows the set up of a new connection via the new device while it is still possible to communicate via the previous

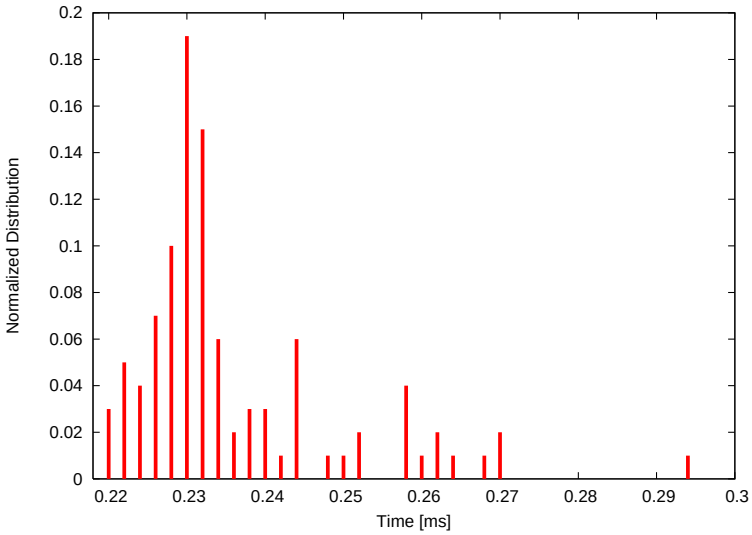

Figure 2.10: Fast handover latency measurement via time-stamps.

device. This process is explained in detail in the following UDP measurement section where packets simultaneously arrive on both interfaces (see Figure 2.11).

Finally, the evaluation measures the overall Fast Handover completion time, which depends on several parameters, such as round trip time, network load, processing time of QoS and AAAC components, as well as access technology. The measurements result in 8 ms and 26 ms overall Fast Handover time for intra- and inter-technology handover, respectively. The only parameter affecting these results within the measurement set-up is the deployed access technology since the QoS and AAAC components have ideal properties and the network is not loaded with other traffic. The time consumption of AAAC is negligible since it involves local processing only on the old AR and new AR, while the processing and RTT of QoS is below 1 ms. The increase of these values by real components would influence linearly the overall Fast Handover time, but not the handover latency.

2.5.3 UDP Measurement Results

For the following measurements, the traffic generator creates a UDP data stream with an inter-packet delay of 5 ms and 500 bytes packet size. Particularly, the small inter-packet delay generates a network load that is similar to the demands of real applications. However, the results are transferable because the high load within this scenario places even more demands on the system, e.g., the loss probability within a small Fast Handover latency increases for decreasing inter-packet delay.

2.5.3.1 Inter-Technology Handover

Figure 2.11 illustrates an Ethernet-WLAN inter-technology Fast Handover process in the presence of a UDP data stream, as observed by the MN as the receiver.

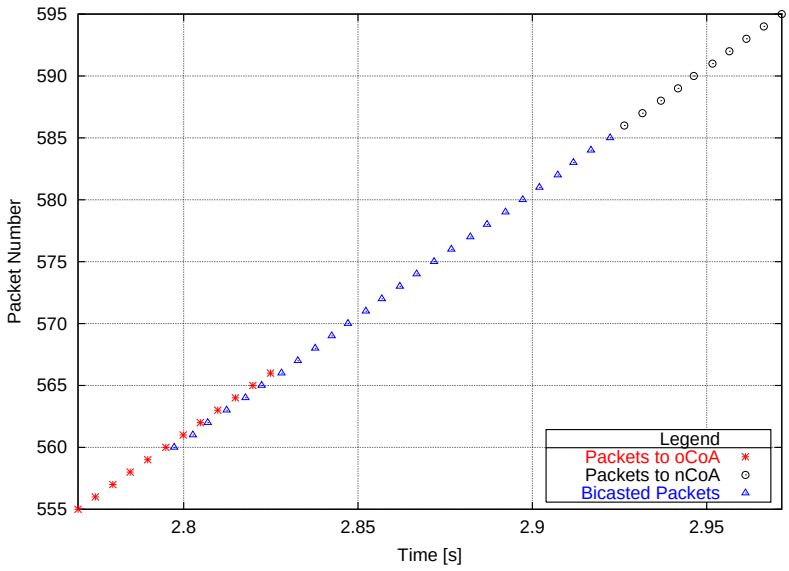

Figure 2.11: Real-time UDP traffic in the presence of Ethernet-WLAN handover, observed by the receiver (MN).

As mentioned before, the MN is able to communicate simultaneously on the different access interfaces. Therefore, the MN receives all packets between 2.8 s and 2.825 s twice: It receives the original data via the old AR on the previous access interface and it receives the bicasted data via the new AR on the new access interface. The MN closes the physical connection to the old AR at 2.825 s. Consequently, it receives only bicasted packets after this point in time. In parallel, the MN updates its HA about the new location with a BU message. The HA updates its binding cache for the MN and changes its routing and tunneling entries, respectively. With the arrival of the BACK, the HA has changed the location entry of the MN. The BACK arrives at 2.925 s, and all consecutive packets are directly addressed to the new point of attachment.

Summarizing, the inter-technology Fast Handover is uninterrupted because the different access interfaces allow simultaneous communication on both interfaces. In this case, the MN receives all packets twice during the handover. The experiment results show that duplicates occur for about 20 ms in average.

2.5.3.2 Intra-Technology Handover

Figure 2.12 shows a WLAN-WLAN intra-technology Fast Handover process in the presence of a UDP data stream, again, as observed by the MN as receiver.

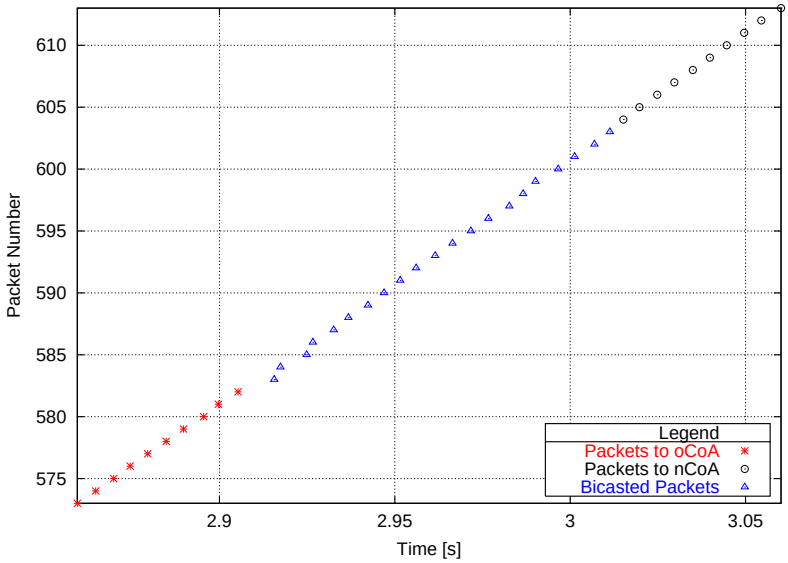

Figure 2.12: Real-time UDP traffic in the presence of WLAN-WLAN handover, observed by the receiver (MN).

Before the handover, the MN receives packets via the old AR. At 2.91 s, the MN performs the physical handover. This physical handover results is an interruption, as evaluated in Section 2.5.2.2. Figure 2.12 shows exemplary the average handover latency of 0.23 ms. After the physical handover, the bicasting mechanisms fills the gap until the HA is updated about the new location of the MN. The MN receives bicasted packets via the new AR, which already arrive at the new AR when the MN physically attaches. With the arrival of the BACK at 3.2 s, the HA has updated its location information of the MN, and packets are directly addressed to the new CoA.

Summarizing, the average handover latency during a WLAN-WLAN intra-technology Fast Handover is 0.23 ms. The bicasting mechanism provides continuous packet delivery until the location and forwarding information at the HA is updated. Consequently, there is no noticeable interruption in real-time audio or video data streams during an intra-technology Fast Handover.

2.5.4 TCP Measurement Results

The following evaluation focuses on TCP performance during a Fast Handover. The traffic generator creates a TCP stream between the CN and the MN. Except for the different transport protocol, the traffic generator uses the same settings as before.

2.5.4.1 Inter-Technology Handover

Figure 2.13 illustrates an Ethernet-WLAN inter-technology handover in the presence of a TCP connection. The plot shows the segments transmitted by the CN in order to analyze the TCP state and congestion control reactions at the sender.

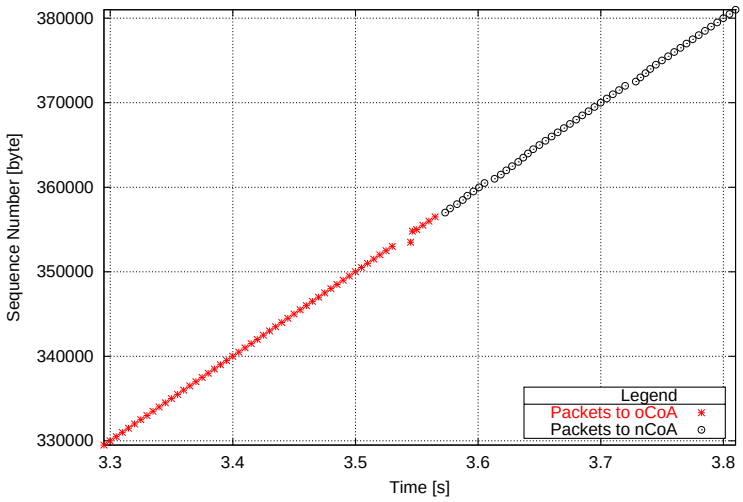

Figure 2.13: TCP connection in the presence of Ethernet-WLAN handover, observed by the sender (CN).

As before, the different physical access interfaces allow simultaneous communication on both interfaces. However, Figure 2.13 does not show duplicate packets because TCP filters and drops duplicates. Therefore, the TCP stream continues during the inter-technology Fast Handover without interruption.

The graph shows a small increase of the inter-packet delay at 3.55 ms. However, this glitch does not originate from handover interruption, but due to a delayed acknowledgment from the MN caused by IP layer re-configuration and duplicate ACK-packet handling processing time by the TCP stack.

Summarizing, the TCP stream is not interrupted by an Ethernet-WLAN inter-technology Fast Handover. TCP filters and drops duplicate packets. The TCP sender does not change the TCP state, and it does not invoke congestion control because there is no packet loss that indicates congestion to the TCP sender. Consequently, an inter-technology Fast Handover does not affect the performance or throughput of a TCP stream.

2.5.4.2 Intra-Technology Handover

Figure 2.14 shows a WLAN-WLAN intra-technology Fast Handover in the presence of a TCP connection. Again, the graph shows the transmitted packets by the CN to analyze the TCP state and congestion control of the TCP sender.

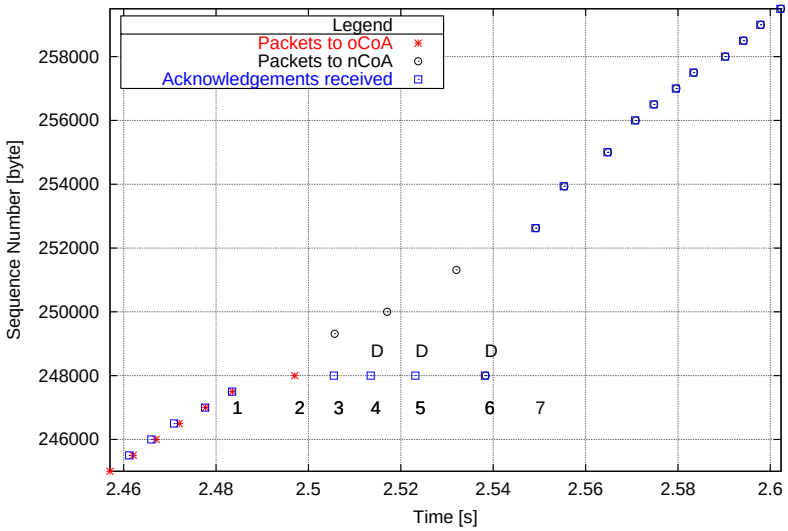

Figure 2.14: TCP connection in the presence of WLAN-WLAN handover, observed by the sender (CN).

During the WLAN-WLAN Fast Handover, as shown in Figure 2.14, one packet is lost. This single packet loss represents the *worst case* of the measurements since in only 2% of the performed measurements one packet is lost. Packet loss of more than one packet never occurs. The following paragraph explains the TCP reaction to the packet loss in detail, referring to the labels in the Figure.

Between label 1 and label 2, the connection to the MN is disrupted, and packet 2 is lost due to the handover. Label 3 marks the acknowledgment for packet 1, which contains a destination option header, updating the CN about the new location of the MN. Afterwards, the CN addresses packets destined for the MN to the new CoA. The acknowledgments 4, 5 and 6 are duplicate acks, which trigger the TCP Fast Recovery algorithm. Thus, the CN re-transmits the lost packet 2, without entering slow start and, thus, without performance degradation. Acknowledgment 7 confirms all data, i.e., including the lost packet and the packets received during the dup-ACK phase. When the sender receives this acknowledgment 7, it exits the TCP Fast Recovery phase and continuous data transmission.

Summarizing, TCP Fast Recovery and Selective Acknowledgments in combination with the Fast Handover scheme avoid *(i)* the invocation of TCP Slow-Start and *(ii)* the redundant retransmission of successfully transferred packets, as shown in the worst case scenario of single packet loss above. Therefore, the Fast Handover approach avoids TCP performance degradation in the presence of handovers.

2.6 Summary and Conclusions

This chapter described the integration of Fast Mobile IPv6 Handovers in an IPv6 mobility environment which integrates AAAC anc QoS in a complete architecture. This architecture has been implemented in this thesis, and in the framework of the EU IST project Moby Dick [94]. The implementation and deployment in test networks facilitate a performance evaluation via measurements in an integrated environment. The evaluation comprises:

- Fast Handover latency measurements, using different methodologies.
- UDP and TCP performance analysis in the presence of Fast Handovers.

The measurement results show that in case of Ethernet-WLAN inter-technology Fast Handover, the handover is not interrupted since simultaneous communication on the different access interfaces is possible. The results show for inter-technology handover a short period of 20 ms on the average where the MN receives duplicates, i.e., the packet via the old AR on the previous interface and the bicasted packet via the new AR on the second access technology. Since most applications detect and suppress duplicates, the UDP connection is seamless in case of Fast inter-technology Handover.

In case of WLAN-WLAN intra-technology Fast Handover the Fast Handover approach provides uninterrupted Ethernet-WLAN inter-technology and in average a handover latency of 0.23 ms for WLAN-WLAN intra-technology handover. The *worst case* result looses a single packet during the handover. Such a single packet loss invokes the TCP Fast Retransmission scheme which retransmits the lost packet without noticeable performance degradation. Particularly, the TCP slow-start algorithm, which significantly decreases the performance, is not triggered. The sender re-transmits the packet, while further packets are continuously transmitted during the Fast Retransmission phase.

Summarizing, the Fast Handover approach provides uninterrupted Ethernet-WLAN inter-technology and in average a handover latency of 0.23 ms for WLAN-WLAN intra-technology handover in an integrated Mobile IPv6 environment, including QoS and AAAC. These handover latency results support uninterrupted real time UDP communications and avoid TCP performance degradation in presence of Fast Handovers. The results provide an important step towards All-IP future networks. Future work will evaluate network controlled FHO (e.g., handover invocation for load balancing), as required by network operators for a successful deployment and network operation.

After the handover latency measurements and TCP / UDP performance evaluation upon handover occurrence in single-hop scenarios, the following chapter moves on to mobile ad hoc networks that facilitate multi-hop communication.

The most prevalent use case for mobile multi-hop communication networks include vehicular ad hoc networks (VANETs).

VANETs represent a special kind of mobile ad hoc networks (MANETs) that facilitate inter-vehicle and vehicle-to-roadside communication without additional or pre-established infrastructure. VANETs facilitate the integration of existing Internet applications into vehicles and enable new applications, e.g., aiming on an increase of safety on the road.

The following two Chapters 3 and 4 focus on transport layer issues in VANETs, such as reliability, for point-to-point and point-to-multipoint applications, respectively.

Chapter 3

Design and Evaluation of a Vehicular Transport Protocol (VTP)

3.1 Introduction

Vehicular ad hoc networks (VANETs) are self-organizing, wireless, multi-hop networks that enable vehicle-to-vehicle and vehicle-to-roadside communication. The main characteristic of VANETs is a high degree of node mobility, resulting in frequent topology changes.

VANETs enable a variety of new applications and facilitate the integration of existing applications into vehicles. These applications pose different requirements on the network and transport layers, such as reliability or in-order delivery of data. The fulfillment of the application requirements in VANETs is challenging due to the unique characteristics of the environment [82]. These environmental challenges for a transport protocol comprise:

- Challenges of the wireless medium, such as high bit error rate (BER) or hidden and exposed node problems.
- In a wireless multi-hop chain, the transmission of a packet interferes and contends with further data traffic in the wireless transmission range, particularly with consecutive packet of the data flow. This effect is known as multi-hop data traffic interference.
- A paradigm change in congestion detection since traditional mechanisms, such as packet loss or retransmission timeout, are not suitable for mobile ad hoc networks.

- The most scarce resource in a VANET is the wireless bandwidth. This bandwidth must be shared among competing flows whereas road specific packets, such as safety information, must be scheduled at a higher priority than other data traffic.

- Highly dynamic network topology.

The performance of a transport protocol in this vehicular environment depends on its ability to deal with these challenges.

The following design of a vehicular transport protocol (VTP) focuses on point-to-point applications, such as media transmission or email which require reliable and in-order data delivery. The demands of these applications are similar to the service provided by TCP in the Internet [66]. However, TCP performs poorly in wireless, mobile ad hoc networks [43, 48, 58].

A variety of TCP extensions aim at performance improvement in wireless, multi-hop networks, e.g., [48, 7, 87, 40]. Since most TCP extensions still remain below an optimal performance, *non-TCP* approaches argue that basic TCP design elements are inappropriate for wireless ad hoc networks, and transport performance can be significantly improved (i.e., compared to TCP extensions) when considering the specific characteristics of the environment [125].

We argue that the unique characteristics of VANETs necessitate the development of a new transport protocol that is specifically tailored to these conditions. Thus, the following VTP design takes the path characteristics of multi-hop communications in VANETs into account, as evaluated in our paper [82]. The key features of the VTP design are:

- The VTP sender uses a rate-based transmission scheme. The transmission rate is determined by a *rate timer* that steadily schedules the transmission of data packets when multi-hop connectivity between source and destination is recognized.

- VTP decouples congestion control from error and flow control, mainly to avoid throughput reduction for non-congestion-related packet loss. In VANETs, packet losses are frequent because of high mobility and the resulting topological changes. These losses must not invoke congestion control.

- VTP uses explicit signaling of available bandwidth from intermediate nodes for congestion control. The estimation of available bandwidth by intermediate nodes uses information from the MAC layer protocol.

- VTP provides reliability via retransmissions of lost packets. Selective acknowledgments (SACKs) report lost packets to the VTP sender. The receivers transmit SACKs in dynamic intervals. It adjusts the interval according to the current transmission rate and the source-destination distance.

- The VTP sender uses statistical knowledge to predict the expected communication behavior of a connection. In absence of acknowledgments, the expected communication duration for the respective source-destination distance assists the rate timer calculation.

Prior to the specification of a VTP, a detailed analysis of the path characteristics in a highway scenario quantifies the expected connectivity and disruption durations and evaluates the metrics packet loss, round trip time (RTT), RTT jitter and reordering. These statistical results directly influence the following transport protocol design. A simulative evaluation of VTP shows the performance improvements compared to TCP.

This work was conducted in the framework of the *Network on Wheels* project, which is supported by the German Ministry of Education and Research (BMB+F). The project investigates key technical questions for VANET communication, including transport protocols and position-based routing [101].

3.2 Background

The background comprises an overview of ad hoc routing protocols, including topology-based and position-based approaches. The choice of an appropriate routing protocol for certain scenarios is important because the performance of the ad hoc network significantly depends on the routing protocol. Since the routing protocol influences the characteristics of the ad hoc network, the design of a transport protocol should be aware of the routing protocol and the best performing protocol should be chosen anyway.

3.2.1 Ad Hoc Routing Protocols

This section surveys topology-based and position-based ad hoc routing protocols and discusses the suitability of the respective proposals for the high dynamic vehicular environment.

3.2.1.1 Topology-Based Ad Hoc Routing Protocols

Topology-based routing protocols establish routes on the basis of topological information about the network. They can be distinguished in proactive, reactive and hybrid approaches.

Proactive algorithms, like the dynamic destination-sequenced distance-vector routing (DSDV) [110] or the optimized link state routing protocol (OLSR) [30], maintain all available routes in the network even if some of the routes are currently not used. Particularly in a highly dynamic scenario, the maintenance of

unused routing information wastes a significant amount of the available wireless bandwidth [34].

In contrast, reactive (or on-demand) routing protocols, such as dynamic source routing (DSR) [67], the temporary ordered routing algorithm (TORA) [102] or ad hoc on demand distance vector routing (AODV) [109], maintain only the routes that are currently in use.

Hybrid routing approaches, such as the zone routing protocol (ZRP) [52], combine proactive and reactive routing components in order to improve routing scalability and efficiency.

All topology-based algorithms establish and maintain end-to-end routes which frequently break in mobile ad hoc networks, as illustrated in Figure 3.1. The route break recovery is time consuming and degrades the routing protocol performance since new routes must be build up on demand.

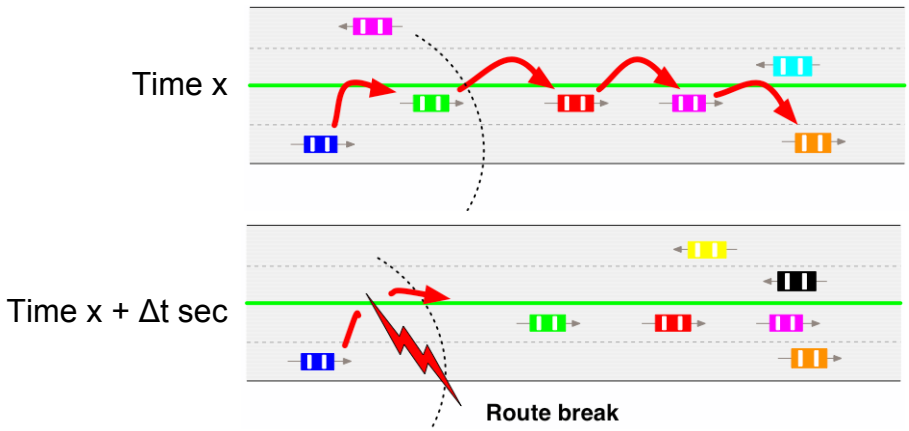

Figure 3.1: Route breaks in the highly dynamic vehicular environment.

3.2.1.2 Position-Based Ad Hoc Routing Protocols

Position-based routing protocols (PBR), as surveyed in [93], eliminate some of the limitations of topology-based routing by using additional geographical information. PBR scales well in the highly dynamic vehicular environment because each forwarding node locally determines the next hop independently for each packet, based on geographical position information. Thus, PBR does not require the establishment or maintenance of routes.

Prior to transmitting the first packet of a connection, the sender requires to determine the geographical position of the destination. Typically, a location service is used for this task.

In mobile ad hoc networks, the focus is on decentralized location services because a centralized approach would require the position of a reachable location server; the reachability of a location server cannot be guaranteed at all times. Typical examples of decentralized location services, such as the *quorum-based* location service [51], the *grid* location service [83] or the *homezone* location service [49, 123], are included in the PBR survey [93].

The geographical position of the destination is included in the packet header and determines the forwarding decisions of the sender and intermediate nodes. This mechanism requires that each node is aware of its own geographical position and the positions of its single-hop neighbors within radio transmission range.

Each node acquires its own position via a positioning system, such as the satellite-based global positioning system (GPS) [57, 69] or other types of positioning services [24, 55]. The positions of the single-hop neighbors within radio range are typically distributed in broadcast messages, termed *beacons*. All nodes periodically broadcast these beacons to their single-hop neighbors. Thus, each node is aware of the position of nodes in its vicinity. The accuracy of this information depends on the beaconing interval.

The PBR forwarding decision at each node depends on the destination's position, as contained in the packet, and the positions of the potential forwarding nodes. Typically, a forwarder selects the next hop that is geographically closest to the destination, using the greedy forwarding strategy of the greedy perimeter stateless forwarding (GPSR) [70, 71] approach, as visualized in Figure 3.2.

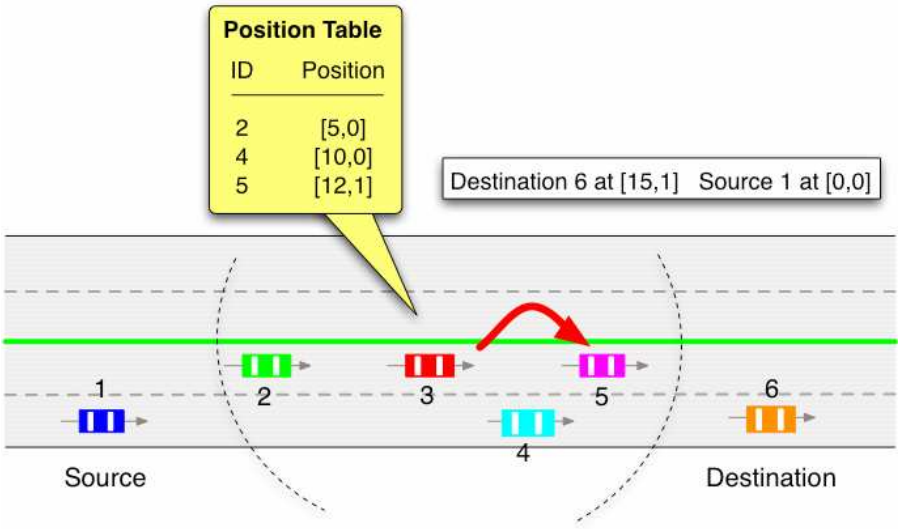

Figure 3.2: PBRV greedy forwarding strategy.

Node (3) is the current forwarder in the communication flow between the source node (1) and the destination node (6). The packet contains the position [15,1] of the destination and node (3) has three single-hop neighbors: Node (2) at [5,0], node (4) at [10,0] and node (5) at [12,1]. Consequently, node (3) selects node (5) as next hop because node (5) is closest to the destination.

Note that intermediate hops may update the position of the destination in the packet header, in case it has more accurate position information than contained in the packet. Thus, the preciseness of position information increases the closer the packet comes to the destination.

In order to determine the next hop, different forwarding strategies are possible. [93] classifies three main packet-forwarding strategies for PBR, such as greedy forwarding and restricted directional flooding.

As explained in the example above, greedy forwarding selects the next hop within radio range that is located closest to the destination. The restricted directional flooding algorithm works in a similar way, but forwards the packet to a set of single-hop neighbors that are closer to the destination. However, greedy forwarding and restricted directional flooding approaches fail if there is no single-hop neighbor that is closer to the destination than the forwarding node itself, but a route still exists. These scenario is illustrated in Figure 3.3.

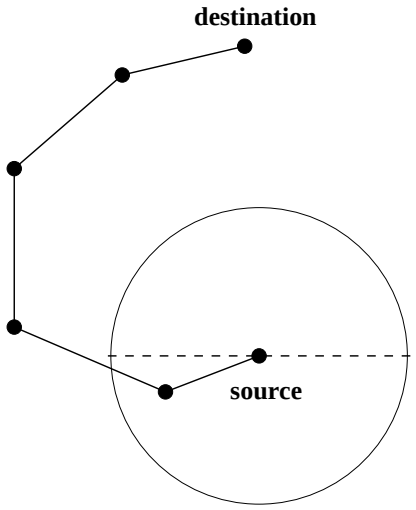

Figure 3.3: Greedy routing failure scenario.

To alleviate these situations, recovery strategies have been proposed, such as the *face-2* algorithm [19] or to select the node with the least backward (negative) progress [127]. However, the latter recovery strategy might lead to routing loops. Further approaches propose not to forward packets that encounter a local maximum [60]. Further details are given in the PBR survey [93].

3.2.1.3 Contention-Based Forwarding

The periodic exchange of beacons in PBR introduces overhead and consumes scarce wireless bandwidth. A position-based routing algorithm which does not require beacon information is proposed in contention-based forwarding (CBF) [46].

CBF broadcasts the packet to all nodes within transmission range of the forwarding node. Each node that is closer to the destination than the forwarding nodes enters a contention period, i.e., initiates a timer. The timeout depends on the distance to the destination. Upon timeout of the node closest to the destination, the node rebroadcasts the packet and silences all other potential forwarders, which cancel their respective contention timers.

Since not all nodes may overhear the rebroadcasting, packet duplication may occur. In order to avoid packet duplications CBF incorporates three different suppression strategies: Basic suppression, area-based suppression and active selection.

The basic suppression scheme stops the packet transmission of a packet only when overhearing the rebroadcasting, as explained above. Thus, packet duplication may occur.

Area-based suppression reduces the probability of duplication because only nodes in a pre-selected geographical area enter the contention. The area is selected in a way that the potential forwarding nodes are within transmission range of each other and that can overhear the rebroadcast packet.

The active selection is inspired by the RTS/CTS mechanism and selects the next hop prior to packet transmission. The current forwarder coordinates the selection process. The active selection scheme avoids packet duplication at the cost of additional control message overhead.

3.3 Related Work

This section surveys the related work on transport layer approaches in wireless and mobile ad hoc networks, as classified in Figure 3.4.

The highest level of classification distinguishes between TCP enhancements and non-TCP proposals. Ad hoc extensions for TCP may separate the connection in a wired and wireless part whereas the majority of approaches respect the end-to-end semantics of TCP. Finally, modifications may affect TCP only or cross-layer approaches involve other protocol layers, such as the network layer.

3.3.1 Transport Challenges in Wireless and Mobile Ad Hoc Networks

Besides the challenges of the wireless communication medium, such as a high bit error rate (BER) or the hidden and exposed node problem, the highly dynamic vehicular environment poses specific challenges to the design of a transport protocol for vehicular networks [98], as follows.

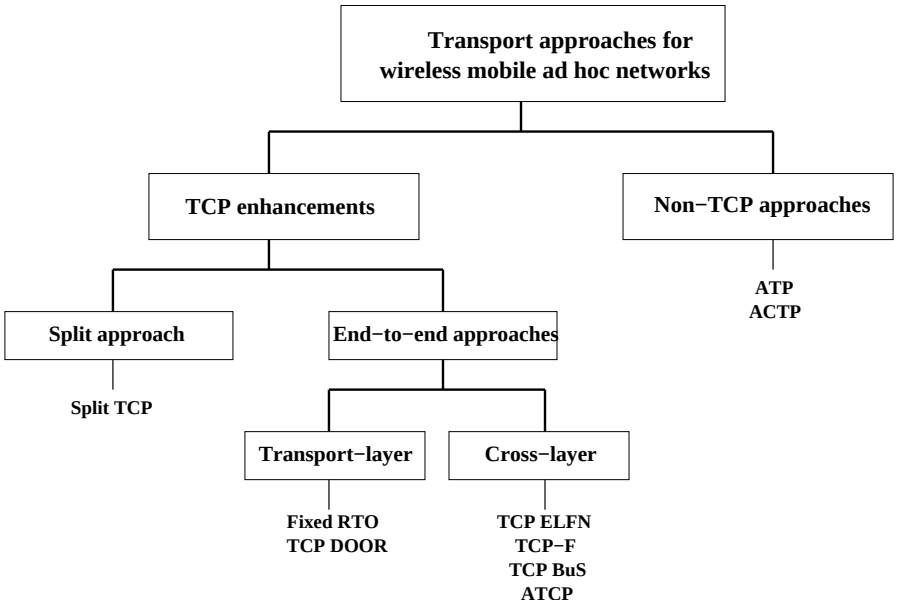

Figure 3.4: Transport layer approaches for wireless mobile ad hoc networks.

- Multi-hop data traffic interference. Vehicular wireless ad hoc networks utilize multi-hop relaying of data packets. A link layer transmission affects all nodes within the wireless transmission range of the sender due to the shared nature of the wireless channel. Particularly, the forwarding of a data packet contends with the transmission of the next packets by the predecessor node. The contention aggravates with further induced traffic of other data streams or the acknowledgments on the reverse path. A transport protocol for vehicular networks should consider these interferences in order to provide for a fair sharing of resources among contending flows.

- Congestion detection. Traditional mechanisms to detect network congestion, such as packet loss and retransmission timeout, are not appropriate in VANETs because a considerable amount of packet loss in VANETs is not due to congestion. Besides congestion, packet loss can occur due to the high bit error rate (BER) of the wireless medium, hidden terminal problem, packet collisions or mobility of vehicles (i.e., route breaks or route changes). Furthermore, the network contention in VANETs is location-dependent because all nodes within radio transmission/interference range contend for the channel. Therefore, the related work in [98] proposes to decouple congestion control from reliability and flow control in order to improve transport protocol performance.

- Bandwidth constraints and data traffic priorities. The most scarce resource in a VANET is the wireless bandwidth. This bandwidth must be shared among contending flows and other data packets whereas road safety information must be scheduled at a higher priority than other data traffic. A vehicular transport protocol must utilize and fairly distribute the available bandwidth and consider the different priorities that impact the performance of the protocol.

- Highly dynamic network topology. The fast movements of vehicles result in continuous topology changes in the ad hoc network. This intense topological change rate demands for specific network protocols because the performance of a routing or transport protocol depends on its ability to adapt quickly to the varying path characteristics. A vehicular transport protocol that considers the specific characteristics of a VANET can provide improved protocol performance.

3.3.2 TCP Performance in Wireless and Mobile Ad Hoc Networks

The transmission control protocol (TCP) [66] is the *de facto* transport protocol standard of the Internet. However, TCP performs poorly in wireless and mobile networks [43, 48, 58] for various reasons. This section provides an overview of TCP performance evaluations and explains the reasons of TCP performance degradation in wireless, mobile networks.

TCP provides reliable and in-sequence byte stream delivery of data to applications in point-to-point communications. In theory, TCP should be independent of the underlying technology and support all kinds of networks. These networks are typically unreliable. In practice, however, the coupling of error and congestion control in the design of TCP tunes the protocol to the characteristics of wired networks: TCP assumes network congestion in case of packet loss, which is almost always true in wired networks.

However, in wireless and mobile ad hoc networks, a significant amount of packet loss is not due to congestion. In this scenario, packets may be lost in addition to congestion e.g., because of:

- Channel errors and the high bit error rate (BER) of the wireless medium,

- Collisions,

- Path asymmetry,

- Route failures and network partitions due to mobility.

The invocation of congestion control in reaction to non-congestion losses results in a performance degradation of TCP [43]. Though the different versions of

TCP (i.e., Tahoe, Reno, NewReno and Vegas), as explained in Section 2.2.3, perform differently in ad hoc networks, all these versions cannot distinguish between packet loss due to congestion and packet loss due to wireless or ad hoc characteristics, as explained above (see also [138]).

The following related work evaluates the TCP performance in wireless and mobile ad hoc networks. The work [53] classifies the reasons for the poor TCP performance in ad hoc networks in *(i)* wireless and *(ii)* mobility and multi-hop related challenges.

(i) The TCP performance in static, wireless single-hop networks is mainly influenced by the characteristics of the wireless network, such as high BER, path asymmetry or hidden and exposed stations. A variety of TCP enhancements address this scenario, e.g., infrastructure-based WLANs [4, 6, 9], mobile cellular networking [10, 20] or satellite communications [2, 36].

(ii) In mobile ad hoc networks, the TCP performance - additionally to the challenges mentioned above - depends on the mobility. Node movement leads to temporal network partitions, route breaks and route changes.

The impact of mobility on the throughput of TCP is evaluated through simulations in [58]. Using the random way-point (RWP) [14] mobility model without pause time, the evaluation shows that for certain mobility patterns the throughput is close to zero whereas other patterns result in high throughput. Furthermore, velocity variations affect the TCP throughput differently, e.g., an increase of the average speed from 2 m/s to 10 m/s results in a significant throughput degradation whereas the increase from 10 m/s to 30 m/s affects the throughput only slightly. The analysis of the simulation traces of low throughput shows that the route recovery of the DSR routing protocol invokes TCP retransmission timeouts and congestion control in case of route breaks. In simulation traces of high throughput, sender and receiver move closer towards each other, so that the previous DSR route remains valid until a new, shorter route is found. In order to prevent congestion control in case of route breaks, the authors propose explicit link failure notification (ELFN) to stale TCP for the duration of a route break.

The paper [3] takes up the mobility related throughput degradation of TCP and defines the TCP throughput in mobile ad hoc networks as a function of node velocity and network load. A significant part of the route re-computation latency depends on the *MAC failure detection latency*. The *MAC failure detection latency* is defined as the amount of time spent before the MAC recognizes a link failure. Finally, [3] identifies an additional problem related to routing and MAC layer, which the authors term *MAC packet arrival latency*. When a link failure along the path is detected, it is reported back to the routing agent on the sending node. In case further nodes also used this link before, the node that detects the link failure has to wait for further packets of these communications before reporting the link failure also to these sending nodes. This delay also contributes to the route re-computation latency.

The performance of TCP Reno over the routing protocols AODV [109], DSR [67] and ADV [18] in mobile ad hoc networks is evaluated by simulations in [37]. The results show that ADV performs well for a variety of mobility patterns and topologies. In order to improve TCP performance over on-demand routing protocols, the work also presents a TCP enhancement proposal termed *fixed RTO*, which is explained in detail in the following survey of TCP enhancements for mobile ad hoc networks.

The work in [86] evaluates TCP performance over multi-path routing, i.e., Split multipath routing (SMR) with maximally disjoint paths [89]. Multi-path routing provides advantages in wireless ad hoc networks, such as reduction in route computation time, high resilience to path breaks, high call acceptance ratio and better security. However, the evaluation shows that TCP performance suffers from multi-path routing. The inaccuracy of the average RTT in multi-path routing leads to more TCP timeouts and the increased out-of-order delivery due to different paths trigger TCP duplicate ACKs, which in turn invoke TCP congestion control.

The following sections survey TCP enhancements and non-TCP approaches, as shown in Figure 3.4 in order to improve transport layer performance in wireless mobile ad hoc networks.

3.3.3 TCP Modifications for Wireless and Ad Hoc Networks

This section surveys TCP enhancements to improve TCP performance in mobile ad hoc networks according to Figure 3.4, including split connections and end-to-end approaches as well as pure TCP approaches and cross-layer solutions.

We first present approaches that respect the end-to-end semantic of TCP and modify TCP only.

The *fixed RTO* proposal [37] modifies the retransmission (RTO) calculation of the TCP sender. The sender does not purely rely on RTT measurements in the network anymore, but uses a heuristic to distinguish between route failures and congestion. In case of two consecutive RTO timeouts, the sender assumes a route failure. The sender retransmits the unacknowledged packet, but does not increase the RTO value, as original TCP would due to its *exponential backoff* mechanism. The RTO timeout value remains fixed until the route is re-established and the outstanding packet is acknowledged. In [37] the authors evaluate the approach simulatively over on-demand routing protocols and consider TCP's selective and delayed acknowledgment options. The fixed RTO approach improves TCP performance for on-demand routing protocols, but is restricted to wireless networks only.

TCP DOOR (detection of out-of-order and response) [135] detects out-of-order delivery of TCP segments and interprets these as an indication of route failures. The detection of out-of-order data delivery can be sender-based or receiver-based. The sender-based identification of out-of-order packet delivery requires additional

information in the TCP duplicated acknowledgments, i.e., *ACK duplication sequence number (ADSN)*. TCP DOOR appends the ADSN as a one-byte option to the duplicated acknowledgments and increases the number for each ACK. The receiver-based approach adds the *TCP packet sequence number (TPSN)* as a two-byte option to each packet to identify out-of-order delivery. Upon detection, the receiver notifies the sender about the out-of-order delivery via the *OOO-bit*, which is explicitly defined in the ACK packet header. When the sender is aware of an out-of-order event, it temporarily disables congestion control and instantly recovers during congestion avoidance. The duration of congestion control disabling depends on the RTT. TCP DOOR is evaluated through simulations and the results show similar performance for sender- and receiver-based out-of-order detection. Thus, the authors recommend the sender-based approach because it does not require explicit out-of-order notifications.

The following approaches respect the end-to-end semantics of TCP and require cross layer interaction between TCP and the network layer.

TCP-F (feedback) [25] relies on network feedback to detect routing failures. When the routing agent of a node along the path detects a broken link, it replies with an explicit *route failure notification (RFN)* to the sender. Upon reception of an RFN, the sending instance of TCP enters a *snooze* state. In this state, TCP-F stops packet transmission of the respective flow, freezes all TCP variables, such as timers and congestion window, and initiates a *route failure timer*. The sender remains in snooze state until either it is notified about route recovery by a *route re-establishment notification (RRN)* or the route failure timer expires. When a new route is found and the sender receives a RRN notification, the TCP-F sender resumes transmission based on the previous congestion window and timers - independent of the fact that the conditions along the new path might have changed. When the route failure timer expires, TCP-F enters the standard congestion control mechanism. The evaluation results in [25] show performance improvements, however, the simulation scenario is very basic.

The *explicit link failure notification (ELFN)* [58] mechanism also uses explicit network feedback to detect route failures, similar to TCP-F. In contrast to TCP-F, ELFN is based on real interaction between TCP and the routing protocol. The ELFN notification is piggybacked on the routing failure notification. It contains sender and receiver addresses, ports and the respective TCP sequence number. Upon reception of an ELFN message, the TCP sender enters a *standby* mode. In this mode, the TCP sender periodically transmits *probe* packets in order to check if the route is restored. In [58] the optimal probe interval is two seconds, but for different scenarios the authors propose to adjust the probe interval according to the last stored RTT value. In case a probe packet is acknowledged, TCP resumes. The evaluation comprises a TCP resume with frozen congestion window and timers, as well as with reset values where the former option provides better results. Though

the evaluation results in [58] show a significant enhancement over standard TCP, further evaluations are required, particularly because [3] and [95] come to different results for high and even for low network load, respectively.

Ad hoc TCP (ATCP) [87] once more uses network feedback, but aims at the detection of route failures *and* tries to make the high BER of the wireless medium transparent to TCP. ATCP inserts an additional layer between TCP and IP. This layer monitors the network state and sets TCP in persist, congestion control or retransmit state. The ATCP layer uses the network information provided by ECN (explicit congestion notification) [68] or the ICMP 'destination unreachable' message to determine the respective state. When receiving the ICMP 'destination unreachable' message, TCP assumes a route failure and enters the persist state. In this state, TCP freezes its congestion window and timers and transmits periodical probing packets until a new route is discovered. The ECN explicitly notifies the sender about network congestion. Consequently, TCP enters its congestion control (i.e., without waiting for timeouts). Upon reception of three duplicated acknowledgments, ATCP enters the retransmit state. This means that the third duplicated acknowledgment is not relayed to TCP. Instead, TCP is set to persist mode and ATCP retransmits the lost segment autonomously, assuming a loss due to the high BER. When ATCP receives the respective acknowledgment, TCP resumes to normal operation. ATCP has been evaluated by implementation in a testbed, emulating different scenarios, such as congestion, packet loss, temporal partitions and reordering. However, the testbed consisted of wired Ethernet connections. Therefore, further evaluation in a wireless and mobile ad hoc environment is required.

TCP-BuS (TCP buffering capability and sequence information) [74] again uses network feedback to detect routing failures and introduces a buffering capability to mobile nodes along the path. TCP-BuS signals route failure and re-establishment via *explicit route disconnection notification (ERDN)* and *explicit route successful notification (ERSN)* messages. The node that detects the link failure and sends the ERDN is called *pivoting node (PN)*. Upon reception of an ERDN message, the TCP-BuS sender stops packet transmission and all packets in transit from the source to the PN are buffered by the PN during the route re-establishment phase. To avoid timeouts during route re-establishment, the retransmission timeout for buffered packets is doubled. Selective retransmission requests by the PN ensure that the sequence of buffered packets is complete. When the route is recovered, the PN relays all buffered packets to the receiver and informs the sender via the ERSN message. Consequently, the sender resumes the transmission. The evaluation in [74] shows that TCP-Bus outperforms standard TCP and TCP-F in different scenarios.

Finally, the split TCP approach provides an example that separates the wired and the wireless part of the end-to-end TCP connection.

Split TCP [79] separates the TCP end-to-end connection into different segments, e.g., according to the characteristics of the network types in between in order to improve throughput and fairness. The node between the split connections is called *proxy*. The routing agent on intermediate nodes assigns proxy functionality to the node, according to its *inter-proxy distance and network parameters*. A proxy intercepts TCP packets from the source or a previous proxy and replies with a *local acknowledgment (LACK)*. The respective proxy is responsible for the reliable forwarding of data to the destination or the next proxy. In order to maintain the TCP end-to-end semantics, an additional acknowledgment is exchanged between destination and source. Split TCP separates end-to-end reliability and congestion control by maintaining two transmission windows at the source, i.e., the *congestion window* and the *end-to-end window*. The congestion window determines the transmission rate along a segment of the path, i.e., the wired path from the source to the proxy or the wireless path from the proxy to the destination, according to the standard TCP congestion window mechanism. The end-to-end window controls the end-to-end path between source and destination. Basically, the congestion window adjust the transmission range in the wired and wireless parts of the network while the end-to-end window may intervene for the overall control, e.g., in case one of the connections is broken. The evaluation of Split TCP shows a performance improvement of about 30% for the inter-proxy distance of three to five hops, at the cost of increased network overhead, buffer sizes and complexity.

3.3.4 Non-TCP Approaches

In contrast to the TCP enhancements, non-TCP transport layer approaches use completely different algorithms for error, congestion and flow control in mobile ad hoc networks, as indicated by the separate branch in Figure 3.4. This design choice is based on the argument that basic design elements of TCP are fundamentally inappropriate for the specific characteristics of mobile ad hoc networks [125]. Non-TCP approaches outperform TCP and TCP enhancement proposal because the specific characteristics of the environment are considered. However, non-TCP approaches lack immediate interoperability to TCP when connected to a fixed network. Interoperability requires additional effort, e.g., gateway proxies for transport protocol translation.

The *application-controlled transport protocol (ACTP)* [92] ensures reliability in the application layer. This light-weight protocol between the application and UDP provides and maintains simple feedback information about the successful delivery of packets, potential loss of a packet in the absence of the respective acknowledgment, a retransmission timer and the lifetime of a packet. The retransmission timer and the packet's lifetime depend on the RTT (i.e., typically four times the RTT). Beyond, ACTP supports priorities of packets when the lower layers provide the respective differentiated services.

The ad hoc transport protocol (ATP) [125] decouples congestion control from reliability and relies on congestion feedback from intermediate nodes along the path to adjust its rate-based transmission scheme.

Intermediate nodes may reduce the requested bandwidth, as contained in every ATP packet, according to their local queuing delay and contention delay. The ATP receiver collects the congestion information and includes the weighted average of this information and the actual available receive buffer (i.e., flow control) in periodic, selective acknowledgments (SACKs). Note that the congestion information in the acknowledgment is derived from the network whereas the flow and reliability information are obtained from the receiver. The SACK period can be dynamically adjusted according to the round trip time.

Upon reception of a SACK, the ATP sender adjusts its timer-based transmission rate to the conditions along the path. The evaluation of the congestion feedback by the sender results in an increase, decrease or maintenance of the current rate (in contrast to TCP which has only decrease or increase phases). The transmission rate also depends on cross-layer information from lower layers. Upon reception of an ELFN, the ATP sender reduces its transmission rate to *probe* packets.

In order to estimate the transmission rate on connection establishment or re-establishment after a route failure, ATP uses a *quick start* mechanism: The ATP sender transmits a probing packet to collect the congestion information from the intermediate nodes along the path. The ATP receiver immediately acknowledges this probing packet. Consequently, the ATP sender can transmit at the maximum possible rate after one RTT.

The vehicular transport protocol, as presented in Section 3.5, adopts some of these basic mechanisms, such as decoupling of error and congestion control, explicit congestion signaling and the utilization of selective acknowledgments. However, VTP differs from ATP in order to respect the specific characteristics of VANETs, as follows.

ATP operates on top of topology-based routing protocols which maintain an end-to-end route. Thus, ATP uses route failure messages from the network layer, and the intermediate nodes maintain per-flow information. In contrast, VTP assumes position-based routing (PBR) as the underlying routing protocol because PBR outperforms topology-based routing protocols in VANETs [44]. With PBR, each node selects the next reachable forwarder that is geographically closest to the destination. Thus, the path that consecutive packets follow may be different due to mobility. Consequently, there are no network messages to assist VTP, and intermediate nodes do not maintain flow information.

Furthermore, ATP does not use retransmission timers. When missing packets are reported in the periodic SACK, the ATP sender instantly retransmits the data. However, the path characteristics in VANETs [82], such as packet loss, reordering, RTT and RTT jitter, suggest for retransmission timers. The VTP retransmission timer cannot rely on the actual RTT because of RTT fluctuations, and SACKs are

transmitted periodically which decreases the accuracy of RTT estimation. Thus, the VTP retransmission timer calculation is based on the source-destination distance and the statistical results of the path characteristics [82].

Before Section 3.5 explains the VTP protocol in detail, the following Section 3.4 evaluates the path characteristics for highway scenarios, as the prerequisite for our VTP design. The evaluation focuses on highway scenarios as the main application domain for VANETs. The accurate functioning in high speed scenarios, e.g., safety related applications in highway scenarios, is important for user acceptance since user should trust the system in life endangering situations. Furthermore, in highway scenarios there might be no road side infrastructure to support the protocol, like this could be used in city scenarios. However, the following design of VTP is transferable to further scenarios, such as cities, by adjusting its parameters, assuming that the underlying routing protocol finds an end-to-end path with a similar performance like in highway scenarios. This is very likely because as an example in cities the existing infrastructure might be used to support the ad hoc network.

3.4 Evaluation of Path Characteristics on Highways

This section investigates the path characteristics that transport protocols experience in VANETs in highway scenarios. These results aid the design of a vehicular transport protocol (VTP). The behavior and performance of a VTP mainly depends on its ability to adapt quickly to varying path characteristics. In the following, analytical and simulative evaluations of connectivity and disruption durations estimate the expected connectivity between communication partners for specific distances. Furthermore, the results quantify packet losses, reordering, round trip times (RTT) and RTT jitter through simulations.

3.4.1 Scenario and Simulation Environment

The scenario, as illustrated in Figure 3.5, simulates varying numbers of vehicles on a 10 km stretch of highway. The spatial distribution of the vehicles and their mobility behavior, i.e., position, direction and speed, are derived from validated highway mobility patterns [45, 80]. The analysis considers different scenarios that have different numbers of lanes in each direction, and varying numbers of vehicles per kilometer.

In the *ns-2* [132] simulation environment, all vehicles are equipped with a single IEEE 802.11b wireless interface providing a radio transmission range of 250 m. Vehicles in radio range can communicate directly. In case the distance between communication pairs exceeds the radio range, but a multi-hop path exists, the vehicles form a self-organizing *ad hoc* network that supports multi-hop communication.

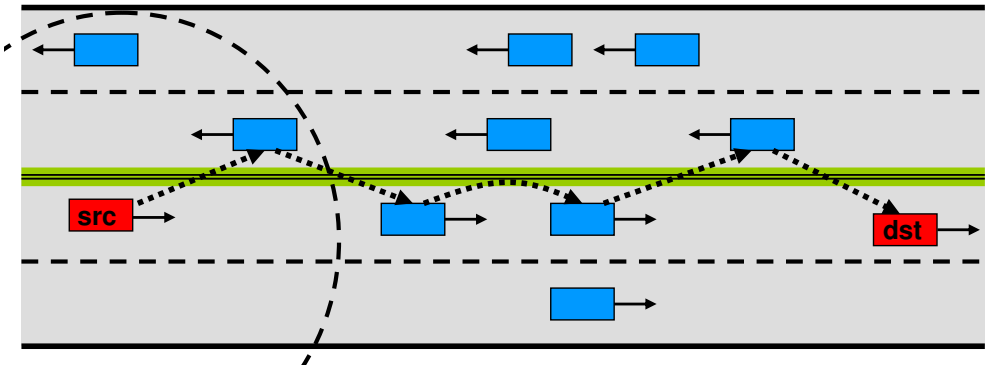

Figure 3.5: Multi-hop inter-vehicle communication in the highway scenario.

Generally, communication in this VANET occurs between random pairs of vehicles distributed throughout the simulated area. However, some parts of the analysis restrict communication to vehicles within specific distances. The number of vehicle pairs that communicate concurrently determine the network load. The simulations investigate 5, 10 and 15 concurrent communications, representing light, medium and high loads. Each communication is a constant-bitrate data transfer using a fixed packet size.

The VANET uses PBR because it outperforms topology-based routing in the highly dynamic vehicular environment [44].

Mobility can also create temporary network partitions that interrupt end-to-end connectivity and cause packet loss. In order to reduce the number of network partitions, oncoming traffic is included when determining next hops [45].

3.4.2 Metrics

This section defines the evaluation metrics that characterize the path characteristics experienced by a single communication instance.

- A *connectivity period* denotes the existence of an end-to-end path between source and destination that enables communication. A *disruption period* denotes the absence of such a path. The *connectivity duration* hence describes the length of a connectivity period whereas the *disruption duration* describes the length of a disruption period. Note that infinite disruptions are not considered.

- The *packet loss probability* describes the likelihood that an individual packet is lost between source and destination, independent of other packets. The *packet loss burst length* describes the number of consecutively lost packets.

- Round trip time (RTT) and RTT jitter. *(i)* The *RTT* describes the time between the transmission of a packet and the reception of the first corresponding acknowledgment. The simulation takes one RTT sample at any given time. *(ii)* The *RTT jitter* describes the difference between two subsequent RTT samples. The *mean RTT* describes the mean across all RTT samples for a given communication.

- Reordered packets are received in a different sequence than they were sent. The *packet reordering probability* describes the likelihood that a packet is reordered, independent of other packets. The *reordering period* describes the time from the reception of the first reordered packet until the originally expected packet arrives. Lost and duplicated packets do not contribute to reordering.

3.4.3 Evaluation Results

This section presents selected simulation results for a highway scenario with two lanes per direction and on the average six vehicles per lane and kilometer. These results are chosen because the scenario is representative for typical weekday road traffic on a German highway. Each sender generates a constant bit rate (CBR) stream of 100 Kb/s. Although some path characteristics are expected to be different in the presence of a transport protocol with congestion control, the CBR streams approximate these environments. The duration of each simulation run is 60 s.

3.4.3.1 Connectivity and Disruption Duration

This section presents the evaluation of connectivity and disruption durations and compares the analytical and simulation results for maximum source-destination distances of 500 m and 2000 m.

First, we compute the connectivity and disruption durations by determining the theoretical availability of a multi-hop end-to-end path in discrete time intervals using global knowledge. These results represent an upper bound for the expected connectivity durations because MAC and physical effects are not considered. Ns-2 simulations that use an ideal MAC validate this first analytical evaluation. Further simulations evaluate the connectivity and disruption duration using the IEEE 802.11 MAC of ns-2.

Figure 3.6 shows the cumulative distribution function (CDF) of the normalized connectivity durations for 500 m and 2000 m source-destination distance.

The analytical evaluation shows that when the communicating nodes remain within a distance of 500 m, 9% of the communications are interrupted within 10 s and 20% are interrupted within the duration of the simulation. Consequently, 91% of the communications remain uninterrupted for 10 s and 80% of the communications continue for the complete duration of the simulation.

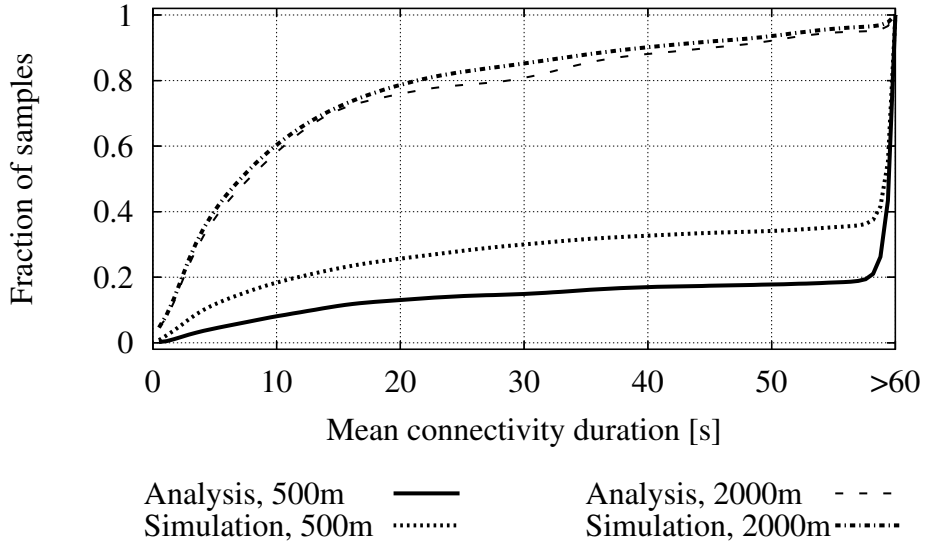

Figure 3.6: CDF of connectivity duration for analysis and simulations.

The results for the IEEE 802.11 MAC in the 500 m scenario in Figure 3.6 show a decrease in connectivity durations. After 10 s, 20% of the communications are interrupted and 38% of the communications are interrupted up to the end of the simulation. These differences between the analysis and the simulations with IEEE 802.11 MAC are mainly due to inaccurate location information in the latter case: The anylysis is based on global knowledge that provides always accurate and up-to-date location information. In contrast, the location information in the simulation is based on beaconing. Thus, the accuracy of the location information depends on the beacon interval. When a vehicle drives out of range before the location information of the sender becomes invalid, the sender may transmit a packet to a node which is not reachable any more, resulting in an interruption.

The connectivity durations significantly decrease for longer distances, as illustrated by the curve for 2000 m maximum distance in Figure 3.6. After 10 s, 54%, after 30 s, 82% of the communications, and after 60 s, 92% of the communications are interrupted. However, the ideal and 802.11 MAC curves converge for 2000 m maximum distance because the interruptions due to routing errors dominate in this case.

Figure 3.7 shows the CDF of the normalized disruption durations for 500 m and 2000 m source-destination distance. These results consider only disruptions of communications that resume. The main result is that the average disruption duration is short, as discussed in the following.

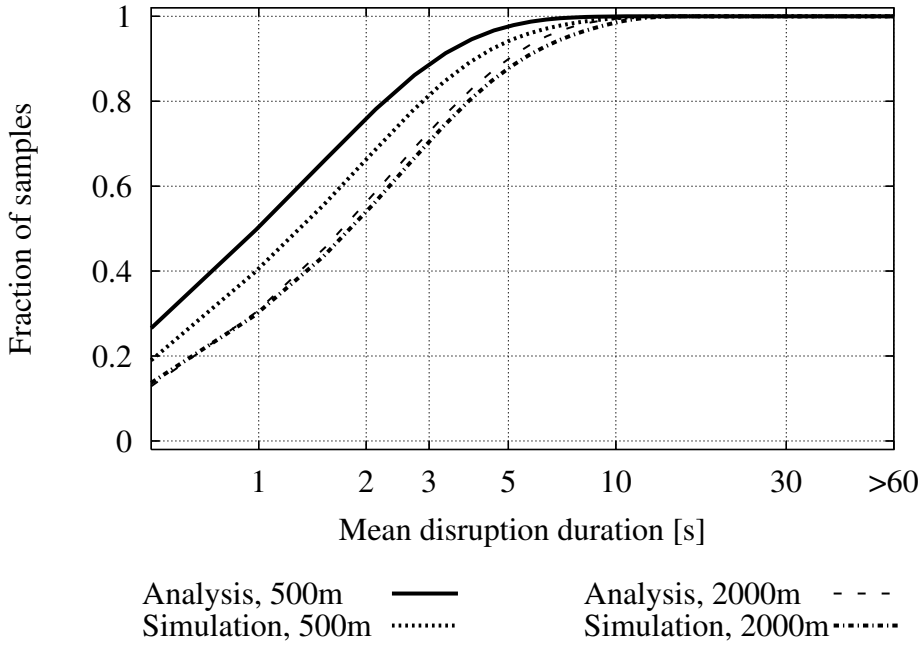

Figure 3.7: CDF of disruption duration for analysis and simulation.

In the analytical results, 92% of the disruptions end after 3 s for 500 m source-destination distance and 75% of the communications resume after 3 s in the 2000 m source-destination distance scenario.

Again, there is a noticeable difference for the 500 m distance when using the IEEE 802.11 model whereas the curves converge for 2000 m distance. In the simulations, for 500 m distance, 85% of the disruptions and for 2000 m distance, 72% end after 3 s.

Summarizing, the connectivity and disruption evaluation results show that for distances up to 2000 m, multi-hop vehicular communication on a highway is possible, and the expected communication duration is in the range of several seconds, depending on the source-destination distance.

3.4.3.2 Packet Loss Probability and Distribution

Packet losses are frequent in vehicular environments because of the high mobility and the resulting topology changes. This section shows the simulation results of loss probability and evaluates the expected number of consecutive losses in a vehicular highway scenario with low data traffic (i.e., five concurrent communication streams).

Figure 3.8 illustrates the loss probability over distance for standard PBR and PBR with *lost link* enhancement[44]. With standard PBR, neighbor table entries time out periodically and can become stale. The *lost link* enhancement aims to reduce packet loss by cross-layer integration, keeping the neighbor table updated based on link layer feedback, such as feedback about packet drops due to unsuccessful transmission after the maximum number of MAC layer retries.

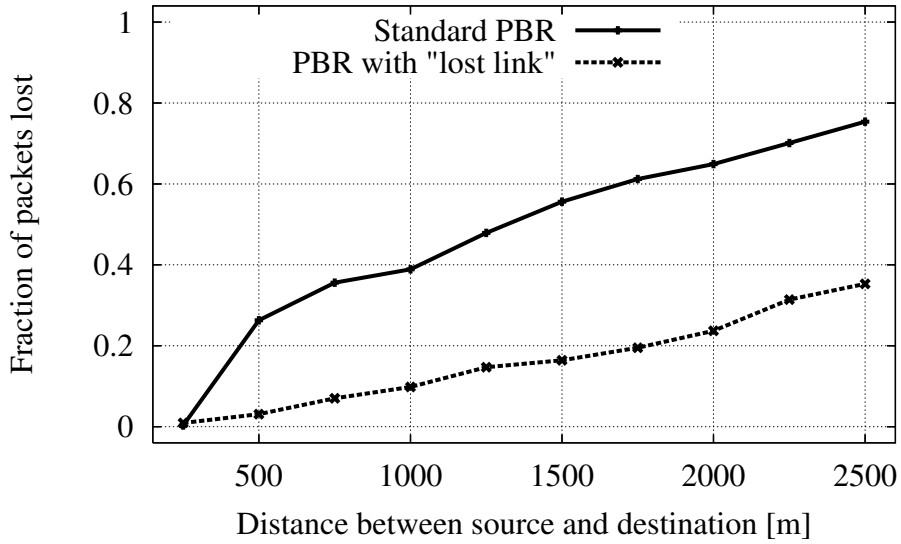

Figure 3.8: Loss probability over distance (standard and lost-link enhanced PBR).

Figure 3.8 shows already a significant loss probability in a scenario with light network load. For both curves, the loss probability up to 250 m is below 1%, due to wireless packet loss in single-hop communication. With standard PBR, the loss probability increases to 26% for 500 m distance because multi-hop communication is required beyond distances of 250 m. Beyond 500 m, the loss probability increases linearly with longer distances.

The curve for *PBR with lost link* in Figure 3.8 shows a completely linear curve of packet loss probability for PBR with *lost link* enhancement. Packet loss is significantly reduced, e.g., down to 3% for a distance of 500 m. However, the reduction of packet loss comes at the cost of increased RTT and RTT jitter because the probing of different neighbors is time consuming (see [81]).

Figure 3.9 illustrates the number of consecutive packet losses over all distances with light data traffic load. In this scenario, 31% of all losses occur as single packet losses, in 54% of the loss events, three consecutive packets are lost, and in 82% of the cases up to ten packets are lost subsequently. This result is mainly independent of the network load. In comparison, the results for a high-load scenario show 29% loss of single packet and 78% of ten subsequently lost packets (i.e., for the com-

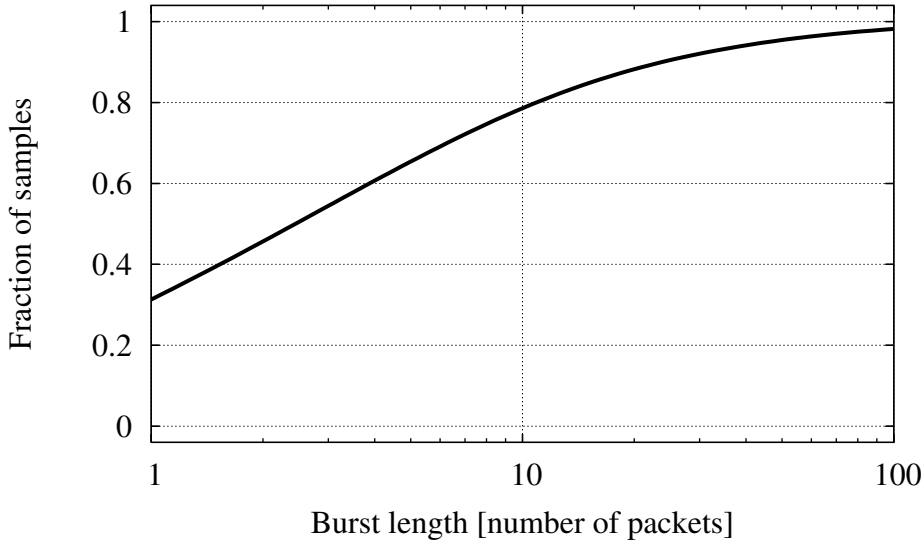

Figure 3.9: Burst length CDF of lost packets over all distances.

plete results see [81]). For high network loads, the number of consecutively lost packets increases mainly because additional queue drops occur, as shown in the loss reason evaluation in [81].

The probability and number of subsequently lost packets affects the error control mechanisms in the VTP design. However, consecutive packet loss of more than ten packets is mainly due to network partitions. The UDP communication continues transmitting during a disruption whereas a VTP should significantly decrease the transmission rate (e.g., send probing packets only).

3.4.3.3 Round Trip Time and RTT Jitter

Traditional transport protocols commonly use a measurement-based estimation of the RTT, e.g., to adjust the transmission window or to determine retransmission timeouts. This evaluation focuses on RTT and RTT jitter in order to determine if this metric is accurate for the use by a transport protocol in VANET environments.

Figure 3.10 shows the medium RTT, lower and upper quartile over distance for 15 concurrent communication streams. For a 500 m source-destination distance, the RTT median is 10 ms and the upper (i.e., 75%) quartile is 19 ms. However, the upper quartile increases significantly for longer distances, e.g., the median for 2000 m distance is 91 ms and the respective upper quartile is 295 ms.

Figure 3.11 shows the evaluation results of RTT jitter for consecutive samples over distance. The median RTT jitter for 500 m distance is 5 ms and the upper

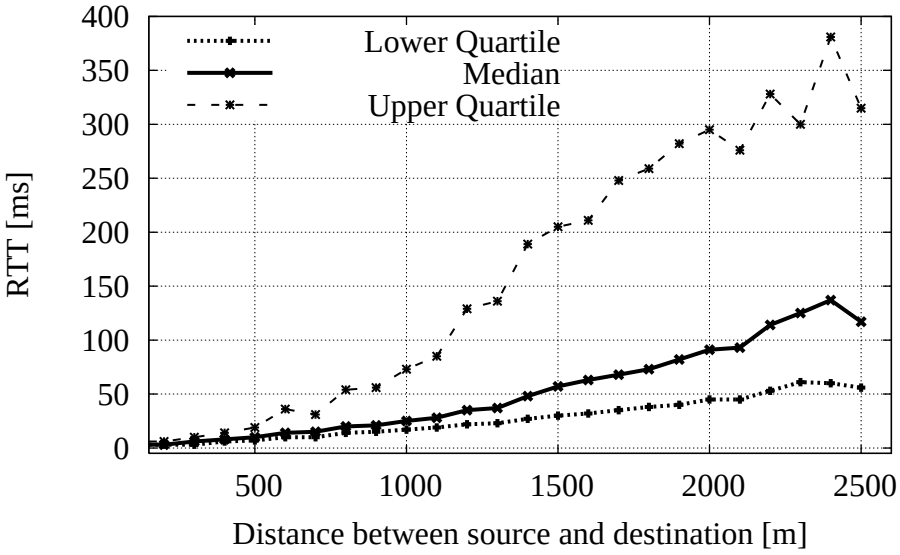

Figure 3.10: Median RTT and quartiles over distance.

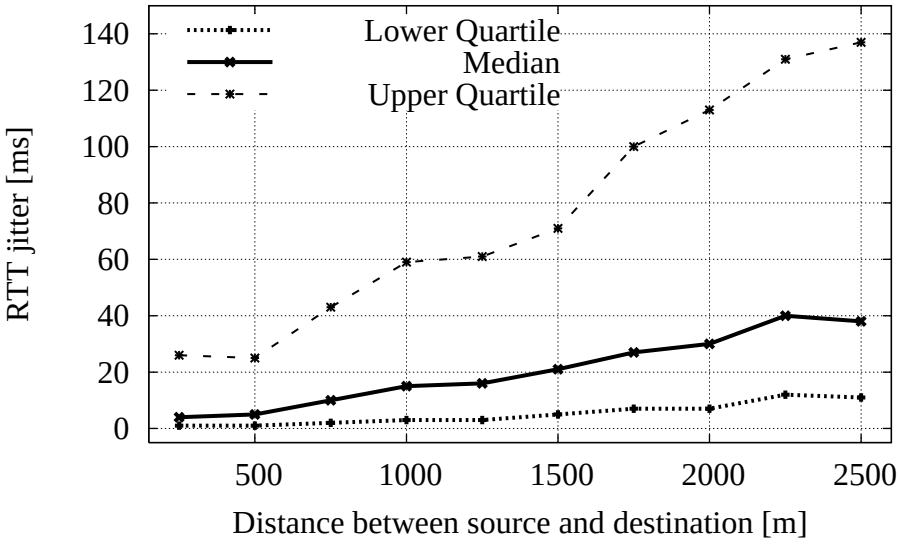

Figure 3.11: Median RTT jitter and quartiles over distance.

quartile is 25 ms. For 2000 m distance, the median RTT jitter is 30 ms and the upper quartile increases to 113 ms. The differences between the median and the upper quartile show that the RTT for consecutive packets also differs significantly. Summarizing, the use of measured RTTs in VANETs for a transport protocol is

problematic. Although the RTT and RTT jitter are acceptably small for source-destination distances below 700 m, higher distances result in extreme fluctuations in RTT, e.g., up to 300% for a 2000 m distance. As a further result, the distance can be regarded as a metric of reliability for the measured RTT value.

3.4.3.4 Packet Reordering Probability and Period

Assuming PBR as the routing protocol, VANETs have a significant probability of packet reordering because each packet might follow a different path towards the destination. Thus, the reordering probability depends on the network load along the paths and the source-destination distance, as shown in Figure 3.12.

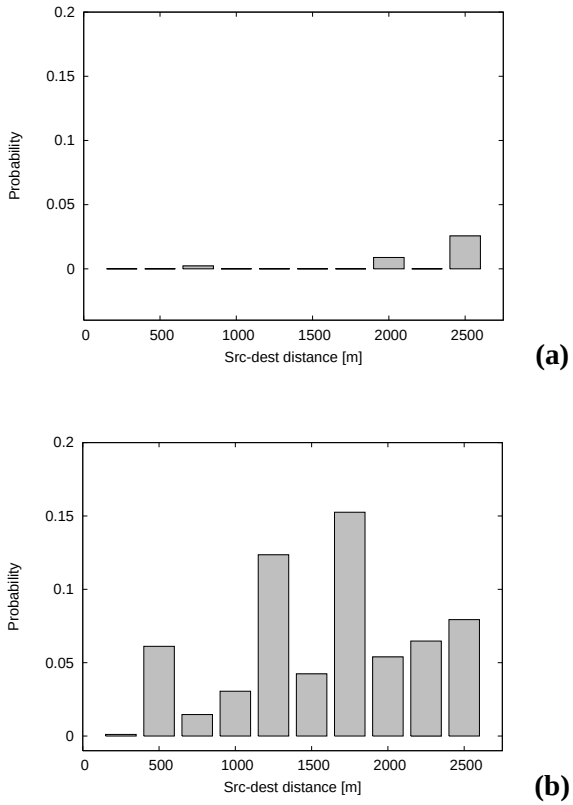

Figure 3.12: Packet reordering probability for (a) 5 concurrent communications (b) 15 concurrent communications.

The reordering probability in light-load scenarios with five parallel streams, as shown in Figure 3.12(a), remains below 1% for most distances. Only for 2500 m distance, the reordering probability reaches 2.5%. In comparison, Figure 3.12(b) shows the reordering probability for high-load scenarios. The fluctuations indicate

the strong dependence on the data traffic distribution in the network. However, generally a reordering probability of at least 15% should be expected in this scenario beyond a distance of 1500 m.

In addition to the reordering probability, the number of subsequently reordered packets and the period of reordering are important for the VTP design. Figure 3.13 compares the reordering period for different network loads in a cumulative distribution function (CDF) graph.

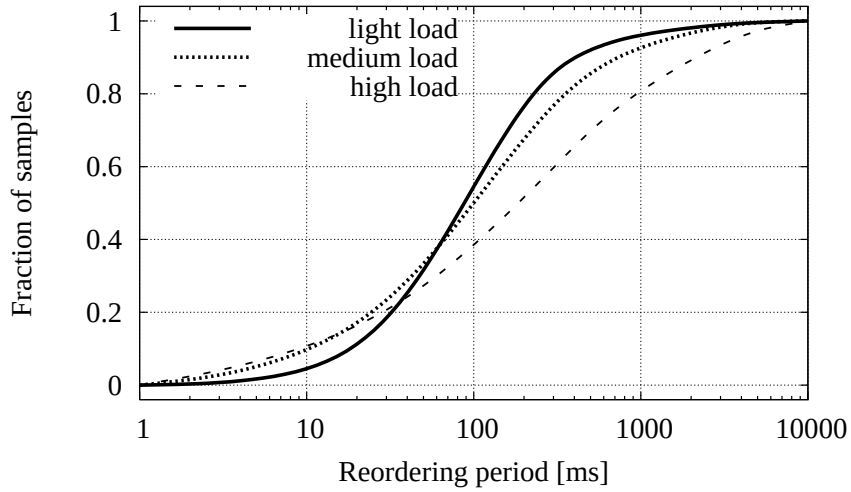

Figure 3.13: CDF of reordering period for different network loads.

In the light-load scenario, 2% of the samples remain in the reordering period for 10 ms, 60% for 100 ms and 96% for 1000 ms. In comparison, in the high-load scenario, 7%, 50% and 93% of the samples remain in the reordering period for 10 ms, 100 ms and 1000 ms, respectively. Note that the reordering probability affects the form of the curves and causes, e.g., the crossing of the curves.

Summarizing, the evaluation of the reordering characteristics aids the timer configurations of a VTP and can avoid redundant repetitions of delayed or reordered packets.

3.4.4 Path Characteristics Evaluation Summary

This chapter has evaluated the communication path characteristics for VANETs in typical highway scenarios, namely connectivity and disruption duration, packet loss, packet reordering, RTT and RTT jitter.

The connectivity and disruption evaluations show that for distances of up to 2000 m, steady communication is feasible. For a distance up to 2000 m, about 40% of the connections remain uninterrupted for 10 s in average. With decreasing distance, the connectivity duration even increases. Disruptions resume in average after 3 s, only marginally dependent on the distance.

The packet loss ratio for a constant packet stream is, however, huge: For a distance of 2000 m, standard PBR shows a packet loss rate of almost two thirds, which can be significantly reduced to 22% when using cross-layer integration (e.g., the *lost link* feature).

Although the RTT and RTT jitter are acceptably small for source-destination distances below 700 m, higher distances results in extreme fluctuations in RTT, e.g., up to 300% for 2000 m distance.

Finally, the reordering probability for a light network load is small (i.e., below 1%), but increases up to 15% for high-load scenarios. The number of consecutively reordered packets and, thus, the reordering period also depends on the network load.

The path characteristic evaluation results assist in the design of a transport protocol for VANETs. The statistical knowledge helps e.g., to estimate if a communication is in connected or disrupted state, supports the distinction between lost and delayed or reordered packets and influences the calculation of the retransmission timer. Note that the utilization of the statistical knowledge naturally involves the novel metric *source-destination distance* as integral part of a VTP. Furthermore, the results advise in basic design decisions, e.g., the packet loss probability results and distribution of losses and reorderings demand for a selective acknowledgment scheme. The following section provides a detailed description of VTP, including the relation to the path characteristic results within the respective mechanisms.

3.5 Vehicular Transport Protocol Specification

This section presents the goals, basic assumptions and key features of the VTP approach and describes the protocol functionality in detail. This includes a description of the transport layer mechanisms, a functional protocol description, state diagrams and the VTP header format. Before, we come back to the assumptions. Additionally to the assumptions of the path characteristics evaluation in Chapter 3.4, we now complete the assumptions and add more specific assumptions for VTP that include the path characteristics results.

3.5.1 VTP Basic Assumptions

This section summarizes the basic assumptions of the communication system for VANETs which represents the execution environment of the VTP protocol.

VANETs are mobile, wireless ad hoc networks that are characterized by a high degree of mobility where vehicles move along streets.

VANETs facilitate inter-vehicle and vehicle-to-roadside communication through wireless networks. Each participating vehicle is equipped with at least one wireless interface. The vehicles use IEEE 802.11 wireless LAN for direct (i.e., single-hop) message exchange between vehicles within each other's transmission range. Typically, the transmission range is 250 m, but it depends on environmental factors, such as obstacles.

VANETs enable multi-hop communication in a self-organized network connected by wireless links in an arbitrary topology. Ad hoc routing protocols provide multi-hop routing capabilities. The high degree of mobility in VANETs advises for position-based routing (PBR) [93] because PBR outperforms topology-based routing in this environment [44]. With PBR, each node selects the next reachable forwarder that is geographically closest to the destination. Thus, the paths that consecutive packets follow may be different due to mobility. VTP assumes a maximum source-destination distance of 2000 m (or eight hops) because the simulative evaluation in [82] shows insufficient expected connectivity duration beyond this distance.

PBR requires that each vehicle is able to determine its geographical position. Thus, each vehicle is additionally equipped with a positioning system, such as the global positioning system (GPS). The GPS also provides a synchronous clock.

In VTP, the protocol layers inside a network node are strongly coupled to enhance network performance. VTP nodes (i.e., sender, receiver and intermediate nodes) use lower-layer information to assist the transport layer. Examples include: *(i)* The inverse of the source-destination distance $1/d$ of a communication pair is part of the ACK and retransmission timer calculation formula. PBR provides the

geographical positions of source and destination for the distance calculation as a cross-layer service. *(ii)* Intermediate nodes estimate the available bandwidth in their vicinity via time measurement of the IEEE 802.11 MAC layer. The MAC layer measures the time between the point in time when a packet is ready to send and the MAC layer acknowledgment (i.e., including channel busy and contention time). VTP divides the packet size by the measured duration in order to estimate the available bandwidth.

According to the path characteristics evaluations, VTP aims at communication end points within a maximum hop-distance of about six to eight hops. Communication between vehicles that are more far away does not provide sufficient communication duration for useful communications. This restriction is in line with related performance measurement results of PBR, e.g., [96]. This assumption makes our transport protocol scalable.

The VTP design assumes similar network load conditions along a geographical path, i.e., a street. Thus, the network feedback about available bandwidth is similar, even if consecutive packets take different paths due to PBR routing. VANETs satisfy this requirement, assuming that the wireless transmission range of all communicating vehicles exceeds the width of the street. Thus, all vehicles can *sense* the data traffic in the vicinity of the respective street segment, as illustrated in Figure 3.14.

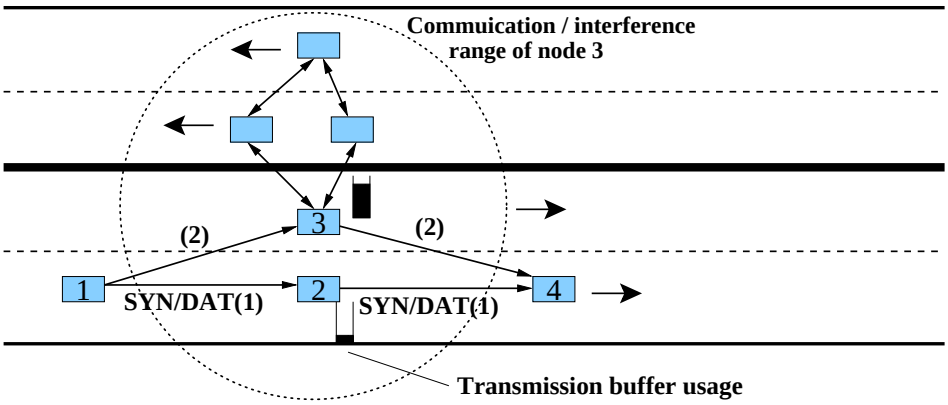

Figure 3.14: Similar network load condition assumption along streets.

Node 2 forwards the first syn/data packet (1) from the source node 1 to the destination node 4. Due to node movement, the routing protocol selects node 3 as forwarder for the consecutive packet (2). However, the feedback (i.e., available bandwidth) provided by node 2 and 3 is expected to be similar because both nodes sense all the data traffic in the area denoted by the circular communication range.

3.5.2 Goals, Key Features and Protocol Overview

The objectives of VTP include the establishment and release of an end-to-end connection, reliable delivery of data packets, in-sequence byte stream delivery of data to the application, flow and congestion control. The major design goals are to maximize the throughput of a connection, thereby preserving fairness among contending flows, and adapting the throughput of the different flows according to the minimum available bandwidth along the path, as it will be described in the following.

In order to achieve these goals, VTP must consider the characteristics of the highly dynamic, wireless vehicular environment, such as packet loss rate, end-to-end delay, delay jitter and reordering impact the transport layer [82]. VTP considers these metrics in its design choices.

In order to provide the services, as described above, VTP relies on the following key features:

- Rate-based transmission,
- Decoupling of error and congestion control,
- Congestion control via explicit signaling in the packet header,
- Selective acknowledgments that are sent in periodic intervals depending on the current transmission rate and the source-destination distance, and
- Use of statistical knowledge (e.g., expected communication/disruption duration) for rate calculation, error and congestion control.

The remainder of this section provides an overview of VTP.

VTP establishes a connection between a sender and a receiver. This connection is full-duplex, but the following description focuses on uni-directional traffic to simplify the explanation. A VTP connection can be either in a *connected* or *disrupted* state, indicating whether multi-hop connectivity exists or not.

In absence of acknowledgments, the VTP sender assesses whether a source-destination path exists or a disruption has occurred. The sender estimates the expected remaining communication duration based on the statistical knowledge of [82]. In case the statistical mean is below a threshold, the sender switches to *disrupted* state.

In *disrupted* state (e.g., the network is partitioned), the VTP sender only transmits periodic *probe* packets without data to detect path recovery. The rate of the probe packets is adjusted according to the statistically expected disruption duration for the given source-destination distance.

The arrival of acknowledgments indicates a *connected* state. In this state, the VTP sender steadily transmits packets at the maximum allowed data rate. This rate is determined by the feedback about available bandwidth along the path from intermediate nodes.

The congestion control of VTP uses explicit signaling of available bandwidth in the header of each data packet to adjust the transmission rate of the sender and avoid congestion. This mechanism decouples congestion control from error and flow control. The VTP sender inserts the minimum of its locally available or desired bandwidth in the VTP header of each packet. Intermediate nodes may reduce this value in the header: Each node measures and maintains the available bandwidth in its vicinity, as explained in Section 3.5.3.3. Furthermore, each intermediate node periodically collects the number of connections on which it forwards and their respective desired bandwidths. This allows an allocation of bandwidth to each flow according to the max-min fairness algorithm [12]. An intermediate node reduces the header bandwidth field in case it exceeds its share of the available bandwidth.

When the VTP packet arrives at the receiver, it contains the minimum of the available bandwidth along the multi-hop path. The VTP receiver maintains this information in a weighted average function to smooth fluctuations. It includes this weighted average of the available bandwidth in the acknowledgments.

The VTP receiver uses selective acknowledgments (SACKs) in order to efficiently report blocks of lost packets. Furthermore, the receiver transmits SACKs in dynamic intervals in order to wait for delayed or reordered packets and to reduce the contention on the wireless channel. The dynamic calculation of the SACK interval considers the current transmission rate and the source-destination distance.

VTP provides reliability via retransmission of lost packets whereas the SACKs inform the sender about received and lost packets. The VTP sender maintains a retransmission timer per connection. When packets are not acknowledged before the retransmission timer expires, the packets are considered lost and scheduled for retransmission. However, the retransmission timeout calculation cannot rely on the typical round trip time (RTT) measurements because of the extreme fluctuation of RTTs in VANETs [82]. Therefore, the VTP sender calculates the retransmission timeout out of the expected SACK interval, the current transmission rate and the source-destination distance.

One of the main objectives of VTP is the aggressive convergence to the maximum possible transmission rate of the connection (while preserving fairness to contending data traffic) to exploit even short connectivity periods. Therefore, VTP uses a *quick start* mechanism on connection establishment or after a disruption: The VTP sender transmits a *syn* packet to establish a connection or *probe* packets during a disruption to check when connectivity resumes. Upon reception of one of

these specific packets, the receiver replies immediately with an acknowledgment which contains the available bandwidth as collected by the *syn* or *probe* packet. Consequently, the VTP sender can initiate its transmission at the maximum possible rate after one RTT. Note that the syn packet already contains data to exploit particularly short connection periods.

Section 3.5.3 explains in detail the respective transport mechanisms of VTP from sender, source and intermediate node perspectives. Before, Section 3.5.1 summarizes the basic assumptions of the ad hoc communication system and the vehicular environment.

3.5.3 Transport Layer Mechanisms

This section explains the VTP transport layer mechanisms in detail. The key features of VTP include: Connectivity state management, rate-based transmissions, explicit congestion signaling and selective acknowledgments in dynamic intervals. Beyond, VTP uses statistical knowledge to predict certain path characteristics of a connection, including the source-destination distance as a metric.

3.5.3.1 Connectivity State Control

A VTP connection can either be in *connected* or in *disrupted* state.

The VTP connection is in *connected* state when steadily SACKs arrive at the sender before the *retransmission timer* expires. Typically, the retransmission timer is two times the acknowledgment interval plus an estimation of the RTT, as explained in detail below. In this state, the VTP sender adjusts its rate according to the feedback about the available bandwidth along the path from intermediate nodes, as contained in the SACKs.

In the absence of SACKs, the sender calculates the expected remaining connection duration, using statistical results for the given source-destination distance. In case the statistical mean is lower than a threshold, the VTP sender switches to *disrupted* state. In this state, the VTP sender stops the transmission of data packets. Instead it transmits periodic *probe* packets (i.e., control messages without data) to check if connectivity resumes. In the absence of data packets, the receiver also switches to *disrupted* state and reduces its acknowledgment rate equally to the rate of probe packets.

The arrival of an up-to-date acknowledgment triggers the transition from *disrupted* to *connected* state. The VTP sender resumes transmission at the maximum possible rate along the path, as contained in the acknowledgment.

3.5.3.2 Rate-Based Transmission

In *connected* state, the VTP sender uses a *rate-timer* to schedule the transmission of the next packet. This timer determines the transmission rate of the respective connection. The sender dynamically adjusts the rate (i.e., the timeout value) according to the available bandwidth along the path, as contained in the SACKs. Based on this network feedback, the VTP sender maintains its rate in three phases, as explained in the following. The available bandwidth feedback in the SACKs triggers the transition between these phases.

- Decrease phase. The VTP sender decreases its transmission rate when the available bandwidth signaled by the intermediate nodes is smaller than the current transmission rate. In this case, the sender immediately reduces its transmission rate to the available bandwidth in order to avoid congestion.

- Increase phase. The VTP sender increases its transmission rate when the available bandwidth signaled by the intermediate nodes is above the current transmission rate plus a threshold. The threshold intends to smooth out small bandwidth fluctuations due to different paths and enables a phase with constant transmission rate. If the signaled available bandwidth is above the threshold, the sender increases its transmission rate only by a fraction of the additionally available bandwidth. The increase is limited such that the sender does not overload the network. Each additionally injected packet uses the available bandwidth multiple times in a wireless forwarding chain when nodes within each other's transmission range forward the packet. Thus, the amount of increase depends on the number of forwarders within transmission range. The greedy forwarding strategy of PBR typically selects the forwarder within its transmission range that is geographically closest to the destination. Consequently, the packet consumes the bandwidth usually twice, resulting in a transmission rate increase at most half of the additionally available bandwidth.

- Constant phase. The VTP sender maintains its current rate when the available bandwidth signaled by the intermediate nodes is above the current transmission rate, but below the threshold.

In *disconnected* state, the VTP sender stops the transmission of data packets. Instead it transmits periodic *probe* packets in order to check if connectivity resumes. The interval of *probe* packets is less than the predicted disruption duration for the given source-destination distance of the statistical results. The VTP receiver continues to acknowledge the last state of received data. However, in *disrupted* state, the acknowledgment interval decreases equally to the interval of the *probe* packets.

3.5.3.3 Explicit Congestion Signaling

The congestion control of VTP uses explicit signaling of available bandwidth along the multi-hop path from intermediate nodes: A VTP sender inserts its locally available bandwidth or, if the connection requires less than the available bandwidth, the desired bandwidth in each VTP packet. Intermediate nodes along the multi-hop path that forward the packet may reduce this value according to their local network load conditions.

For this purpose, each node maintains an estimation of the currently available wireless bandwidth in its vicinity, as explained in detail in the remainder of this Section, and observes the utilization of its sending queue. Beyond, each intermediate node periodically collects the number of connections it forwards and their respective desired bandwidths. Thus, the intermediate nodes distribute the available bandwidth to the flows according to the max-min fairness algorithm [12]: Each flow gets the same share of the available bandwidth. Flows that request less than their share are fully served, and the remaining bandwidth is distributed among flows, requesting more than their share. When a flow requests more than its share, the intermediate node reduces the bandwidth field in the VTP header before forwarding the packet. Note that this mechanisms scale due to our assumptions that each nodes monitors only it's (single-hop) vicinity and communications will take place between nodes that are not more far away than six to eight hops (see assumptions 3.5.1). This *localization* of communication makes our algorithms scalable.

When the packet arrives at the VTP receiver, it contains the minimum available bandwidth along the path, which defines the maximum transmission rate for the sender. The VTP receiver accumulates the available bandwidth information in a weighted average function, as shown in Equation 3.1. The weighted average intends to smooth varying information of the available bandwidth in consecutive packets. These variations may occur since consecutive packets routed by PBR may arrive via different paths. However, such variations of available bandwidth information are small, due to the basic assumption of similar network load in street segments along the path.

$$BW_{av} = BW_{av\_prev} * \beta + BW_{pkt} * (1 - \beta) \tag{3.1}$$

The VTP receiver calculates the weighted average of the available bandwidth BW_{av} via the previous weighted average BW_{av_prev} and the actual value BW_{pkt}, as contained in the packet. The factor β determines the responsiveness to variations. It intends to smooth small fluctuations in subsequent packets.

The receiver includes the weighted average BW_{av} of the available bandwidth in the selective acknowledgment.

Wireless Bandwidth Measurement In order to provide explicit congestion feedback, each node must measure the current network load in its vicinity. Therefore, VTP adapts the IEEE 802.11 bandwidth measurement mechanism of [27] and [28]. This method measures the time *ts* when a packet is ready to send until the respective MAC acknowledgment is received at *tr*, as illustrated in Figure 3.15.

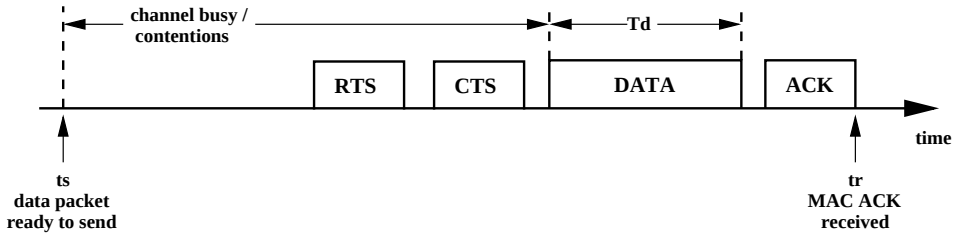

Figure 3.15: IEEE802.11b timestamps for wireless bandwidth measurement.

The IEEE 802.11 MAC protocol uses *carrier sense multiple access with collision avoidance (CSMA/CA)* [31] to coordinate the transmission of packets in a distributed coordination function (DCF). A node that prepares a transmission first senses if the channel is idle. If so, this sender issues a *request-to-send (RTS)* message or, otherwise, it backs off for a random interval before sensing the channel again. Upon reception of a RTS, the receiver replies with a *clear-to-send (CTS)* message in order to allow the transmission and silence all other nodes within transmission range. When receiving the CTS, the sender transmits the actual data packet which is confirmed by the receiver with an MAC layer ACK.

The measured interval includes the channel-busy time, contentions, the DCF time and the duration of the actual transmission. According to [72], [1], the resulting, actual throughput can be calculated, as shown in Equation 3.2.

$$TP = \frac{S}{(t_s - t_r)} \tag{3.2}$$

The variable TP represents the measured throughput, S represents the packet size and t_s and t_r are the timestamps, as explained above.

[27], [28] further generalize the throughput calculation formula to a normalized representation (i.e., normalized to arbitrary reference packet sizes). However, Equation 3.2 is adequate for the following simulative evaluation because the simulation scenario considers fixed packet sizes.

3.5.3.4 Selective Acknowledgments in Dynamic Intervals

The error control of VTP uses selective acknowledgments (SACKs) to provide reliability. SACKs combine negative and positive acknowledgments by confirming received packets and reporting missing blocks in between. A *cumulative ACK* field

provides the maximum segment number up to which all packets were continuously received. The *SACK blocks* report missing packets between successfully received packets. This covers single packet loss as well as bulks of consecutively lost packets.

The SACK blocks in VTP are options attached to the VTP acknowledgments. The maximum number of SACK blocks that a VTP receiver reports in one single ACK is restricted in order to respect the maximum packet size and keep the ACKs small. Note that further information of missing packets that did not fit into a SACK will be included in the subsequent SACK. The maximum number of SACK blocks depend on the scenario, as shown by the following example:

For the theoretical data channel rate of 2 Mbps and a packet size (i.e., maximum transfer unit MTU) of 1500 bytes, the sender transmits a maximum of 166 packets per second. Assuming a typical ACK interval of 0.25 s, the sender transmits up to 41 packets within a single interval. The path characteristics analysis [82] shows a loss probability of 61% at a 2000 m source-destination distance, which is the maximum distance in the VTP assumptions. This leads to an expected loss of 25 packets. Note that the loss probability includes single as well as block losses. SACK reports consecutive losses in one block. Consequently, a maximum number of 25 SACK blocks is useful in this scenario. In case there are more blocks of packets lost than the maximum number of SACK blocks within one acknowledgment, VTP always reports the first losses in the sequence of the data stream.

The VTP receiver acknowledges a connection establishment packet or a probe packet (i.e., after a disruption) immediately. Beyond, the receiver sends SACKs in dynamic intervals because *(i)* to acknowledge each received packet separately would increase contention and network load and *(ii)* this scheme accounts for delayed or reordered packets. The calculation of the SACK interval is based on the actual transmission rate and the source-destination distance, as shown in Equation 3.3.

$$ack\_t = \frac{1}{trans\_rate_{av}} * \frac{K}{d^n} \tag{3.3}$$

The acknowledgment timeout ack_t is calculated by the inverse of the average transmission rate $trans\_rate_{av}$ (i.e., inter-packet delay), a scaling constant K and the distance d. The exponent n expresses an exponential impact of the distance. It is chosen as 0.5 to account for the erratic increase of loss probability in short distances (i.e., below 500 m) and the linear increase beyond, as evaluated in [82].

Equation 3.3 ensures that the SACK interval linearly increases with the transmission rate, but exponentially decreases with the distance.

Similar to the transmission rate, the VTP instance maintains the weighted average of the SACK timer, as shown in Equation 3.4. Again, the factor β determines the responsiveness to variations and will be adjusted through a simulative study.

$$ack\_t_{av} = ack\_t_{av\_prev} * \beta + ack\_t_{curr} * (1 - \beta) \tag{3.4}$$

The SACKs inform the VTP sender about losses and successfully received packets in discrete time intervals, according to the ACK-timer above. This ACK information impacts the retransmissions: When a packet is not acknowledged before the retransmission timer expires, it is considered lost and scheduled for retransmission. Thus, the retransmission timer is responsible to distinguish between lost, delayed or reordered packets on the sender side. A lost packet should be retransmitted as fast as possible, but superfluous repetition of delayed or reordered packets should be avoided.

However, the *traditional* method of retransmission timer calculation that relies purely on the measured RTT and RTT jitter is not appropriate in VTP for the following reasons: *(i)* The statistical results in [82] show that depending of the source-destination distance RTT fluctuation up to 300% occur frequently. *(ii)* The periodic acknowledgments in discrete time intervals circumvent a *per-packet* RTT measurement. The RTT measurement *per-ACK* results in decreased accuracy. *(iii)* VTP estimates the RTT and jitter by duplicating the (one-way) transmission delay of periodic acknowledgments which in imprecise due to the asymmetry on the forward and reverse path.

The VTP retransmission timer considers the actual ACK interval and an estimation of the RTT by comparing measured RTT and statistically expected RTT. In order to estimate the current RTT, the VTP sender also (i.e., like the VTP receiver) maintains an acknowledgment timer, as shown in Equations 3.3 and 3.4. The retransmission timeout waits at least one ACK interval plus the estimated RTT. When the sender has continuous data ready for transmission, the retransmission timer may wait for two SACKs plus the estimated RTT in order to account for delayed packets that are acknowledged in the subsequent SACK. However, the retransmission timer should not consider more than two ACK intervals because the increasing probability of a disruption may prevent retransmissions [82]. The number of ACK intervals which are considered in the retransmission timer calculation depends on the scenario and might use the source-destination distance as metric.

Furthermore, the sender estimates the RTT by comparing RTT measurements and statistics. The sender measures the transmission delay and maintains a smoothed RTT and RTT variation average value, similar to TCP. Upon reception of the first acknowledgment, the sender initializes the RTT_{av} with the absolute measured value according to the timestamp in the packet and the $RTTVAR_{av}$ according to the statistically expected variance for the given source-destination distance. Subsequent measurements maintain the RTT_{av} and $RTTVAR_{av}$, as shown in Equation 3.5.

$$\begin{aligned} RTTVAR_{av} &= (1-\beta) * RTTVAR_{av\_prev} + \beta * |RTT_{av} - RTT_{meas}| \\ RTT_{av} &= (1-\alpha) * RTT_{av\_prev} + \beta * RTT_{meas} \end{aligned} \tag{3.5}$$

Equation 3.6 shows the complete VTP retransmission timeout calculation.

$$\begin{aligned} retrans\_t &= N * ack\_t_{av} + max(RTT_{meas}, RTT_{stat}) \\ & \quad with \\ & \quad RTT_{meas} = RTT_{meas\_av} + RTTVAR_{meas\_av} \\ & \quad RTT_{stat} = RTT_{stat\_av} + RTTVAR_{stat\_av} \end{aligned} \tag{3.6}$$

The factor N is either one or two and determines the number of acknowledgments to be considered in the retransmission timeout calculation depending on the scenario, e.g., the source destination distance. In addition to the average ACK interval, the calculation adds the maximum of the measured and the statistical RTT for the given source-destination distance. This algorithm considers the ACK interval and the RTT in order to distinguish between delayed and lost packets.

3.5.3.5 Fairness

In VTP, fairness means that an intermediate node equally distributes the available bandwidth to contending flows that are routed via this node, according the *max-min* fairness algorithm [12].

In order to fairly distribute the available bandwidth, each intermediate node must *(i)* measure the available bandwidth in its vicinity. Section 3.5.3.3 describes this measurement of available wireless bandwidth in detail. *(ii)* Intermediate nodes require to know the number of flows and their respective bandwidth demands. However, the nodes cannot maintain a per-flow state, e.g., because path changes occur frequently due to mobility. Therefore, each VTP node periodically collects the number of flows and their bandwidth demands. The period depends on the de-

gree of node mobility which results in topology changes. The high degree of node mobility in VANETs (e.g., high relative speed of opposing traffic) requires a small period. A simulative study will determine the period and optimize it for specific scenarios, such as highways, in our future work.

3.5.3.6 Connection Management

VTP sender and receiver establish an end-to-end connection. This connection is full-duplex to allow request-response actions and exchange data in both directions. However, the following description focuses exemplary on uni-directional data traffic for simplicity. Both end points maintain a reliability scoreboard of successfully received and lost packets, retransmission and acknowledgment timer per connection.

The VTP connection establishment and termination use three-way handshakes. However, an explicit connection establishment takes time and introduces overhead, which is particularly problematic for short-lived connections of wireless, mobile VANETs. Therefore, VTP includes data in the connection establishment packet. Optionally, in the failure case where the connection establishment packet is not acknowledged before the retransmission timer expires, subsequent connection re-establishment might not contain data. This option reduces the network load when the probability that a path is disrupted is high. Thus, it avoids waste of scarce wireless bandwidth.

3.5.4 Functional Protocol Description

This section explains the protocol functionalities in detail and illustrates the message and information exchange of the participating entities (i.e., sender, receiver and intermediate node(s)) related to the transport mechanisms.

3.5.4.1 Connection Establishment

Figure 3.16 illustrates the signaling of a successful and an erroneous VTP connection establishment via 3-way handshake.

Figure 3.16(a) illustrates the successful VTP connection establishment. The sender transmits a syn packet, which includes data, to establish the connection and invokes a retransmission timer. When the respective acknowledgment and reverse-syn arrives before the retransmission timer expires, the connection is established: The sender re-schedules the retransmission timer, acknowledges the connection establishment request from the destination and starts the transmission of data at the allowed transmission rate along the path, as included in the acknowledgment.

In Figure 3.16(b), the first syn packet is lost and, thus, the connection establishment is erroneous. The retransmission timer expires since no acknowledgment arrives within the expected interval. Upon retransmission timer expiration, the sender

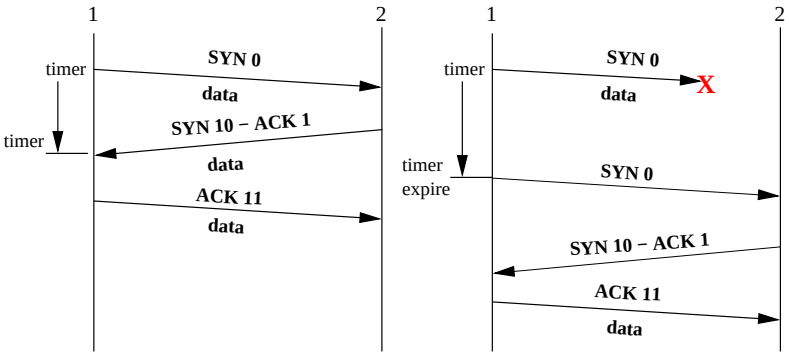

(a) Successful. (b) Erroneous on first SYN.

Figure 3.16: VTP 3-way-handshake connection establishment.

retransmits the syn packet, reschedules the retransmission timer and increases the retransmission counter. This counter restricts the maximum number of retries before the connection establishment terminates with an error. The Figure illustrates the option that in a failure case the retransmission of a connection establishment packet does not contain data. In case the connection establishment packet is lost, the sender assumes with a high probability that the path is disrupted. The subsequent connection establishment packet does not contain data in order to preserve wireless bandwidth. When the sender receives an acknowledgment, the connection establishment continues like in the successful case, as described above.

3.5.4.2 Reliability

VTP provides reliability via retransmissions of lost packets. In order to identify lost packets, the VTP receiver transmits selective acknowledgments in dynamic intervals. The VTP end-systems maintain a *scoreboard* of acknowledged and unacknowledged packets. The VTP sender stores transmitted but not yet acknowledged packets. The VTP receiver keeps track of successfully received packets and transmits the selective acknowledgments according to its scoreboard, as explained in Section 3.5.3.4.

Upon arrival of SACKs before the retransmission timeout, the VTP sender removes the acknowledged packets from its scoreboard and re-schedules the retransmission timer. When the retransmission timeout occurs before the respective packets are acknowledged, the sender considers these packets as lost and schedules them for retransmission. Packets in the *retransmission queue* are scheduled with a higher priority than *new* packets. Thus, when the rate timer expires and the retransmission queue contains packets, these packets are transmitted first. A

VTP connection must not close before all packets are transmitted and acknowledged or terminate with an error, e.g., in case the connection does not resume after a disruption.

Figure 3.17 illustrates the end-to-end error control of VTP for the cases *(i)* when a packet is delayed (11) and *(ii)* when a packet is lost (13) by intermediate nodes in the network.

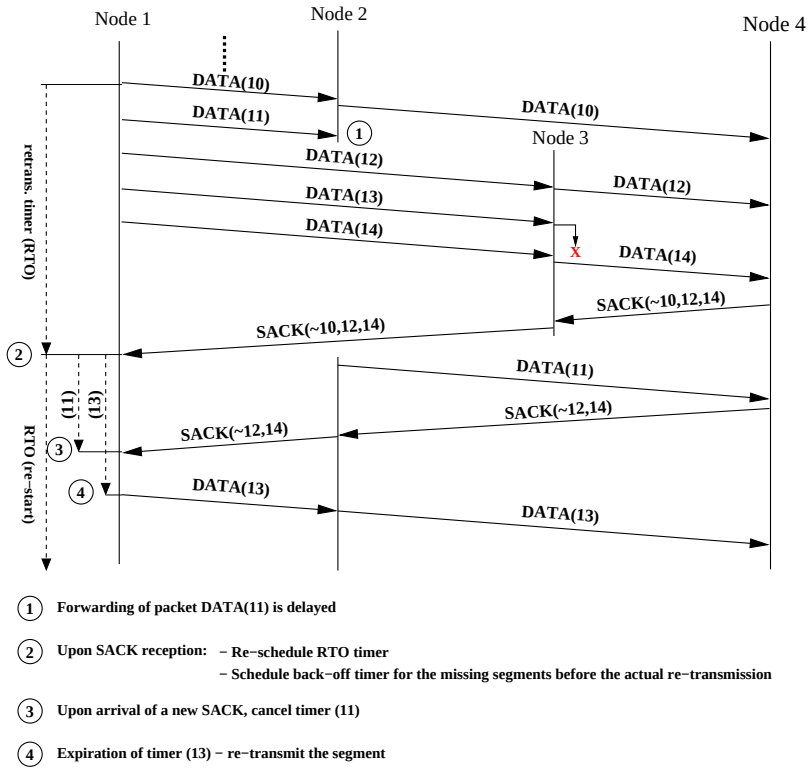

Figure 3.17: VTP error control.

During the acknowledgment timer interval, the receiver maintains received and missing packets in its scoreboard, i.e., in the first interval, it receives the packets (10), (12) and (14). Upon acknowledgment timeout, the receiver sends a selective acknowledgment, which includes a cumulative and selective acknowledgment field, as follows. The cumulative field in the first SACK in Figure 3.17 informs the sender that the receiver successfully and continuously received all packets up to packet number (10). Beyond, the SACK option reports the reception of packet (12) and (14), i.e., this implies that the packets (11) and (13) are missing. Particularly in situations with high data traffic, selective acknowledgments provide important information for efficient retransmission of the lost packets only.

When the first SACK arrives at the VTP sender, it re-schedules the retransmission timer and schedules the missing packets for retransmission. In the mean time,

the receiver gets the delayed packet (11). The receiver includes this delayed packet in the subsequent SACK upon the next acknowledgment timeout. In the second SACK, the cumulative ACK field includes all packets up to number (12) and the SACK option reports the reception of packet number (14), i.e., the only missing packet is number (13).

When the sender learns from the second SACK that the missing packet (11) successfully arrived at the receiver, the sender removes packet (11) from the retransmission queue. However, packet (13) is still missing after two acknowledgment intervals (inclusive the RTT) and, thus, the sender retransmits packet (13), which is considered as lost.

3.5.4.3 Congestion Control

In the *connected* state, the VTP sender adapts its transmission rate according to the available bandwidth along the path, as explicitly signaled by the intermediate nodes.

Each node measures and maintains the available bandwidth in its vicinity, as explained in Section 3.5.3.3. The sender inserts its available bandwidth in the VTP packet. Intermediate nodes may reduce the bandwidth in the VTP packet, in case the available bandwidth measured by the respective node is less than the value in the packet or the fill status of the transmission queue of the intermediate node is above a threshold. That implies that VTP is installed on all nodes of the ad hoc network and the intermediate nodes support the end-to-end VTP connection with their bandwidth information. Procedure 1 shows the comparison of available bandwidth on intermediate nodes.

Procedure 1 Available and requested throughput comparison on the intermediate node.

if ($requested\_tp < available\_tp$) **then**
  Forward the packet without modification
end if
if (($requested\_tp > available\_tp$) **then**
  Update the throughput header field and forward the packet
end if

When a VTP packet arrives at the receiver, it contains the minimum of the available bandwidth along the path. The VTP receiver accumulates and maintains the available bandwidth in a weighted average function, as explained in Section 3.5.3.3. The VTP receiver inserts the weighted average of the available band-

width in the SACK. Upon reception of a SACK, the VTP sender increases, decreases or maintains its transmission rate. The transmission rate is the inverse of the inter-packet delay. The inter-packet delay is the inverse of the the available bandwidth, as shown in Equation 3.7.

$$inter - packet - delay = \frac{1}{BW} \tag{3.7}$$

As indicated before, the VTP congestion control uses three phases to adjust the sender's transmission rate: Decrease, increase or a constant transmission rate. Procedure 2 shows this mechanism.

Procedure 2 Congestion control by transmission rate adjustment.

if ($new\_rate < current\_rate$) **then**
  $current\_rate = new\_rate$ // Decrease transmission rate
end if
if ($new\_rate > (current\_rate - \delta)$) **then**
  $current\_rate = current\_rate - \frac{(current\_rate - new\_rate)}{k}$ // Increase transmission rate partially
end if

When the available bandwidth along the path, as reported by the intermediate nodes, is lower than the current transmission rate, the VTP sender immediately decreases its rate. The rate is adapted to the conditions along the path, i.e., the transmission rate is set equal to the available bandwidth, as shown in the first if-statement in Procedure 2.

When the available bandwidth along the path is more than the current rate, but less than the threshold δ, the sender keeps the current rate. The threshold δ enables a stable transmission rate by avoiding fluctuations due to frequent, small variations in the available bandwidth.

When the available bandwidth along the path is greater than the current rate plus δ, the sender increases its transmission rate. However, the increase is only a fraction of the additional available bandwidth because the rate increase amplifies with the forwarding of additional packets by the nodes within transmission range (i.e., induced traffic). This increase state is shown by the second if-statement in Procedure 2. The factor k determines the amount of rate increase. The factor can be constantly set to four or five according to the theoretical maximum nodes within transmission range or it can be dynamically calculated according to the distance in hops over the radio transmission range.

In absence of acknowledgments (e.g., due to temporal network partitions or congestion), the VTP receiver estimates the expected remaining connection duration based on statistical results. In case this remaining time is lower than a threshold, the VTP sender switches to *disrupted* state. In this state, the VTP sender stops transmission of data and periodically transmits *probe* packets in order to check when connectivity resumes. The interval of probe packets is determined by the statistical disruption duration for the given source destination distance. Figure 3.18 illustrates the transition from *connected* to *disrupted* state and the respective transmission of probe packets.

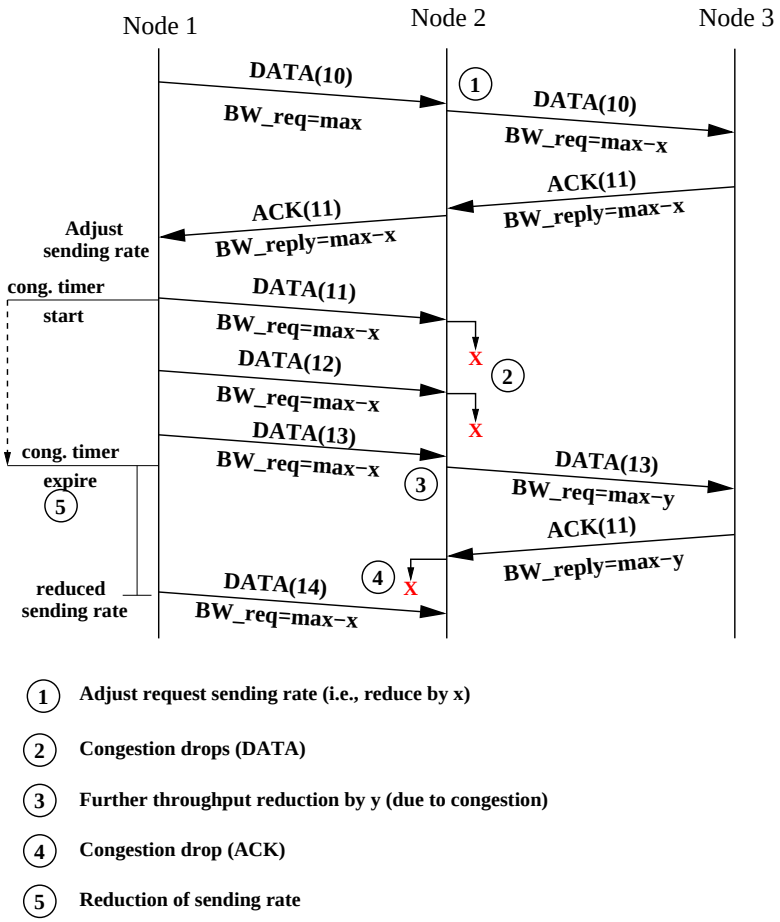

Figure 3.18: VTP connection state management in the absence of acks.

The arrival of an acknowledgment indicates resumed connectivity or resolved congestion. Thus, the VTP sender switches from *disrupted* to *connected* state. The VTP sender resumes its data transmission at the available rate, as indicated in the acknowledgment. Figure 3.19 shows exemplary this state transition and the VTP recovery after network congestion.

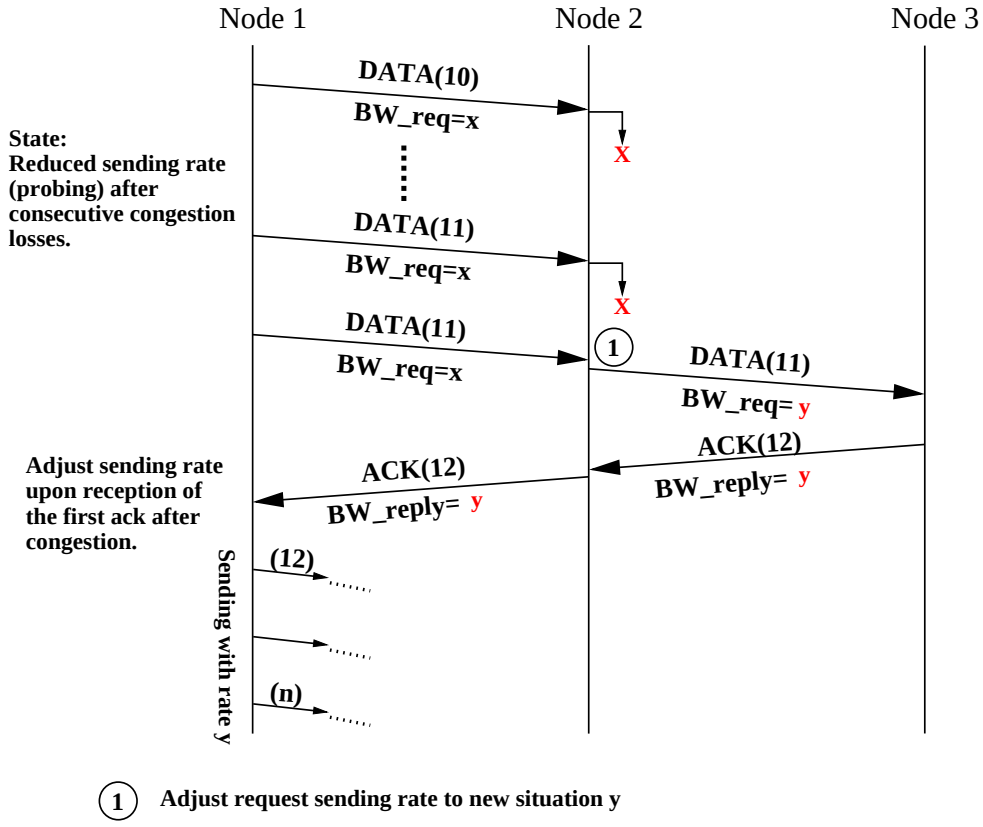

Figure 3.19: VTP recovery after congestion or network partition.

3.5.4.4 Flow Control

The VTP receiver informs the sender about its remaining receive-buffer capacity in order to avoid buffer overflows. The VTP receiver advertises its available receive-buffer size in the flow feedback of the acknowledgment.

The VTP sender must respect the flow control window and must not transmit more packet than the buffer of the receiver can accept. This restriction is independent of error or congestion control. The receiver might, e.g., temporarily stop the transmission of packets although bandwidth is available.

3.5.5 VTP State Transition Diagrams

This Section explains the VTP mechanisms inside a node via state transition diagrams for the VTP sender and receiver separately.

3.5.6 VTP Sender State Transition Diagram

Figure 3.20 shows the VTP sender state transition diagram.

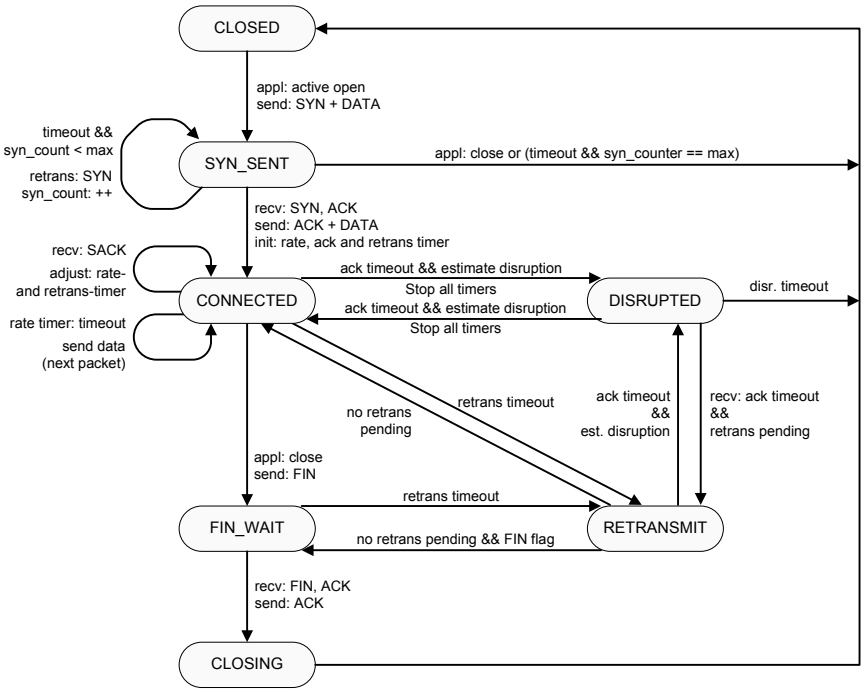

Figure 3.20: VTP sender state transition diagram.

The states are:

CLOSED represents the state when no connection is established.

SYN_SENT state waits for a connection confirmation after the sender has transmitted the connection request *SYN*.

CONNECTED is reached after a successful 3-way-handshake. In this state, the sender transmits a data packet on rate timer expiration. Thus, the rate timer determines the transmission rate. The sender maintains the rate timer based on network feedback, as contained in SACKS.

DISRUPTED represents a state when no connectivity between source and destination exists. In absence of acknowledgments, the sender predicts the remaining connectivity duration based on statistical knowledge. In case, the expected remaining connectivity is below a threshold, the sender switches to disrupted state. In disrupted state, the sender does not transmit data packets (i.e., it cancels all timers). Instead, it periodically transmits probe packets in order to determine when connectivity resumes.

RETRANSMIT state retransmits lost packets. The sender considers a packet as lost when it is not acknowledged before retransmission timer expiration. In this case, the packet is scheduled for retransmission. Upon rate timer expiration, lost packets are retransmitted instead of *new* data packets in this state.

FIN_WAIT represents the state when the application closes the connection, but packets in transit need to be delivered or acknowledged. The connection must not close before all packets are received and acknowledged.

CLOSING is reached when all packets are exchanged. This state resets all timers and variables of the respective connection.

3.5.7 VTP Receiver State Transition Diagram

Figure 3.21 shows the VTP receiver state transition diagram.

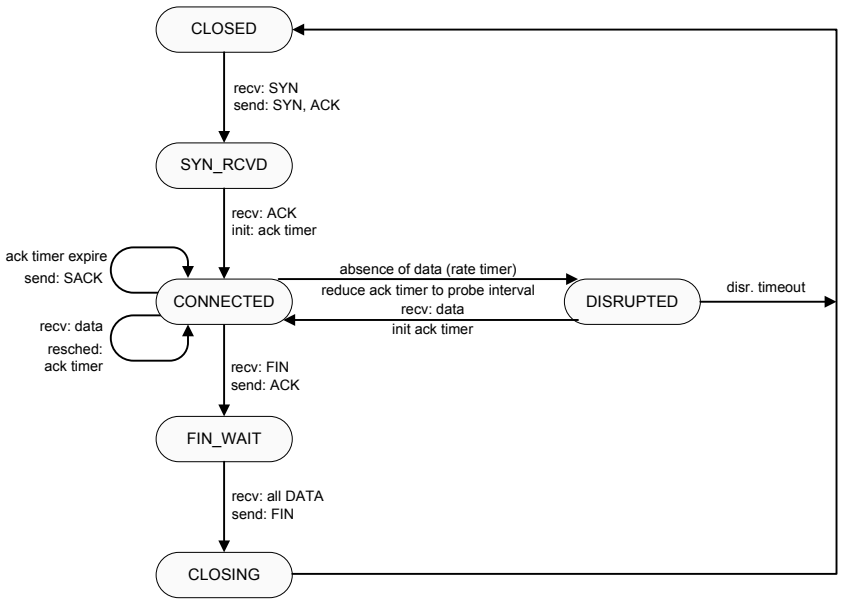

Figure 3.21: VTP receiver state transition diagram.

The states are of the diagram in Figure 3.21 are:

CLOSED represents the state when no connection is established. In this state, the receiver waits for connection requests from remote VTP peers.

SYN_RCVD represents the state after reception of a SYN packet. The receiver initializes the variables, confirms the connection establishment and waits for the respective acknowledgment to confirm the connection.

CONNECTED represents the state when a connection is established. The receiver maintains successfully received packets in a scoreboard and transmits SACKs in dynamic intervals.

DISRUPTED represents a state when no connectivity between source and destination exists. In absence of data packets, the receiver reduces the ACK interval, similar to the probe packet interval.

CLOSE_WAIT represents the state when the peer (i.e., sender) closes the connection, but outstanding packets are still in transit or missing. The connection must not close before all packets are received and acknowledged.

CLOSING is reached when all packets are exchanged. This state resets all timers and variables of the respective connection.

3.5.8 VTP Header Format

The only explicit VTP signaling is required for connection establishment and termination. These control messages utilize the same VTP header as all data packets, as shown in Figure 3.22. The VTP header includes all the relevant transport layer information.

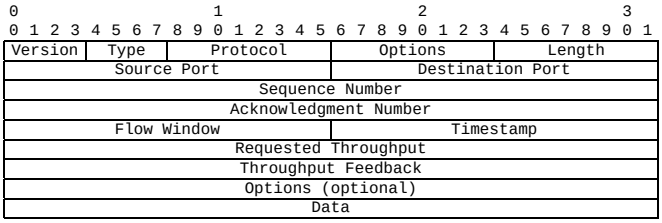

Figure 3.22: VTP header format.

Version: 4 bits This field specifies the used version of VTP.

Type: 4 bits The type field contains a code to identify the content of the message, such as connection establishment or termination, data, acknowledgment or the combination of data and acknowledgment in a duplex connection.

Protocol: 8 bits This field indicates the next-level protocol used in the data part of the packet.

Options: 8 bits This field defines if options are appended to the header and specifies the type of options in case.

Length: 8 bits This field contains the length of the VTP header in bytes. The length of the VTP header is variable because of options, such as selective acknowledgment blocks. The minimum, fixed part of the VTP header is 28 bytes.

Source Port: 16 bits The source port number identifies the connection.

Destination Port: 16 bits The destination port number identifies the connection.

Sequence Number: 32 bits The sequence number identifies the portion of the data flow contained in the packet. The number represents the first data octet in the segment. When the type field indicates that the packet establishes a connection, the sequence number specifies the beginning of the data flow.

Acknowledgment Number: 32 bits The acknowledgment number represents the cumulative acknowledgment and indicates up to which segments all packets are consecutively received, i.e., which packet is expected to arrive next.

Flow Window: 16 bits The flow window advertises the remaining capacity of the receiver's buffer. Independent of the transmission rate, the VTP sender must not transmit more data than the receiver can accept in order to avoid buffer overflow at the end-point.

Timestamp: 16 bits This field indicates the sending time of a packet (i.e., data or acknowledgment). The end-points use this field for RTT estimation. As mentioned in the assumptions, each vehicle is equipped with GPS that provides a synchronous clock in the ad hoc network, as required for the RTT estimation.

Requested Throughput: 32 bits This field specifies the throughput requested by the sender. Intermediate node access and adjust this field to adjust the available bandwidth along the path.

Throughput Feedback: 32 bits This field contains the accumulated average throughput replied by a VTP receiver in the acknowledgment.

Options: Variable Options may be appended at the end of the fixed part of the VTP header, as indicated in the options and length fields. All options are included in the checksum. The currently defined options are shown below.

VTP Header Options

Currently, the only VTP options are the selective acknowledgment blocks that are appended to an acknowledgment in case the VTP receiver identifies non-contiguous packets in the flow. The missing blocks are identified by the left and right edges of the missing segment as shown in Figure 3.23.

The maximum number of SACK blocks is 20. In case the receiver identifies more than 20 non-contiguous blocks, it reports always the first 20 blocks in the flow in the selective acknowledgment.

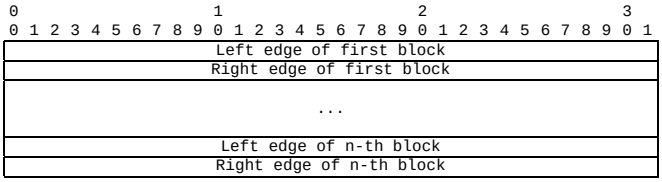

Figure 3.23: VTP header selective acknowledgment option.

3.6 Simulative Evaluation

This section evaluates the performance of VTP through simulations and compares VTP and TCP performance, such as throughput and fairness, in static and mobile wireless environments. The following subsections describe the simulation scenario, explain the metrics and present the simulation results.

3.6.1 Scenario and Simulation Environment

The simulation environment of the transport protocol evaluation in wireless and mobile environments is equal to the simulation environment of the path characteristics analysis, as described in Section 3.4. Thus, the remainder of this section briefly repeats and summarizes the simulation scenario and settings.

The evaluation uses the network simulator ns-2 [132] and comprises static and mobile scenarios.

In the static scenarios, source and destination nodes are positioned in predefined distances. The mobile scenarios consider moving vehicles on a highway, according to the movement pattern of [45]. These patterns represent realistic unidirectional mobility patterns that are typical on German highways and can be combined to bidirectional scenarios via the *hwgui* tool [80]. They include the spatial distribution of vehicles on the highway for different densities of vehicles, i.e., different number of lanes or different average number of vehicles per kilometer and lane. The simulations consider a 10 km stretch of highway.

All vehicles are equipped with a wireless IEEE 802.11b wireless interface that covers a radio transmission range of 250 m. Vehicles that are closer to each other than the radio range can communicate directly. The vehicles form an ad hoc network that enables multi-hop communication when the distance between source and destination exceeds the radio range. In our simulations we employ PBR as routing and either VTP or TCP as transport protocols.

In the static scenario, the communication pairs are positioned in predefined distances. Vehicles in between provide a multi-hop forwarding chain. The respective communication pair establishes one or two simultaneous data flows, e.g., to evaluate throughput or fairness.

In order to evaluate the transition between the *connected* and *disrupted* states in VTP, a connection is established in a static scenario where the connection is tem-

porally disrupted. This disruption is caused by the movement of a single node, i.e., the receiver moves out the radio coverage area of its predecessor for the duration of the disruption.

In the mobile scenario, communication pairs are randomly chosen throughout the simulation area. The evaluation classifies the results according to specific distances. The simulations compare the VTP and TCP transport protocols, transmitting continuously data at the maximum possible rate. Again, the simulations include oncoming traffic in order to reduce temporal network partitions.

3.6.2 Metrics

VTP aims at maximizing the throughput per connection with reliable and in-order delivery of data to applications, including flow and congestion control. However, the maximum throughput of a connection includes preserving fairness to contending data traffic. Consequently, the simulative evaluation of VTP considers the metrics *throughput* and *fairness*, as defined in the following.

- The *throughput* measures the bits per second as transmitted by the sender.
- The *fairness* evaluation observes the throughput of contending flows of simultaneous connections.

Furthermore, the simulations evaluate the reliable transmission of packets in time-sequence graphs, reflecting the sequence number of transmitted packets over time.

3.6.3 Simulation Results

This section presents the simulation results of the VTP performance evaluation and compares the results to TCP performance. The following subsections classify the results in static and mobile environments. The evaluation in static environments comprise throughput and fairness results and shows the VTP transition between connected and disrupted state when the connection is temporally disrupted. The mobile scenarios select randomly a communication pair per simulation run. The results show and compare the average throughput of the runs, classified according to the maximum distance of the communication pairs.

Again, the section provides detailed explanations for selected scenarios. The complete results are attached in Appendix A.

3.6.3.1 Performance Evaluation in Static Environments

The performance evaluations in static environments include throughput and fairness evaluations for single-hop and multi-hop communication for one and two

flows, respectively. These simulations compare the VTP performance to TCP performance as reference. Furthermore, the transition from connected to disrupted state and vice versa in VTP is shown upon the occurrence of a disruption.

Performance Evaluation in Static Environments with One Data Flow

This evaluation compares the throughput of VTP and TCP in static single-hop and multi-hop environments.

Figure 3.24 shows the VTP throughput over time in a single-hop scenario with two nodes located within each other's radio transmission range (i.e., in a distance of 250 m).

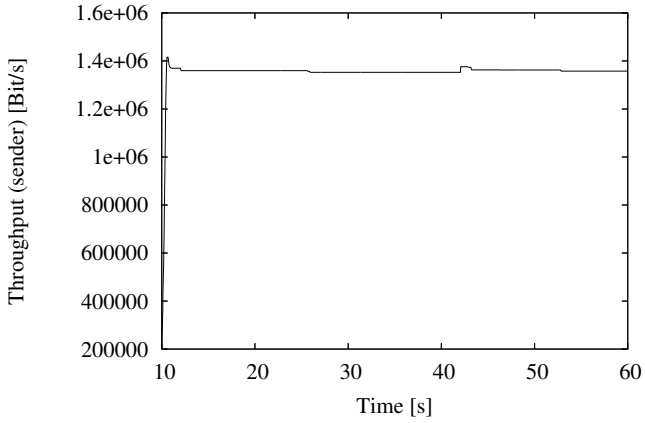

Figure 3.24: VTP throughput over time for one data flow in a static single-hop scenario with 250 m source destination distance and ($k = 3$,$\delta = 0.05$).

Upon connection establishment at 10 s, the transmission rate aggressively converges to the maximum available throughput of approximately 1.4 MBit/s and maintains this throughput almost statically for the complete simulation duration. The result shows exemplary an optimal VTP setting of δ = 0.05, representing an increase threshold of 20%, and a rate of increase of 1/3 of the additionally available bandwidth (see specification in Section 3.5.4.3). A variation of these settings results in fluctuations either around the maximum throughput or in an overall throughput decrease, as shown in Appendix A.

In contrast, TCP cannot maintain a stable throughput even in this static single-hop environment. Figure 3.25 illustrates the congestion window in 3.25(a) and throughput in 3.25(b) of TCP in this scenario. These results are in line with the results in related works, such as [43, 58, 48].

Due to round trip time variations in the wireless environment, the congestion

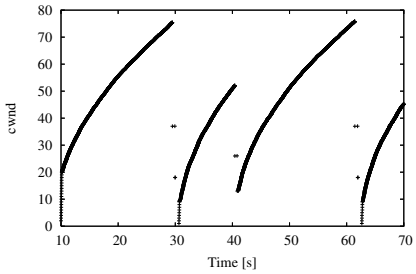

(a) TCP congestion window in a static single-hop scenario.

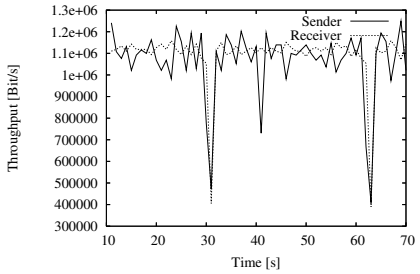

(b) TCP throughput in a static single-hop scenario.

Figure 3.25: TCP congestion window and throughput over time in a static single-hop scenario.

window fluctuates continuously. The congestion window even decreases to zero (e.g., at 31 s and 63 s), resulting in significant fluctuations and drops in the throughput over time.

Furthermore, VTP utilizes the wireless bandwidth more efficiently than TCP. VTP maintains an almost constant throughput of 1.4 MBit/s, whereas TCP reaches an average throughput of 1.12 MBit/s. Besides the efficient congestion control via network feedback, another reason for the better performance of VTP is the cumulative acknowledgment scheme. Figure 3.26 compares the acknowledgment numbers of VTP and TCP in a clipping of one second.

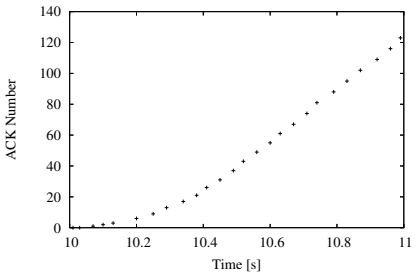

(a) VTP acknowledgment time-sequence graph.

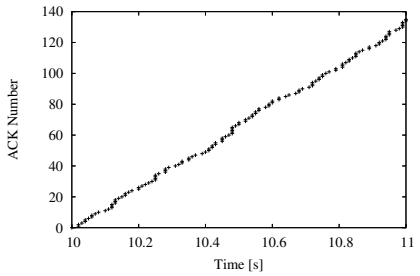

(b) TCP acknowledgment time-sequence graph.

Figure 3.26: VTP and TCP acknowledgment time-sequence graph (clipping).

The comparison of VTP acknowledgments in Figure 3.26(a) and TCP acknowledgments in Figure 3.26(b) shows that the cumulative acknowledgment scheme of VTP results in a decreased number of acknowledgments. Consequently, the wireless channel is less loaded and contention decreases. VTP uses the additionally available bandwidth for the throughput of data.

The throughput results, as explained in detail for the single-hop scenario above, are also valid for multi-hop connections. The complete results up to six hops are available in Appendix A. As an example, Figure 3.27 shows exemplary the VTP throughput over time for three hops and two different k - δ combinations.

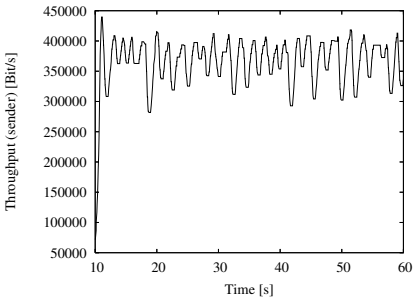

(a) VTP throughput for one data flow in a static three-hop scenario for k=3 and δ=0.05.

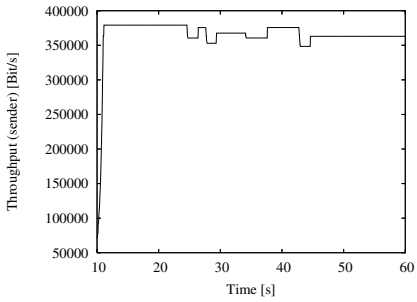

(b) VTP throughput for one data flow in a static three-hop scenario for k=5 and δ=0.2.

Figure 3.27: VTP throughput over time for one data flow in a static three-hop scenario for different k - δ combinations.

As mentioned before, k and δ determine the maximum throughput and the level of fluctuations. Figure 3.27(a) shows the throughput over three hops for k = 3 and δ = 0.05. This setting achieves an average throughput of 403 kBit/s, compared to an average throughput of 376 kBit/s for k =5 and δ = 0.2, as shown in Figure 3.27(b). However, the latter setting avoids the small fluctuations of the first scenario.

In contrast, Figure 3.28 shows the throughput of TCP over three hops in a static scenario.

Similar to the single-hop scenario, the TCP throughput fluctuates continuously fluctuates around the average of 355 kBit/s, which is significantly lower than the VTP average throughput.

Summarizing, unlike TCP, VTP provides stable throughput for single-hop and multi-hop connections.

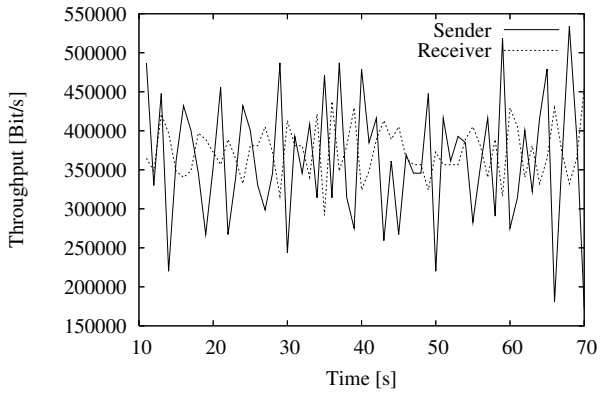

Figure 3.28: TCP throughput over time for one data flow in a static three-hop scenario.

Performance Evaluation with Disruption

This section evaluates the transition from *connected* to *disrupted* state and vice versa in VTP. The scenario is basically static, but to emulate a disruption the receiver moves out of transmission range and back between 15 s and 20 s. Figure 3.29 shows the throughput upon occurrence of a disruption in a single-hop scenario with the communication pair at a distance of 250 m.

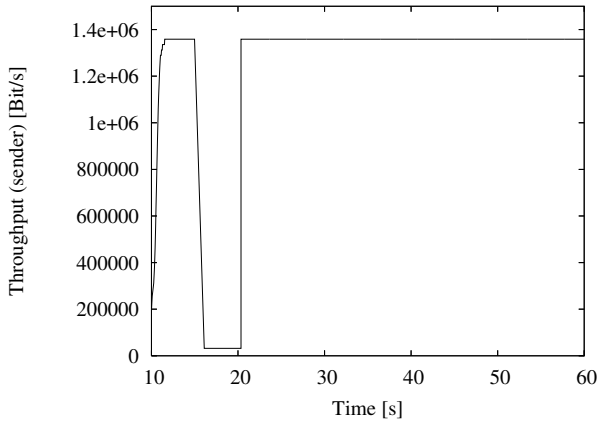

Figure 3.29: VTP throughput over time with disruption for one data flow in a single-hop scenario.

Upon detection of the disruption at 15 s via the absence of acknowledgments, VTP throttles its transmission rate to a very low *probing rate*, which intends to discover a resume of connectivity. When an acknowledgment indicates that connectivity is reestablished at 20 s, VTP recovers to the maximal possible rate within

one RTT. Figure 3.30 shows the respective time-sequence graph of the disruption in a single-hop environment. The x-range of the graph is restricted to 30 s in order to focus on the disruption phase.

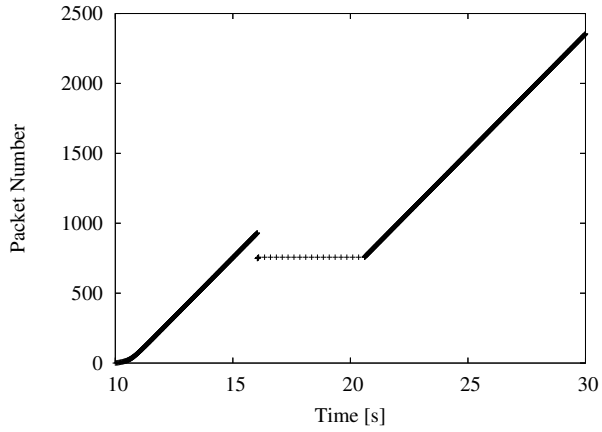

Figure 3.30: VTP time-sequence graph in case of disruption for one data flow in a single-hop scenario.

When the sender switches to *disrupted* state after the disruption at 15 s, it resets the sequence number to the first unacknowledged packet. The sender retransmits this packet at the probe rate until network connectivity resumes and the packet is acknowledged after 20 s. Upon the reception of the respective acknowledgment, the sender switches back to *connected* state and continues to transmit packets with ascending sequence numbers at the maximum rate.

Once more, these results are generally valid for single-hop and multi-hop connections whereas k and δ determine the level of fluctuations like in the static evaluations. Figure 3.31 shows an example of the disruption scenario for three hops. Figure 3.31(a) uses the settings k = 3 and δ = 0.05 whereas Figure 3.31(b) uses the settings k = 5 and δ = 0.2. The complete results of the disruption scenario are attached in Appendix A.

Summarizing, VTP reduces its transmission rate in a disruption, e.g., after a temporary network partition. VTP transmits probe packets at a low transmission rate to detect when connectivity resumes. Upon reception of an acknowledgment in response to a probe packet, VTP increases its transmission rate immediately to the maximal available bandwidth after a disruption.

Fairness Evaluation in Static Environments with Two Data Flows

VTP distributes the available bandwidth equally to competing data flows, according to the max-min fairness algorithm [12], as explained in Section 3.5.3.5. Figure 3.32 shows the throughput distribution of two data flows that both transmit at their maximum data rate in a single-hop scenario.

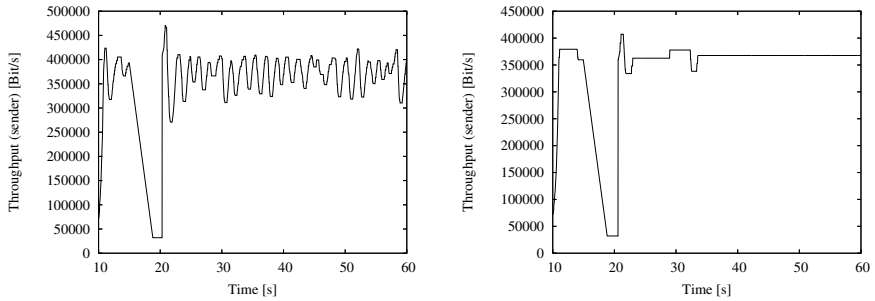

(a) VTP throughput over time with disruption for one data flow in a three-hop scenario with k=3 and δ=0.05.

(b) VTP throughput over time with disruption for one data flow in a three-hop scenario with k=5 and δ=0.2.

Figure 3.31: VTP throughput over time with disruption for one data flow in a three-hop scenario for different k and δ.

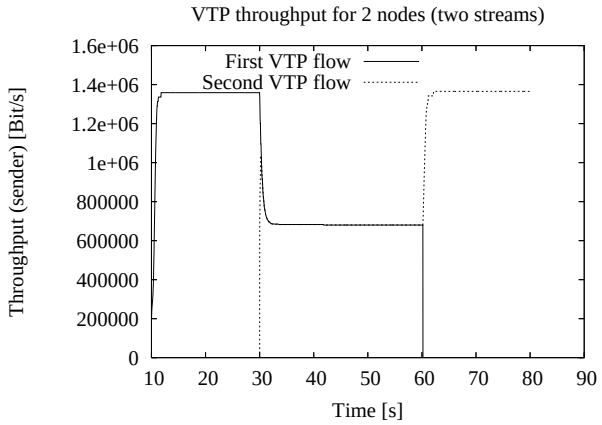

Figure 3.32: VTP throughput over time for two competing data flows in a single-hop scenario.

Between 10 s and 30 s, the first flow uses the complete available bandwidth of 1.4 MBit/s. When the second flow starts at 30 s, the bandwidth is equally shared among the flows. When the first flow ends after 60 s, the second flow immediately increases to the maximum throughput.

Figure 3.33 illustrates the VTP fairness via the time-sequence graph.

When the second flow starts at 30 s, the slope of the time-sequence curve of the first flow decreases, such that the first and second flow increase in parallel. Consequently, both flows transmit the same amount of data. When the first flow stops at 60 s, the slope of the second curve doubles and reaches the slope of the first flow before.

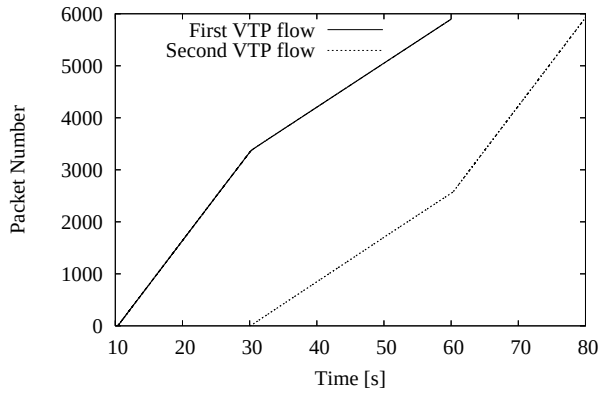

Figure 3.33: VTP time-sequence graph for two competing data flows in a single-hop scenario.

In comparison, Figure 3.34 shows the TCP throughput for two competing flows in a single-hop scenario.

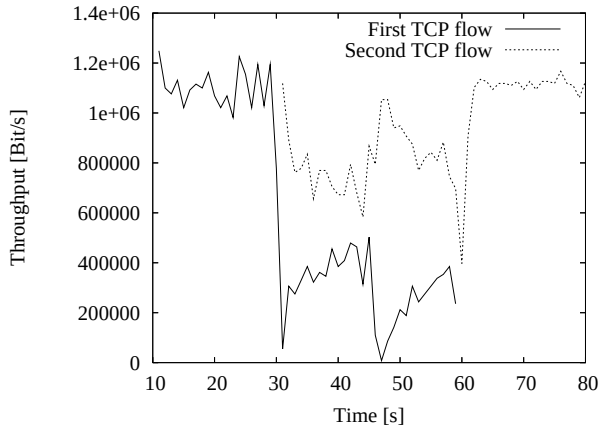

Figure 3.34: TCP throughput over time for two competing data flows in a single-hop scenario.

The graph shows that even in a static, wireless, single-hop scenario, TCP does not provide fairness. The second flow uses most of the available bandwidth while the throughput of the first flow reaches almost zero at 31 s and drops to zero at 48 s. The TCP congestion window decreases to zero several times, as shown in Appendix A. One of the main reasons for this unfairness is the mutual invocation of the MAC and TCP timers, as evaluated in detail in related work about fairness or MAC - TCP interaction, such as [78] or [3].

Figure 3.35 represents this unfairness among two simultaneous TCP flows in the time-sequence graph.

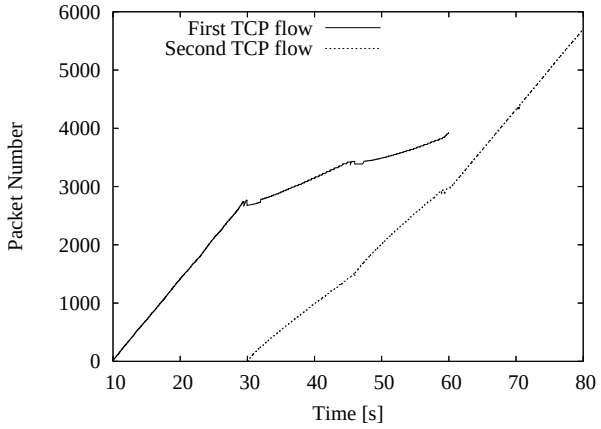

Figure 3.35: TCP time-sequence graph for two competing data flows in a single-hop scenario.

When both TCP flows transmit simultaneously between 30 s and 60 s, the slope of the second flow is higher than the slope of the first flow.

Again, the fair throughput distribution of VTP in the single-hop scenario is transferable to multi-hop connections. Figure 3.36 shows exemplary the throughput distribution between two competing flows in a three-hop scenario for two k - δ combinations. The full set of results is attached in Appendix A.

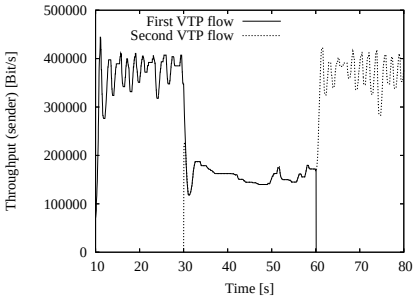

(a) VTP throughput for two data flows in a static three-hop scenario for k=3 and δ=0.05.

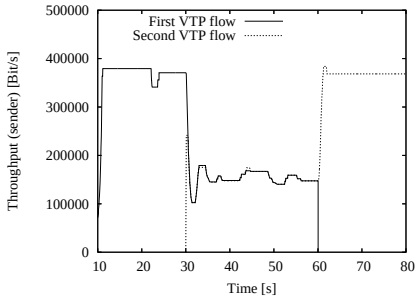

(b) VTP throughput for two data flows in a static three-hop scenario for k=5 and δ=0.2.

Figure 3.36: VTP throughput over time for two data flows (fairness) in a static two-hop scenario for different k - δ combinations.

Figure 3.36(a) shows the VTP throughput for k=3 and δ=0.05, and Figure 3.36(b) illustrates the VTP throughput for k=5 and δ=0.2. Similar to the throughput evaluation for one flow, the k-δ combination determines in this fairness evaluation the average throughput versus the fluctuations around this average.

Summarizing, unlike TCP, VTP provides fair distribution of throughput among coompeting flows.

3.6.3.2 Performance Evaluation in Mobile Highway Environments

This section presents the VTP and TCP performance results in mobile highway environments.

The vehicles drive along a 10 km highway stretch according to the validated movement patterns of [80] which represent typical weekday road traffic on German highways. The simulations consider the different road traffic densities, as provided by the movement patterns. The movement patters are divided into bidirectional cuts of 60 s that determine the maximum simulation duration, similar to the path characteristic analysis in Section 3.4.

The communication pairs are randomly chosen out of the vehicles that remain inside the evaluated highway section for the whole duration. The sender establishes an FTP connection to the receiver. Data is continuously ready to send.

The results show and compare the mean throughput and standard deviation per simulation run, as perceived by the receiver. The following figures show exemplarily the mean VTP and TCP throughput for the bidirectional scenario with two lanes per direction (lpd) and six nodes per lane, kilometer (npkm) and for VTP two different k and δ settings. The results are classified according to the maximum distance between source and destination during the simulation run.

Figure 3.37 shows the average throughput over source-destination distances for VTP in a scenario with 2 lpd, 6 npkm, k = 3 and δ = 0.05.

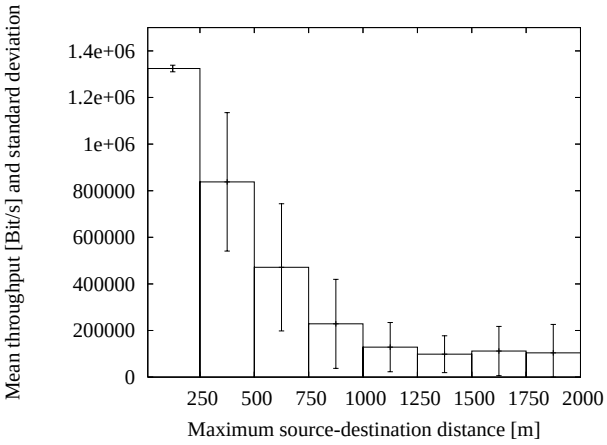

Figure 3.37: Average VTP throughput in a mobile bidirectional highway environment with 2lpd and 6npkm and k = 3, δ = 0.05.

Figure 3.38 illustrates the same scenario for k = 5 and δ = 0.2.

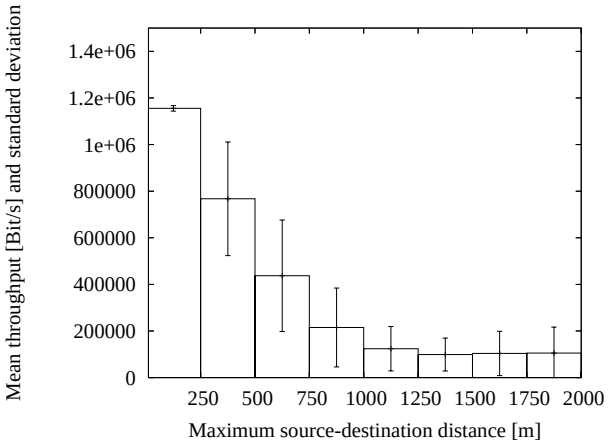

Figure 3.38: Average VTP throughput in a mobile bidirectional highway environment with 2lpd and 6npkm and k = 5, δ = 0.2.

In comparison, Figure 3.39 shows the respective TCP throughput. The complete results can be found in Appendix A.

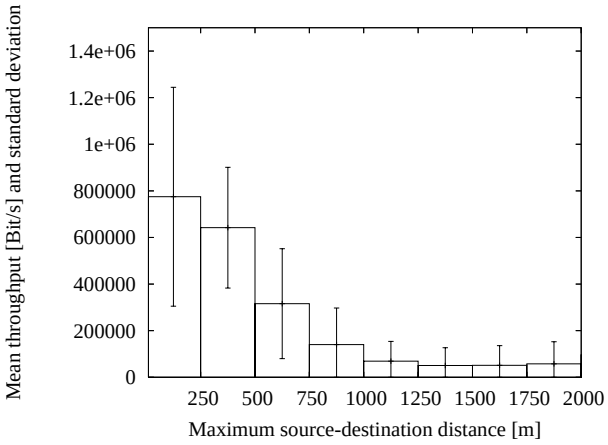

Figure 3.39: Average TCP throughput in a mobile bidirectional highway environment with 2lpd and 6npkm.

Summarizing, VTP performs better in mobile highway scenarios. The average gain of VTP is about 70%, depending on the source-destination distance. Beyond that, the standard deviation of VTP is in general smaller. As an example, when the communication pair remains in single-hop distance closer than 250 m, the mean VTP throughput is about 1.3 Mbit/s, which is up to 70% above the average throughput of TCP within this small distance. In this case the standard deviation of VTP is about 14 kbit/s whereas the standardard deviation of TCP is about 469 kbit/s. The performance gain of VTP is valid for all scenarios, as shown in Appendix A.

3.7 Summary and Conclusion

Vehicular ad hoc networks (VANETs) enable multi-hop vehicle-to-vehicle and vehicle-to-roadside wireless communication in a self-organized ad hoc network. The movement of vehicles represents the main characteristic of VANETs. Vehicular mobility results in frequent topology changes. Novel routing protocols, such as position-based routing (PBR), are tailored to this environment since PBR forwards packets *per hop*, and no end-to-end route is required. Consecutive packets might follow different paths. These unique characteristics affect end-to-end connections, in particular transport protocols. This chapter analyzes the path characteristics that transport protocols experience in highway scenarios, such as connectivity and disruption duration, packet loss, reordering, round trip time (RTT) and RTT jitter. These results aid in the following design of a vehicular transport protocol (VTP) which is evaluated through simulations.

The evaluation of the path characteristics investigates the metrics connectivity and disruption duration, packet loss probability and characteristics, packet reordering, RTT and RTT jitter.

The connectivity evaluation results show that steady communication is feasible for source-destination distances up to 2000 m. For a distance of 2000 m, about 40% of the connections remain uninterrupted for 10 s on the average. With decreasing distance, the connectivity duration even increases. Disruptions resume after 3 s at the latest, only marginally dependent on the distance.

The packet loss ratio for a constant packet stream is huge: For a distance of 2000 m, standard PBR shows a packet loss rate of almost two thirds which can be significantly reduced to 22% when using cross-layer integration.

Although the RTT and RTT jitter are acceptably small for source-destination distances below 700 m, higher distances result in extreme fluctuation in RTT, *e.g.*, up to 300% for a 2000 m distance.

Finally, reordering ratios for light loads are small (below 1%), but increase to 15% for medium data loads.

These unique characteristics of VANETs necessitate the development of a novel transport protocol.

The evaluation results of the path characteristics influence the design of a vehicular transport protocol (VTP) that is tailored to the unique properties of VANETs. VTP aims at maximizing the throughput of a connection while preserving fairness to competing data traffic. The objectives of VTP include the establishment and release of an end-to-end connection, reliable delivery of data packets, flow and congestion control. The reliability mechanisms of VTP must cope with frequent packet losses, reordering, high RTT and high RTT jitter. The performance of VTP mainly depends on its ability to adapt quickly to varying path characteristics.

The key features of VTP are:

- The VTP sender uses a rate-based transmission scheme. The transmission rate is determined by a *rate-timer* that steadily schedules the transmission of data packets when multi-hop connectivity between source and destination is assumed.
- VTP decouples congestion control from error and flow control, mainly to avoid throughput reduction for packet loss not related to congestion. In VANETs, packet losses are frequent because of high mobility and the resulting topological changes. These losses must not invoke congestion control.
- VTP uses explicit signaling of available bandwidth from intermediate nodes for congestion control. The estimation of available bandwidth by intermediate nodes uses information from the MAC layer protocol.
- VTP provides reliability via retransmissions of lost packets. Selective acknowledgments (SACKs) report lost packets to the VTP sender. The receiver transmits SACKs in dynamic intervals. It adjusts the interval according to the current transmission rate and the source-destination distance.
- The VTP sender uses statistical knowledge to predict the expected communication behavior of a connection. In absence of acknowledgments, the expected communication duration for the respective source-destination distance assists the rate timer calculation.

A simulative study evaluates the throughput and fairness of VTP in static and mobile wireless environments and compares these metrics to the performance of TCP.

VTP maintains a constant transmission rate and reacts quickly to disruption or congestion, based on feedback (or absence of feedback) from intermediate nodes. Selective acknowledgments inform the sender about received and missing packets in order to provide reliability by retransmissions. VTP uses statistical knowledge to predict connection behavior, such as expected communication duration, and adapts its transmission rate accordingly.

As a main result, VTP provides reliable end-to-end connections and outperforms the varying throughput and unfairness of TCP by maintaining a steady throughput above the average throughput of TCP.

The future work will adapt and evaluate VTP in city scenarios, assuming that VTP can achieve similar performance like in highway scenarios when the underlying routing protocol maintains similar packet delivery ratios. VTP will be implemented in the framework of the NoW project and evaluation via measurements will be performed. Finally, interoperability to TCP is required in order to allow access to the Internet or fixed networks at the roadside. This can for example be achieved by installing translation proxies at road side access points or tunnel VTP in TCP over the fixed network.

The previous chapter designed a vehicular transport protocol for unicast communication. Beyond these point-to-point applications, one of the main goals of VANETs is the increase of road safety by reliable point-to-multipoint distribution of safety information to endangered vehicles.

Typically, safety applications require the efficient and reliable distribution of information to vehicles inside a geographically restricted target area over time. The information should be *kept alive* in the target area for the lifetime of the safety event, i.e., particularly vehicles that enter the target area after the initial message is distributed must be informed.

The following Chapter 4 designs and evaluates a *time-extended reliable geographical flooding* algorithm that provides reliable and efficient distribution of information in a target area over time.

Chapter 4

Information Distribution in a Geographical Area in Vehicular Ad Hoc Networks

4.1 Introduction

Vehicular ad hoc networks (VANETs) facilitate enhanced traffic safety by means of wireless multi-hop communication in a self-organizing network. In fact, the increase of safety on the road is the main objective of VANETs, beside passenger information and entertainment. VANETs enable active safety applications, such as hazard warning or extended brake lights, by extending the driver's horizon and warn affected traffic about potential dangers as early as possible. The single-hop and multi-hop distribution of safety messages ensures that relevant information is transmitted and consumed in the local area where it is needed.

The distribution of safety information poses specific challenges on reliability and efficiency:

(i) Safety messages should be delivered to affected vehicles only (e.g., vehicles approaching the hazard). Typically, safety events in VANETs are bound to a restricted geographical region, termed *target area* (TA). Geographically-scoped flooding (GeoCast) addresses all vehicles in a geographical region [62, 91]. However, GeoCast does not provide any means for reliability and causes redundant message repetitions (it is usually based on a flooding algorithm where each vehicle forwards each message once). Therefore, additional algorithms for reliable and efficient distribution of safety information in a geographical target area are needed.

(ii) The high degree of vehicle mobility in VANETs cause that vehicles continuously leave and enter the target area. The movement results in frequent topology changes. Vehicles that enter the target area after the message has been initially distributed miss the safety information. It must be ensured that those vehicles are also informed within the lifetime of the event, i.e., VANETs require reliability over time.

(iii) Messages are sent over error-prone wireless links. In case a message is lost, redistribution is required. However, the rebroadcasting should be restricted to the local, single-hop neighbor scope in order to avoid redundant multi-hop retransmissions. Furthermore, re-flooding should not be used preemptively but must only be utilized when a message loss is indicated (e.g., in absence of acknowledgments).

In summary, the design of algorithms for *reliable* and *efficient* (i.e., avoiding redundant retransmissions) distribution of safety information in geographical target areas is challenging.

Traditionally, reliability is defined as the guaranteed delivery of messages from a source to a single or multiple receiver(s). Reliability assures that messages arrive uncorrupted and in-sequence at their destination(s). In order to meet the requirements of traffic safety applications in VANETs, this chapter extends the classical reliability definition by spatial and time components:

Reliability of safety information in VANETs requires the reliable distribution of an information to all vehicles inside a geographical target area during the lifetime of a safety event. This explicitly includes the distribution of the safety information to vehicles that enter the target area after the information has already been issued.

This chapter surveys and evaluates the related GeoCast approaches (e.g., with and without temporal caching of messages) and presents the *time-extended reliable geographical flooding (TERGF)* algorithm to provide reliable and efficient distribution of safety messages in VANETs, according to the extended reliability definition above.

4.2 Background

This section provides an overview on the GeoCast protocol which addresses nodes in a geographical area.

4.2.1 GeoCast

The GeoCast algorithm [62, 77, 91] provides geographical addressing and the delivery of messages to vehicles inside a specific geographical region, termed *target area*. GeoCast is commonly classified as a multicast protocol where the geographi-

cal location of the vehicles determines the membership to the multicast group. Thus, all vehicles must be aware of their geographical position by means of a positioning system, such as the global positioning system (GPS) [57, 69]. A GeoCast message contains the definition of its target area by geographical coordinates (i.e., latitude and longitude) and a geometrical shape. The target area is typically coded by geographical positions and geographical shapes, as shown in the following:

- Point,
- Circle, defined by (center point, radius),
- Rectangle, defined by (two points, height), and
- Polygon, defined by ($point_1$, $point_2$, ..., $point_{n-1}$, $point_n$, $point_1$).

GeoCast distinguishes two phases in the distribution process:

(i) When the sender is not located inside the target area, the GeoCast message is first forwarded towards the target area by means of standard, unicast routing (e.g., PBR). This phase is referred to as *line-forwarding*.

(ii) When the message reaches the first vehicle inside the target area or the sender is located inside the target area, the GeoCast is distributed (e.g., flooded) through the network inside the geographical boundaries. This phase is termed *area-forwarding*. Since safety messages typically concern the immediate vicinity of the event, GeoCast for safety applications assumes that the originator of the message is located inside the target area, as we assume in the remainder of this section.

The literature [139] classifies existing GeoCast approaches into three categories:

(i) Flooding-based protocols, such as location-based multicast (LBM) [76] and the Voronoi diagram-based GeoCast [122], use flooding or a variant of flooding to route and distribute GeoCast packets.

(ii) Routing-based protocols, such as mesh-based GeoCast routing protocol (MGRP) [16], GeoCast adaptive mesh environment for routing (GAMER) [23] or GeoTORA [75], establish routes from the source to the vehicles in the target area via explicit control messages.

(iii) Cluster-based protocols, such as GeoGRID[84] or the obstacle-free single / multi-destination geocasting protocol (OFSGP/OFMGP) [26], geographically *partition* the network into several disjoint and equally sized regions. Each region assigns a cluster-head for executing the information exchange.

4.3 Related Work

This section surveys flooding-based and multicast-based schemes for improving the reliability and efficiency in wireless networks.

4.3.1 Flooding Approaches in Wireless Networks

The literature classifies existing flooding approaches in four categories, which are: Simple flooding, probability-based flooding, area-based flooding and flooding based on the knowledge about neighboring vehicles [136].

Simple flooding. With *simple flooding*, also referred to as *blind flooding*, a node simply rebroadcasts a message exactly once. The nodes use e.g., the source ID and packet ID to identify already received messages. This distribution process continues until the message has traversed the network. Hence, every node that receives a message forwards this message to all its neighbors although they may have already received this information. As a result, messages are duplicated, and bandwidth is wasted. However, this algorithm achieves a high probability of reliability, achieved by (redundant) repetitions of messages.

Probabilistic-based flooding. Probabilistic-based flooding, e.g, [99], is similar to simple flooding, except that nodes only rebroadcast a message with a certain probability. In networks with a high density of nodes, a high reception probability can be achieved while saving scarce wireless bandwidth because multiple nodes share similar radio transmission ranges. However, in networks with a low density of nodes, the reception probability decreases, and not all nodes receive the message.

The *counter-based scheme* in [99] uses the inverse relation between the number of times a packet is received by a node. The nodes calculate the probability that it can reach additional neighbors with a rebroadcast. Upon reception of a new message, the node initiates a counter and starts a timer which is randomly chosen. As long as the timer continues, the counter is incremented for each redundantly received packet. Upon timer expiration, the packet is only rebroadcast in case the counter is less than a pre-defined threshold. Otherwise, the node drops the message.

Area-based flooding. Area-based flooding algorithms use distance information in the decision process whether a packet should be rebroadcast or not. The node may use geographical position knowledge of a positioning system or it could estimate the distance via signal strength measurements. When the receiving node is close to the sender, the additional area covered by a retransmission would be small whereas a receiver far away covers more additional area. The *distance-based scheme* and *location-based scheme* in [99] estimate the additional coverage area in this way.

A node using the distance-based scheme compares the distance between itself and all neighbors from which it received the message. Upon reception of a previously not received message, the node initiates a timer and caches redundant packets. When the timer expires, all nodes compare their distance to the source to a threshold, and the packet is only rebroadcast in case a node is closer than a pre-defined distance.

Location-based schemes use a more precise estimation of the expected additional coverage area, which relies on geographical positioning. Each node must be able to determine its geographical position. Each packet contains the geographical position of its sender or forwarder. When a node receives the packets, it calculates the physical distance and the additional coverage area it could reach. Again, it compares the result with a pre-defined threshold in order to decide if the node rebroadcasts the packet.

Neighbor knowledge flooding. Using *self-pruned flooding* [85], a node broadcasts a message on the wireless link that includes a list of its single-hop neighbors. Every node that receives the message compares the list of nodes in the message (except the node itself) with its own neighbor list (NL). In case the node has further neighbors that are not in the list (i.e., $NL_{(rec)} > NL_{(msg)} \bigcup (sender)$), it replaces the list of neighbors in the message by its own single-hop neighbors and rebroadcasts the message. While this algorithm achieves a similar level of reliability as simple flooding, it reduces the overhead significantly since the approach avoids message duplications to neighbors in overlapping wireless regions.

The time-extended reliable geographical flooding (TERGF) algorithm of this chapter uses the basic idea of flooding with self-pruning, although in a different context. The TERGF algorithm distributes information rather than packets, i.e., the TERGF design assists content-based aggregation of information to reduce bandwidth consumption. Beyond, TERGF combines self-pruning with GeoCast and extends the algorithm by an acknowledgment scheme in order to achieve full reliability. For self-pruned flooding, the node includes only those neighbors into the list in the message that are located inside the geographical target area. If a receiving node compares the list with its own neighbors, it also excludes the neighbors that are not inside the target area. As a result, a node rebroadcasts a message only if it reaches further nodes inside the geographical target area.

The *scalable broadcast algorithm (SBA)* [105] uses two-hop neighbor knowledge. In order to establish this knowledge, the nodes periodically exchange *hello messages* that contain the neighbor list of its predecessor. When a node receives a packet, the receiver compares its own neighbors with the sender's neighbors in order to determine if a rebroadcast would reach additional nodes.

The *dominant pruning* approach [85] also utilizes two-hop neighbor knowledge. In this approach, the sender proactively selects the one-hop neighbor(s)

which should forward the message. Only selected nodes are allowed to rebroadcast the message. When a node receives a message, it checks if its address is included. If so, it rebroadcasts the message and uses a modified version of the *greedy set cover* algorithm [88] in order to select the next level forwarding nodes.

The *multipoint relaying* mechanism [112], which is part of the *optimized link state routing (OLSR)* [30] protocol, is similar to dominant pruning. Based on a two-hop neighbor knowledge, the sender determines *multipoint relays (MPRs)* that are responsible for the redistribution of the message.

The *ad hoc broadcast protocol (AHBP)* [106], *CDS-based broadcast algorithm* and the lightweight and efficient network-wide broadcast (LENWB) [124] are similar to the multipoint relaying approach, but differ in the calculation effort to determine the forwarding nodes. Details are given in [136].

4.3.2 Passive Acknowledgments in Wireless Networks

The *passive acknowledgment scheme* uses the shared character of the wireless medium, as follows. In case a message is forwarded on a wireless channel via multiple hops, the sender (or forwarder) listens for the forwarding (i.e., retransmission) of the packet by the respective successor node. A node can only forward correctly received messages. Therefore, the successor must have received the message correctly when the sender can *overhear* the forwarding of the message. Consequently, the sender interprets the overheard message as a *passive acknowledgment*.

The time-extended reliable geographical flooding (TERGF) algorithm, as presented in the remainder of this chapter, adopts this scheme, though in a different context. When a packet needs to be rebroadcast by one of the next hops, the sender / forwarder interprets the rebroadcast as a passive acknowledgment. However, in case there is no need to forward a message, an explicit acknowledgment is required. Section 4.5 explains the TERGF algorithm in detail, including the adaptation of the passive acknowledgment scheme.

4.3.3 Reliable Multicast Communication

Multicast communication in traditional packet switched networks faces similar problems to the GeoCast-type of communication in VANETs. First, transport protocols based on positive / negative acknowledgments (ACK/NACK) are not applicable for large multicast groups. As the number of multicast receivers grows, the amount of back traffic overwhelms its capacity to handle them (i.e., ACK/NACK implosion). Second, if losses occur uncorrelated on different parts of the multicast tree, data may need to be sent multiple times to satisfy all receivers.

In schemes for reliable multicast, such as RMTP [103] or SRM [42], designated receivers collect status information from receivers and retransmit lost data packets on request of other receivers. Hence, these nodes provide both local re-

transmission of data and aggregation of signaling traffic for retransmission. Both reliable multicast schemes are explained in the following sections.

The main differences of reliable multicast approaches to the time-extended geographically scoped reliability provisioning are:

(i) The addressed nodes inside the target area do not join a multicast group. Instead, the nodes are identified with respect to their geographical position.
(ii) There are no multicast distribution trees for data forwarding or ACK/NACK notification.
(iii) The broadcast characteristic of the wireless medium in multi-hop ad hoc networks facilitates overhearing of messages.

4.3.3.1 Reliable Multicast Transport Protocol (RMTP)

The *reliable multicast transport protocol (RMTP)* [103] provides reliable (i.e., sequenced and lossless) delivery of a data stream from one sender to a group of receivers in the Internet. The RMTP design follows a multi-level hierarchical approach. RMTP groups receivers into a hierarchy of local regions with a *designated receiver (DR)* in each region. Each receiver acknowledges data to its corresponding DR, which in turn relays acknowledgments to the next level of DRs until the acknowledgments reach the original sender. This mechanism avoids an avalanche effect of acknowledgments, known as the acknowledgment implosion problem. DRs cache data and retransmit them upon request which decreases the end-to-end latency in case of losses. RMTP uses a packet-based selective retransmission scheme in order to increase the throughput.

An extension to RMTP is the *reliable multicast file transfer protocol (RMFTP)*. This application-level protocol uses TCP for bi-directional control messages and RMTP for the uni-directional data transmission. RMFTP facilitates server-based *push* and client-based *pull* transmissions.

4.3.3.2 Scalable Reliable Multicast (SRM)

The SRM protocol [42] provides a framework for scalable, reliable multicast. Different multicast applications have widely different requirements, e.g., with respect to reliability, the number of sources or replication of data strategies. These differences affect the design of a multicast protocol, which should dynamically adapt to the specific requirements, but leave as much functionality and flexibility to the application as possible.

SRM is based on the IP multicast protocol [35]. In IP multicast, a data source simply transmits to the group's multicast address. In order to receive data, each receiver *joins* the multicast group in the local sub-network. SRM enhances the multicast group concept by maximizing information and data sharing among all members. Thus, each member is responsible for its correct reception of all the

data. Furthermore, SRM follows the design of TCP/IP by adopting the best-effort data delivery model and building reliability on an end-to-end basis.

SRM dynamically adapts its control parameters to the observed network performance. This allows applications on top of SRM to adapt to a wide range of group sizes, topologies and link bandwidths, while maintaining robustness and high performance.

4.4 Temporal Caching of GeoCast Messages

The multi-hop forwarding and distribution of messages relies on the availability of appropriate single-hop neighbors: In unicast, the absence of a single-hop neighbor that is closer to the destination results in a routing failure. In this case, either routing recovery strategies can find an alternative route or the packet is dropped when no alternative route can be found in a predefined time interval, e.g., in low density or highway scenarios. In GeoCast, insufficient connectivity (e.g., in low density scenarios) results in network partitions. Thus, the message cannot reach all vehicles inside the target area. The high mobility of VANETs causes frequent changes in the network topology, which may also result in temporal network partitions. Particularly, scenarios with a low density of equipped vehicles are affected.

One approach to improve the reliability of GeoCast message distribution - particularly suited for low density scenarios - is to add a cache for GeoCast packets. The following section presents this concept of *store-and-forward*.

4.4.1 Store-and-Forward Concept

The basic idea of *store-and-forward* is to add a queue for GeoCast packet in the network layer of every vehicle and to cache GeoCast packets for a certain amount of time. A vehicle can retransmit the packet out of the cache when needed, depending on the classification below (e.g., periodical rebroadcasts or retransmission upon detection of a new neighbor).

Particularly in scenarios with a low density of vehicles equipped with a communication system, the retransmission of a cached packet can increase the packet delivery ratio. The repetition of a cached message can inform vehicles that e.g., could not be reached during the initial distribution of the message, due to a temporal network partition or in case the vehicle enters the target area at a later point in time. The following section 4.4.2 presents a typical target scenario.

According to the line- and area-forwarding classification, as defined in Section 4.2.1, GeoCast caching strategies differ as follows:

The goal of store-and-forward in the line forwarding mode aims at preventing packet loss in case of route failures towards the target area. Thus, packets are only cached when no neighbor closer to the destination is available. When the

network layer is notified about newly discovered neighbor(s), e.g., by means of PBR beacons, which are more close to the destination, the vehicle forwards the previously *unroutable* packet.

The intention of store-and-forward in the area-forwarding mode is to keep the information alive for a validity time / lifetime of an event. Thus, *each* GeoCast packet is cached for a pre-defined interval. In this scenario, the vehicles either rebroadcast the message periodically or rebroadcast it upon detection of new neighbor(s).

The remainder of this chapter focuses on store-and-forward when the packet is distributed in the area-forwarding mode. Safety applications in VANETs focus on this mode to improve reliability of safety messages because a safety information is valid only inside a restricted geographical area in the vicinity of the message originator. A vehicle sensing a hazard situation addresses the related warning message to affected vehicles (i.e., vehicles approaching the hazard) in the restricted geographical area surrounding the hazard only. Consequently, there is no need to transport the message to the target area in line-forwarding mode.

Though store-and-forward improves the reliability, it also increases the network load due to redundant retransmissions of messages. The following parameters impact the performance of store-and-forward and can be tuned for optimization:

- Cache size,
- Maximum number of retransmissions,
- Timer management in the cache:
  - Restriction of the maximum inter-packet delay for retransmissions,
  - Periodic validity verification of cached packets (i.e., cache clean up).

In the framework of the network on wheels (NoW) project [101], a basic version of store-and-forward is implemented, as reported in the following.

4.4.2 Target Scenario

Figure 4.1 illustrates an exemplary VANET target scenario with low penetration of equipped vehicles.

Vehicle A experiences a hazard situation, broadcasts a respective GeoCast warning message and simultaneously caches the message. Since the vehicles do not have connectivity to each other, as indicated by the circular transmission ranges, the initial message does not reach any other vehicle. When vehicle B approaches and enters A's communication range, vehicle A retransmits the warning message to the endangered vehicle B. In other words, A fetches the message from its cache and rebroadcasts it. Upon message reception, vehicle B warns the driver and, again,

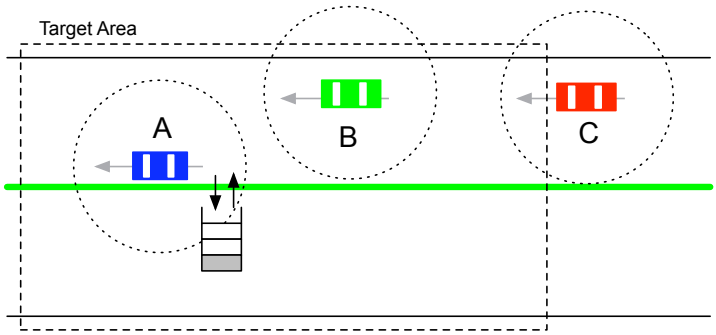

Figure 4.1: Exemplary scenario for GeoCast with store-and-forward.

forwards and caches the safety message. When vehicle C enters the target area at a later point in time, in is informed as soon as it has wireless connectivity to vehicle B.

In another scenario, the *physical transport* of a message, e.g., via oncoming traffic, can assist to improve reliability, as shown in Figure 4.2. The oncoming vehicle B receives the safety message of vehicle A and caches it while continuing its way. The message is physically transported and repeated when vehicle B passes by vehicle C. Thus, vehicle C is informed as early as possible.

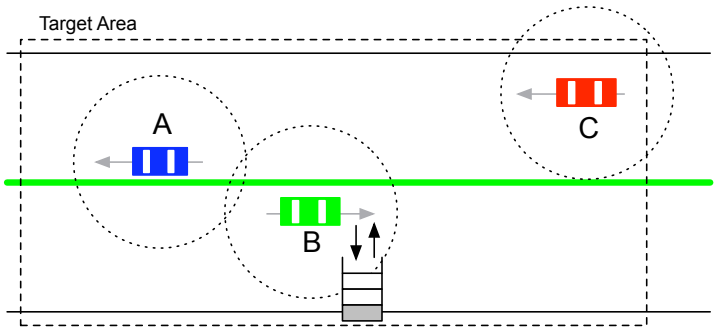

Figure 4.2: Physical transport of a GeoCast message via store-and-forward.

4.4.3 Implementation Report

The position-based router implementation of the NoW project [101] provides data delivery for topologically-scoped broadcast messages, position-based unicast and GeoCast packets. The store-and-forward concept is implemented and tested as an extension of this communication system software.

The GeoCast store-and-forward implementation adds a queue to the PBR network layer. Each time a vehicle generates or forwards a GeoCast message and sends the packet to the wireless interface, the store-and-forward implementation adds a copy of the packet to this queue. Since the PBR GeoCast implementa-

tion only accepts and forwards GeoCast packets by vehicles inside the target area, caching is restricted to vehicles that area located inside the target area upon packet arrival. Figure 4.3 illustrates the integration of the GeoCast cache schematically.

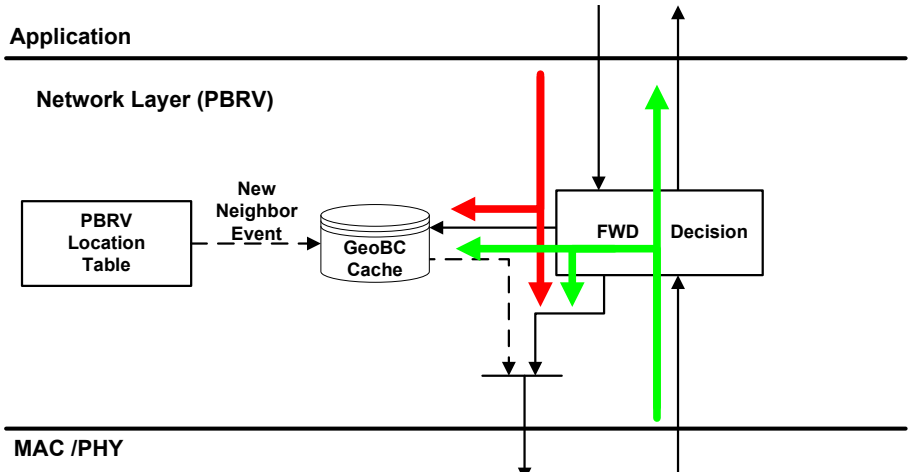

Figure 4.3: Main functional blocks for GeoCast with store-and-forward.

The cache is implemented as a *First In - First Out (FIFO)* queue. The queue size is determined by a variable in the configuration file of the NoW router, i.e., it is adjustable before starting the PBR router, but it remains fixed during runtime. The default size of the cache is ten packets, which is sufficient for scenarios with a low density of equipped vehicles. The *drop-on-overflow* management is implemented so that the most outdated packet is dropped first. In the current, basic version of the implementation, packets are stored until a new neighbor is detected without time limit. For system deployment, a timer management should be activated to remove outdated packets from the queue in order to enhance efficiency. However, such a timer was not desired for the demonstration implementation in order to be more flexible in during the demonstration.

The detection of a new neighbor inside the target area triggers the rebroadcast of the respective packets from the cache. The neighbor table of the PBR routing protocol manages single-hop neighbor information and informs the store-and-forward queue about new neighbors and their respective geographical positions. Remember that PBR relies on the periodic exchange of beacons between single-hop neighbors which announce the presence and geographical position of a vehicle. PBR manages the local neighborhood information in the *neighbor table*. Upon the arrival of a beacon, PBR searches the neighbor table for a corresponding entry. In case a respective entry is detected, PBR updates this entry, e.g., timestamp or aggregated target area. Otherwise, a new entry is added. The creation of a new entry in the neighbor table indicates that a new neighbor is reachable within the radio coverage area. PBR relays this information to the store-and-forward implementation. The store-and-forward scheme checks if this new neighbor is located inside

the target area of the cached packets. If so, the respective packet(s) are rebroadcast out of the queue.

The functionality of the GeoCast caching has been successfully presented in a demonstration of the INVENT project [114] which relies in the inter-vehicle communication (IVC) system of the FleetNet and NoW projects.

Store-and-forward increases the packet delivery ratio, particularly in scenarios with low densities of equipped vehicles. The algorithm rebroadcasts cached messages, e.g., when temporal network partitions resume when further vehicles reach the dangerous area or when the physical transport of packets due to the movement of the vehicles reaches additional vehicles.

However, the maintenance of long transmission queues in case of multiple events decreases the network performance. Beyond, the reliable distribution of messages over time via store-and-forward significantly increases the network load, particularly for a high density of participating vehicles. Store-and-forward may rebroadcast many messages redundantly because the algorithm cannot distinguish different messages for the same event or identify already informed vehicles, as shown by the following examples:

(i) When several vehicles approach and detect the hazard, each vehicle generates a safety warning message based on its sensor information. Consequently, multiple safety messages for the same event circulate in the target area. Since store-and-forward is not aware of the content of a message, it loads the queues and distributes in the network many different messages with the same content and information. A *bursty* transmission of messages increases contention and increased collision rate in the access medium and results in longer delays for the message distribution.

(ii) Store-and-forward rebroadcasts messages upon detection of new neighbors inside the target area, regardless whether the new neighbor is already informed about the respective safety event. If so, the rebroadcasting of the message is redundant. Particularly in scenarios with a high density of equipped vehicles, the redundant rebroadcast of messages is significant and decreases the network performance. As an example, vehicles that drive in the opposite direction of a traffic jam may physically transport the packet and rebroadcast it frequently when detecting new neighbors. However, the message has already been distributed through the traffic jam before. Thus, each equipped vehicle driving opposite to the traffic jam will frequently repeat the already known information, while driving on and detecting new neighbors. The network is continuously loaded with redundant information.

The reliable and efficient distribution of information in a geographical target area, independent of the density of equipped vehicles, demands for a more ad-

vanced algorithm. Such an algorithm should consider message content and maintain the information state of vehicles in its vicinity inside the target area. The *time-extended reliable geographical flooding* algorithm, as presented in the following section, aims at this reliable and efficient distribution of information in a geographical target area over time.

4.5 Time-Extended Reliable Geographical Flooding (TERGF)

The *time-extended reliable geographical flooding (TERGF)* algorithm aims at reliable and efficient distribution of information in a geographical target area over time in an ad hoc communication network. Reliability, as defined in TERGF, refers to the distribution of information rather than individual packets, in contrast to traditional packet-switched networks. TERGF provides efficient reliability by combining the following mechanisms:

1. Reliable distribution of safety information (in contrast to the *traditional* packet-based reliability) via a safety information manager.

2. Flooding of safety information based on the combination of GeoCast and self-pruning, extended by acknowledgment-based single-hop reliability.

3. Passive acknowledgment to detect single-hop losses.

4. Redistribution of safety information in case of single-hop losses.

5. Redistribution of safety information when vehicles enter the target area for the lifetime of the safety event.

Principally, vehicles maintain states of safety events and communicate to distribute these states. Though the distribution of information uses GeoCast as the basic addressing scheme, the communication of safety messages is restricted to single-hop broadcasting. A vehicle that receives a safety message updates its local safety information state, including aggregation of information. Only if this receiver can identify further, not-informed neighbors, it generates a new safety message to further distribute this information. The forwarded message may be different from the received message. Particularly, a new safety message may be generated at a later point in time, e.g., when additional vehicles enter the target area. Consequently, the information is efficiently distributed and *kept alive* for the lifetime of an event in the target area.

The following sections explain the TERGF algorithm in detail, provide a simulative evaluation of TERGF and compare the metrics *information distribution ratio* and *redundant packet rate* to the standard GeoBC approach.

4.5.1 Assumptions about the TERGF Communication System

This section describes the vehicular communication system, as required for the distribution of safety information. Furthermore, it defines an information structure to maintain safety information and explains the distribution of safety information in the geographical area based on the TERGF algorithm.

4.5.1.1 Safety Information Structure

With respect to traffic safety applications, VANETs represent a distributed system where a global state about the environment is distributed in a geographical area. A VANET node maintains a local state that represents a subset of the global state. Since information is continuously generated and associated with a lifetime, states can be aggregated, and a state entry can also disappear. In order to keep local state information up-to-date, vehicles communicate with other vehicles in their spatial neighborhood.

The TERGF algorithms assumes that the local state in a vehicle can be described by an abstract safety information structure, as shown in Figure 4.4. This information structure is generic, such that it can be easily adapted to any type of traffic safety applications. In its basic shape, the structure defines the relevant information of a safety event or hazard situation. Extensions would be specific information related to the type of event or additional information, such as alternative routes.

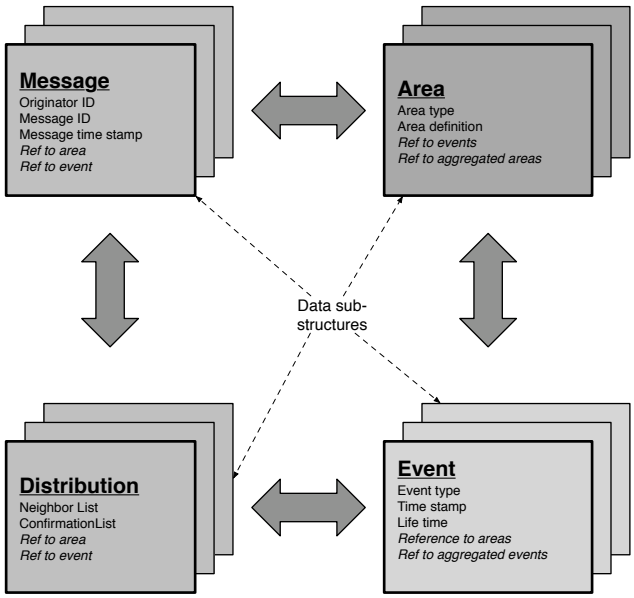

Figure 4.4: Safety information structure.

The overall structure comprises *message*-related, *event*-related, *area*-related and *reliability*-related information, organized as sub-structures. Message-related

information refers to originator ID and message ID that uniquely identify a message in the network, including the timestamp of a particular message. Event-related information comprises event type and life time of the event. Area-related information determine the geographical area of an event in terms of area type, position and size[1]. As the TERGF specification in section 4.5.2 will present, reliability-related information contains the state that is needed by the reliability algorithm.

The sub-structures organize the contained elements in tables, such as lists or hash tables. An element includes pointers to other elements in other sub structures, e.g., an element in the message sub-structure points to elements in the event and area sub-structures. Multiple elements in one sub-structure may point to a single element in another sub-structure. This linking allows two different events to be associated with a single area.

The proposed structure facilitates aggregation, e.g., with respect to area and event. Multiple elements of the area sub-structure can be aggregated to a new element. Typical examples, as shown in Figure 4.5, are:

(i) Multiple areas with the same event type can be aggregated to a new element that comprises the adjacent geographical areas, linked together.

(ii) Multiple events associated with the same area can be fused into a new element that comprises all of these events.

In both cases, a new element in the respective information sub-structure is generated and linked with the original entries. Consequently, the elements in a sub-structure are in a tree-like relation with an aggregated element as the root of the tree and the original entries as leaves. Since each element is associated with a lifetime, an aggregated element refers to the minimum lifetime of the original elements.

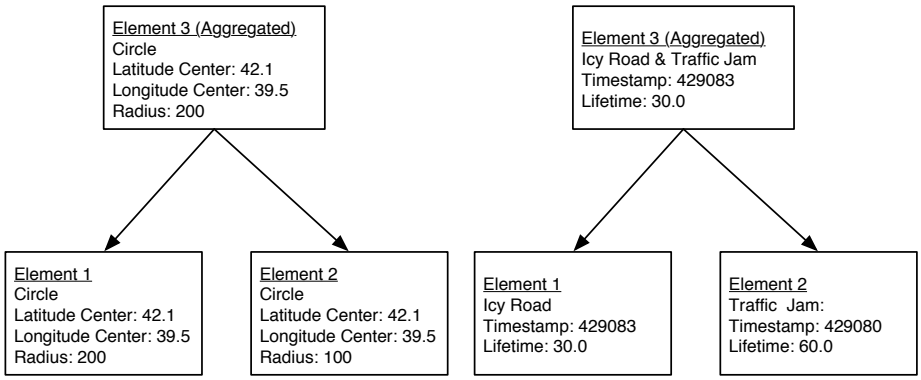

(a) Elements of area sub-structure. (b) Elements of event sub-structure.

Figure 4.5: Exemplary aggregation of sub-structure elements.

[1]E.g., a circle is defined by its center (as latitude and longitude) and radius.

4.5.1.2 Distribution of Safety Information

In order to maintain the safety information structure, vehicles exchange information. The messages address all vehicles in the geographical target area associated with the event. However, as mentioned before, each receiver individually decides whether to redistribute the information or not. The *safety information manager* in each vehicle controls the communication: It receives all safety messages, creates or aggregates elements in the safety information structure and decides whether to transmit or redistribute information. Furthermore, it maintains timers and removes outdated elements from the safety information structure. In addition, it provides the interfaces to required local information or the one-hop neighbor list of the network layer, as shown in Figure 4.6.

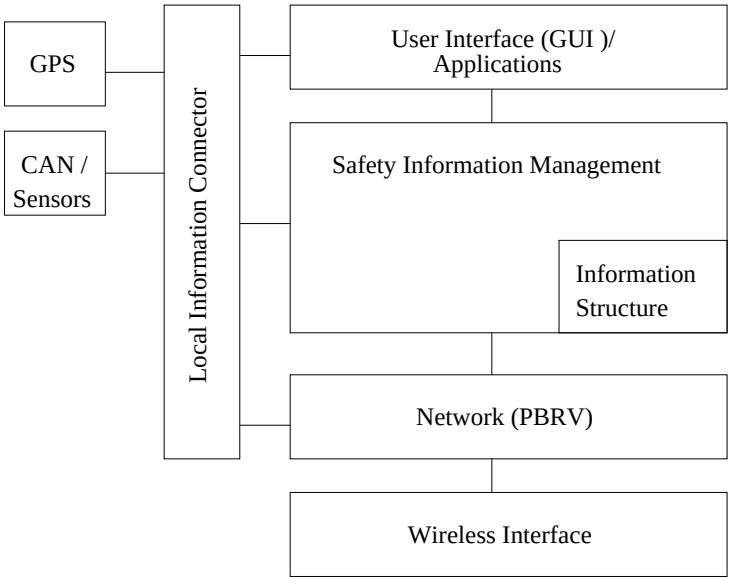

Figure 4.6: Safety information manager.

In principle, vehicles exchange information by means of single-hop broadcast messages. A vehicle issues a message if the safety information manager receives an appropriate local event by the car sensors and updates the information structure accordingly. On reception of a safety message, the safety manager in the receiving vehicle updates the information structure that in turn triggers a procedure for information aggregation and the generation of a new safety message if needed. As a result, the safety message is distributed via multiple wireless hops. The distribution of the information is geographically scoped: Every receiver checks whether it is located inside the target area in order to decide whether to accept and forward or to drop the message. In the basic version, the message is forwarded when the receiver is located inside the target area while extensions check whether the receiver has further neighboring vehicles located inside the target area. In case no such neighbor(s) exist(s), no further message is issued.

4.5.2 TERGF Definitions and Description

This section presents the basic TERGF algorithm for the efficient, reliable and aggregated distribution of safety information in a geographical target area over time, i.e., considering the fluctuation of vehicles in the target area. The following section first defines reliability, as required for the distribution of safety information in VANETs.

4.5.2.1 Reliability Definitions for Broadcast and Flooding

Reliability means that a message transfer to one or more receivers is guaranteed. However, different definitions exists for reliability of broadcast message transfer. This section first defines *simple reliable flooding (SRF)*, as follows:

Definition 1: Simple Reliable Flooding.
A message is reliably transmitted if every node receives the information at least once.

The following definition of *simple reliable geographical flooding (SRGF)* extends the previous SRF definition by a spatial component:

Definition 2: Simple Reliable Geographical Flooding.
A message is reliably transmitted if every node located inside the geographical target area receives the information at least once after the message has been initially distributed.

Finally, the following definition of *time-extended reliable geographical flooding (TERGF)* additionally considers a temporal component:

Definition 3: Time-Extended Reliable Geographical Flooding.
An information[2] *is reliably distributed if every node located inside the geographical target area within a duration of time $T = t_2 - t_1$ receives the information at least once. The timestamp t_1 represents the point of time when the safety event occurs, and timestamp t_2 defines the point of time when the safety event expires. Consequently, T refers to the lifetime of the event.*

Note that the time-extended reliable geographical flooding definition three explicitly includes the information of vehicles that enter at a later point in time, i.e., after the safety message has been initially distributed.

While reliability according to definition one and two can be achieved by existing algorithms, as surveyed in Section 4.3, the TERGF algorithm focuses on the latter definition of reliability.

[2] It is worth noting that the definition of TERGF focuses on the distribution of information rather than the packet / message-based definitions of SRF and SRGF.

4.5.2.2 TERGF Algorithm Description

The algorithm works as follows: A vehicle (originator) that wishes to broadcast a message creates a list of its single-hop neighbors that are located inside the geographical target area. It adds this list and the coordinates of the target area to the message and broadcasts it. Furthermore, the originator maintains a *confirmation list*. It adds the neighbors included in the message to the confirmation list and starts two timers: One timer for the lifetime of the event T_e and another re-transmission timer T_{rt}.

A vehicle (forwarder 1) that receives the message and has the safety event already registered in its safety information structure replies an explicit acknowledgment (ack). Another vehicle (forwarder 2) that has not yet registered the safety event in its safety information structure executes the following procedure: It creates the entry in the safety information structure and compares the list of neighbors in the message with its local neighbor list. In case forwarder 2 has more neighbors than the neighbor list included in the message (i.e., excluding its own address), it generates a new message, adds its own single-hop neighbor list and broadcasts the message. This message, however, can be different from the previous.

If the originator receives (i.e., overhears) a message from one of its neighbors with the same ID and geographical area, it interprets this message as implicit acknowledgment and marks the correct reception by this vehicle in its confirmation list. In case the retransmission timer expires before the originator receives an implicit or explicit acknowledgment, the message is rebroadcast. The maximum number of rebroadcasts is limited by the lifetime of the event or a maximum retransmission counter.

The algorithm explicitly covers the case when a new vehicle enters the target area and the transmission range of an informed vehicle. Upon a *new neighbor event* of a vehicle that is not yet registered as informed in the safety information structure, the safety information manager re-generates the respective message and redistributes the safety information based on the algorithm as explained before.

Figure 4.7 illustrates the basic algorithm for reliable distribution of safety messages. Vehicles are drawn by circles, and solid lines indicate single-hop connectivity between vehicles. In the example, the originator A distributes a safety information via the (geo-)broadcast message (1) which includes the addressed neighbors B, C and D. Vehicle E drops the message since it is located outside the target area and not addressed in the header. The potential forwarders C and D do not redistribute the information because there are no further (i.e., additional) neighbors in the target area. Therefore, vehicles C and D reply an explicit acknowledgment (ack) to originator A. Since vehicle B is aware of a further neighbor, namely vehicle F, it generates message (2) and again includes the list of its current, single-hop neighbors. When vehicle A overhears message (2), it interprets the message as a passive (i.e., implicit) acknowledgment and marks the respective entry in its confirmation list.

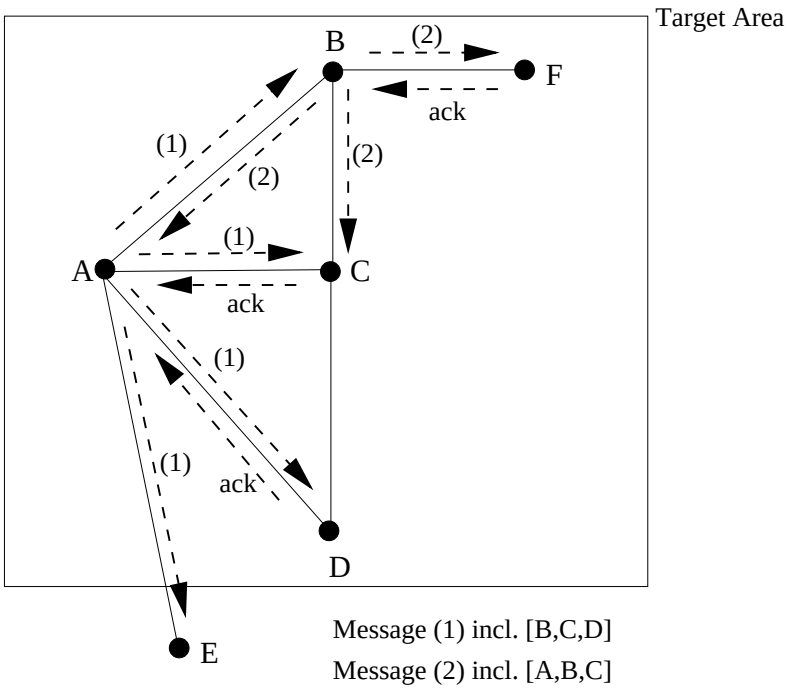

Figure 4.7: Basic algorithm for reliable distribution of safety information.

4.5.3 TERGF Extensions

4.5.3.1 Append Option for Neighbor List

Using the append option for the neighbor list, a forwarding vehicle appends those of its single-hop neighbors to the neighbor list in the message that are not yet included. This is in contrast to the basic algorithm where the single-hop neighbors of the forwarding vehicle replaces the neighbor list in the message.

The benefit of this option is that the vehicles in the target area have a more accurate view on the informed vehicles inside the target area. If vehicles move within the target area, a redistribution of messages to already informed vehicles can be avoided. However, as a tradeoff, the increased accuracy comes at the cost of increased message sizes.

4.5.3.2 Backoff Timer for Redistribution of Safety Information

This extension introduces a timer before redistributing information. Thus, it improves efficiency by avoiding duplicate messages when two or more forwarders have the same further neighbor in common. Figure 4.8 illustrates this scenario.

The potential forwarders B and C receive a safety message from vehicle A. In the basic version, both forwarders would rebroadcast the information in order to inform the further neighbor, namely vehicle D. However, this results in a duplication of the broadcast message about. In order to optimize the forwarding of information

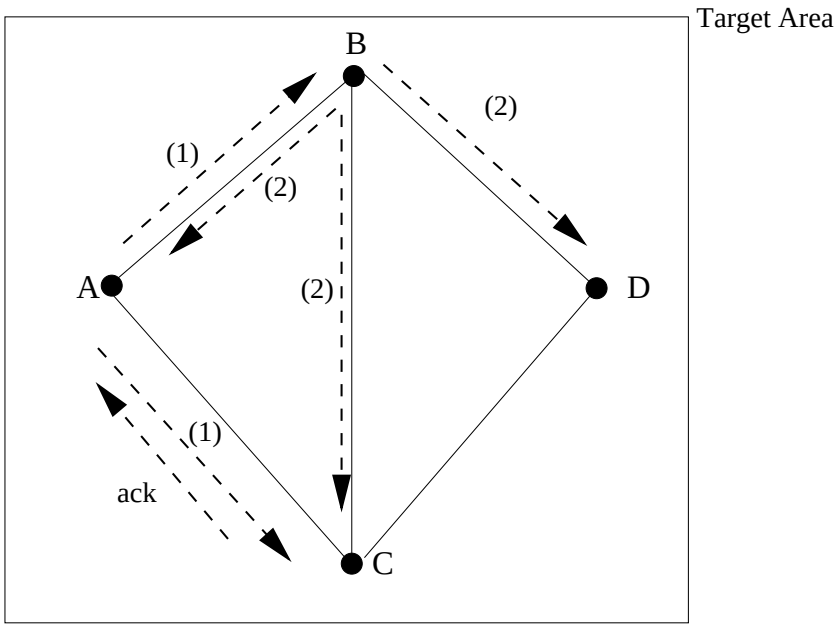

Figure 4.8: Extended TERGF algorithm: Backoff timer.

in this scenario, a single forwarder is selected by means of a (small) contention timer: All forwarders initiate a contention timer before forwarding the message. The initialization either might be random or e.g., according to the position in the target area. The vehicle with the shortest timer, i.e., vehicle B in the example, forwards the information. Vehicle C overhears this message and compares it with its entries in its safety information structure and its local neighbor list. Vehicle C realizes that the information is relayed to common neighbor vehicle D. Consequently, it suppresses the message and resets the contention timer. Since vehicle C does not forward the information, it replies with an explicit acknowledgment to vehicle A.

4.6 Simulative Evaluation

This section presents the simulative evaluation of the TERGF algorithm for the distribution of information in a geographical target area over time, including a comparison to standard GeoBC. It comprises a description of the simulation scenario and environment, as well as the definition of the metrics topological change rate (TCR), information distribution ratio and redundant packet repetition ratio. The final subsection explains the simulation results for a typical scenario in detail.

4.6.1 Scenario and Simulation Environment

The study uses the network simulator ns-2 [132] and considers moving vehicles on a highway, according to the validated movement patterns of [45, 80]. Again, the

highway scenario is chosen as the most dynamic and, thus, the most challenging scenario. We believe that the algorithm can be adapted to city scenarios which is part of the future work. These basic settings are similar to the previous simulations of the vehicular transport protocol (VTP) in Chapter 3, using the same spatial distribution of vehicles and the same mobility behavior. The simulations focus on a 10 km stretch of unidirectional highway and consider different densities of vehicles in scenarios that have different numbers of lanes in each direction and varying numbers of vehicles per kilometer.

All vehicles are equipped with an IEEE 802.11b wireless radio interface that covers a transmission range of 250 m. Vehicles that are located within each other's transmission range can communicate directly whereas communication between vehicles outside of each other's transmission range requires multi-hop forwarding of packets. In order to evaluate the information distribution in a geographical target area, each simulation run selects a vehicle as originator that determines the size and position of the rectangular target area. This selection is arbitrary, but ensures that the target area remains inside the boundaries of the overall simulation area. The target area sizes in the simulation vary between 50 m and 2000 m and always cover the complete street width.

Figure 4.9 illustrates the movement and communication scenario, as described above.

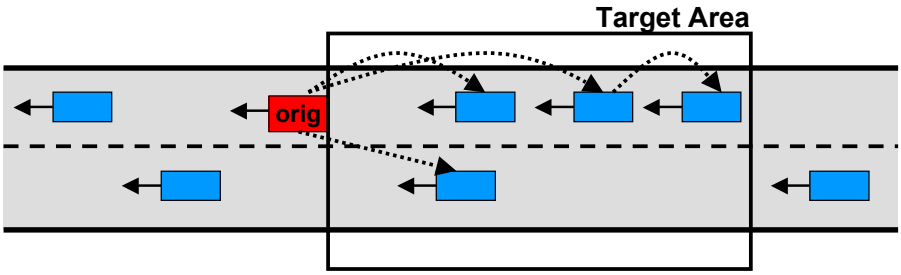

Figure 4.9: Information distribution in a geographical target area in the highway scenario.

The information distribution uses either TERGF or GeoBC. When the target area size is smaller than the radio transmission range, a direct, single-hop broadcast reaches all the nodes inside the target area. In case the target area size is larger than the radio transmission range, multi-hop forwarding is required. However, even multi-hop forwarding cannot guarantee to inform all nodes, e.g., due to temporal network partitions that interrupt the connectivity to all nodes in the target area. The target area might be *in front* or *behind* the originating vehicle, representing different safety applications, such as emergency vehicle approaching from the back and the vehicles in front which will be affected next have to be warned or traffic jam ahead and all vehicles behind the traffic jam should be warned, respectively.

The vehicles move along the highway over time, resulting in frequent topology changes and fluctuation of nodes in the target area. Similar to PBR [93], the vehicles use beaconing to maintain the presence and position of their single-hop neighbors. Since beacons include the geographical position of the sending vehicle, the detection of vehicles that enter the target area uses beaconing. The beacon interval is set to 0.3 s to achieve a high accuracy of neighbor information. The beacon expiry interval at which a vehicle removes neighbor entries when no further beacon has been received is 3.5 times the beacon interval. Furthermore, the transmission time of beacons is randomly jittered to avoid collisions.

The simulation considers safety applications where the vehicle that senses the hazard continues moving, such as extended brake light or icy road. In contrast to safety applications where the vehicle stops (e.g., breakdown or traffic jam), the scenario with ongoing movement represents a *worst case* for the information distribution since all informed vehicles might temporally leave the target area.

The simulations us a message size of 1000 Bytes, allowing to include safety-relevant information, such as event target area related information or even alternative routes. The overall simulation time per run is restricted to 60 s by the given movement patterns. The lifetime of a safety event is set to 30 s, to assure that the network (e.g., neighbor tables) is established before initiating the information distribution.

4.6.2 Metrics

This section defines the metrics for the simulative evaluation of information distribution in a geographical target area.

Topological change rate. Topology change rate (TCR) describes the dynamics of the vehicles and the resulting topological changes over time. The simulations accumulate the average number of vehicles that leave and enter the geographic target area versus the total number of nodes in the target area over time.

Information distribution ratio. Information distribution ratio denotes the ratio of informed vehicles over all vehicles in the geographic target area over time. A vehicle is classified as informed when it receives a packet with a respective safety information at least once. In contrast, all vehicles that are located in the target area contribute to the total number of vehicles in the denominator of the ratio. The simulation observes the information distribution ratio over time. Thus, the fluctuation of vehicles is considered, but the ratio considers only vehicles that are located inside the target area at the respective time interval.

Redundant packet repetition ratio. Redundant packet repetition ratio denotes the average number of redundant packets over average number of total packets in the geographical target over time. A packet is classified redundant when it does not reach uninformed nodes (i.e., either it reaches only already informed nodes or there are no neighbors within transmission range). Consequently, a packet that reaches at least one uninformed node is classified as required and contributes to the total number of messages, but not to the redundant counter. Note that packet transmission is broadcast, and a single transmission might reach multiple receivers. Thus, the transmission of a packet on the wireless interface is considered for this metric.

4.6.3 Simulation Results

This section presents the simulation results for the distribution of safety information in a geographical target area in a highway scenario.

At first, the simulations quantify the topological change rate (TCR) of vehicles entering and leaving the target area for different target area sizes. Afterwards, the section presents the simulation results for the metrics information distribution ratio and total versus redundant packet retransmissions for GeoBC and TERGF, as well as a comparison of both algorithms.

Due to space restrictions, the complete results for all vehicle densities and target area sizes cannot be shown in detail, but are available in Appendix B. This section concentrates on a typical weekday traffic scenario with two lanes per direction (lpd) and six nodes per kilometer (npkm). Furthermore, the simulations consider safety applications that address vehicles in a target area *behind* the originator (i.e., opposite to the driving direction), such as hazard warnings or extended brake lights, as well as applications that address vehicles *in front* of the originator, such as an emergency vehicle approaching. The results in this section focus on scenarios with the target area *behind* the originator. The results related to the target area *in front* of the originator are included in Appendix B. Finally, the results assume that the originator continues moving, as explained in Section 4.6.1. All simulation results are derived from the available 60 movement pattern samples per scenario.

4.6.3.1 Topological Change Rate Simulation Results

This section presents the TCR results for two different target area sizes which represent the number of vehicles that leave and enter the geographic target area over time, as defined in the metrics in Section 4.6.2. Figure 4.10 shows the TCR over time for the target area lengths 50 m and 1000 m. These results utilize a sampling interval of 0.5 s, as the highway movement patterns adjust the position of the vehicles in discrete time steps of 0.5 s accordingly.

Figure 4.10(a) illustrates the TCR over time for a target area size of 50 m, which fluctuates between 0.6 an 1.2. Note that in this scenario the TCR reaches

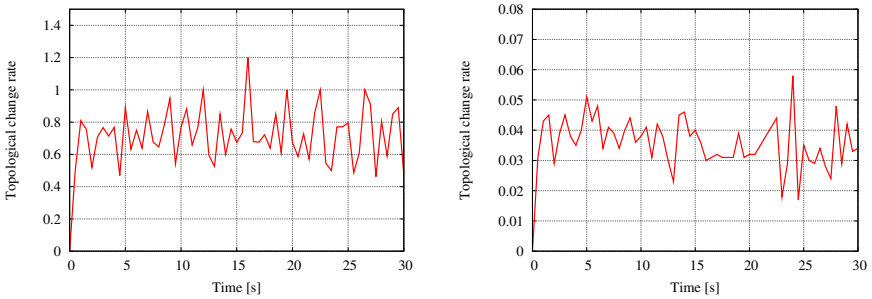

(a) Topological change rate for target area length 50 m.

(b) Topological change rate for target area length 1000 m.

Figure 4.10: Topological change rate for different target area sizes.

beyond *one* that is the number of vehicles entering and leaving the target area is greater than the number of nodes inside the target area. This happens particularly in scenarios with a small geographical target area since the overall number of nodes inside the area is small. As an example, at time 16 s, there are only five vehicles in the target area. In the subsequent interval, four vehicles enter and two vehicles leave the target area. Consequently, the rate is six vehicles that enter and leave over five vehicles that remain in the area, which results in a TCR of 1.2.

Figure 4.10(b) shows the TCR over time for a target area of 1000 m, which fluctuates between 0.03 and 0.05. With the increase of the target area size, the TCR significantly decreases because the total number of vehicles in the target area increases. The average total number of vehicles is 0.6 in a target area of 50 m is whereas it is 12 in a target area of 1000 m. Detailed information and graphs about the average number and standard deviation of total vehicles for different target area sizes are attached in Appendix B.

The TCR results of this section show that frequent fluctuations of vehicles in the target area occur whereas the amplitude of the TCR depends on the target area size and the associated total number of vehicles in the area. These results validate the demand for a reliable distribution of information over time in VANETs in order to inform vehicles that enter the target area after the initial issuing of the message.

The following sections present the simulation results for GeoBC and TERGF information distribution respectively, before comparing both algorithms.

4.6.3.2 Geo-Broadcast Simulation Results

This section presents the simulation results for the information distribution ratio and total versus redundant packet rate for the GeoBC algorithm, according to the specified scenario of 2 lpd, 6 npkm and the target area behind the originator.

The section categorizes the results in target area sizes smaller and greater than the radio transmission range for both metrics, due to the effects of single-hop or multi-hop distribution of messages in the target area, as the following example indicates.

In case the target area is smaller than the radio range, the initial message by the originator typically reaches all vehicles, and rebroadcasts by vehicles in the target area are redundant.

In case the target area is greater than the radio range, even multi-hop forwarding of messages might not reach all vehicles in the target area due to temporal network partitioning. Consequently, GeoBC may not reach an information distribution ratio of 100% in this scenario, even after the initial distribution of the message.

Information Distribution Ratio

The information distribution ratio collects the number of informed vehicles over all vehicles in the target area, as defined in Section 4.6.2. This section presents the information distribution rate over time and the cumulative distribution function (CDF) for target area sizes smaller and greater the radio transmission range.

Target Area Size Smaller Than Radio Range

Figure 4.11 shows a sample of an information distribution ratio over time for a target area size of 50 m, using GeoBC as the distribution algorithm. Note that the figure intentionally depicts a single sample because the average information distribution ratio in this scenario modifies the shape of the curve to a linear decrease due to varying durations when all nodes are informed in the different simulation runs. In some specific samples, there is no vehicle in the target area at the beginning whereas the majority of samples show a sharp decline of the information distribution ratio, as illustrated, but with different durations.

In Figure 4.11, the original message reaches all vehicles in the target area (i.e., single-hop), and the information distribution ratio is 100% immediately after the message transmission. At time 0.8 s, a new vehicle enters the target area. Since GeoBC does not retransmit the message, the information distribution ratio drops to 50%. At time 1.9 s, the information distribution ratio drops to zero since all informed vehicles leave the target area. In GeoBC, the information cannot be kept longer inside the target area than informed vehicles remain there.

The average time that the GeoBC algorithm is able to maintain an information distribution ratio of 100% for a target area size of 50 m is 2 s.

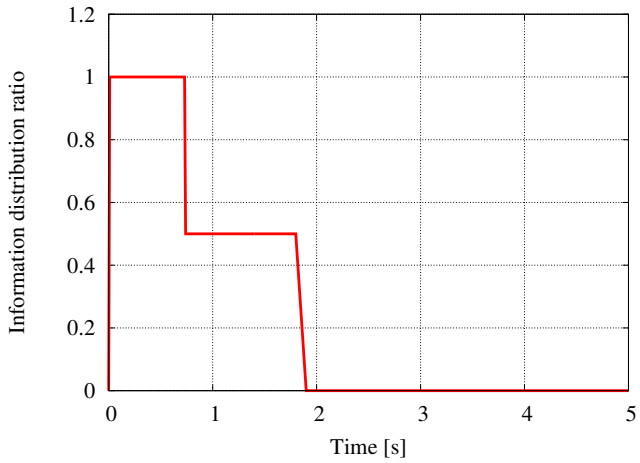

Figure 4.11: GeoBC information distribution ratio sample for target area length 50 m.

Furthermore, Figure 4.12 evaluates the simulation results of the information distribution ratio in a cumulative distribution function (CDF). The CDF describes a statistical distribution. The value in a CDF at each possible outcome represents the probability of receiving exactly that outcome or a lower one.

In this section, the CDF is used to predict the maximal information distribution duration, i.e., that all vehicles (or a percentage of the vehicles) remain informed. Therefore, the information distribution CDF accumulates the duration of the samples when the information distribution ratio decreases below 100%. Thus, the results should be interpreted as the time until the information is lost, i.e., not all vehicles (or the percentage of vehicles) are informed anymore. When the CDF reaches one, the information distribution ratio is zero, i.e., the information is not anymore in the target area in all cases of the simulations.

For the target area size of 50 m, after two seconds, only in 2% of the samples all vehicles are informed. The probability that all vehicles are informed for more than four seconds is zero. With increasing target area size, the information periods increase. For a target area size of 100 m, 10% of the samples keep all nodes informed up to three seconds. The information distribution ratio drops below 100% after 6 s. For a target area size if 250 m, the probability to keep all nodes informed up to three seconds is 35%, and none of the samples can keep all nodes longer informed than 8 s.

Summarizing, GeoBC is able to inform all vehicles that are in the target area upon the transmission of the message via single-hop broadcast when the target area is smaller than the radio transmission range. When further vehicles enter the target area, the message is not rebroadcast. Thus, the entering vehicles are not informed,

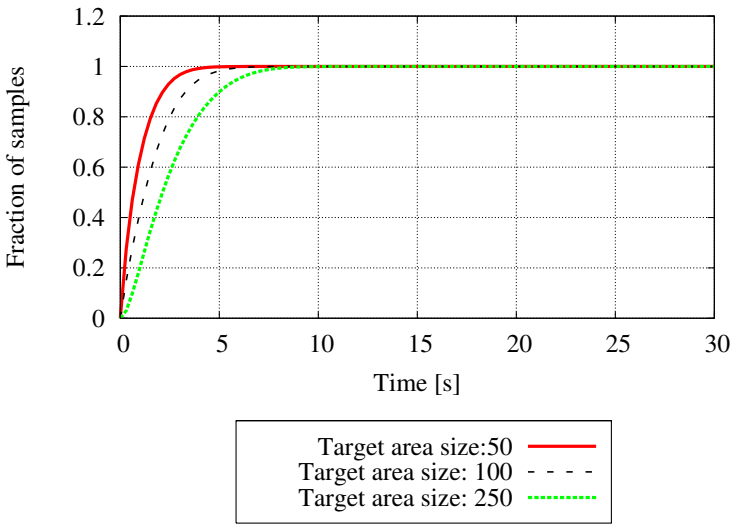

Figure 4.12: CDF of the loss of GeoBC information distribution to all nodes for different target area lengths below radio range.

and the information distribution ration decreases. When all informed vehicles leave the information area, the information is completely lost, i.e., no vehicle is informed anymore. Due to the high TCR for small target areas, the probability that all nodes remain informed is below three seconds, depending on the actual target area size.

Target Area Size Greater Than Radio Range

In case the target area size exceeds the radio transmission range, multi-hop forwarding (i.e., rebroadcasting) is required in order to reach all vehicles inside the target area. Even with rebroadcasts, in some specific scenarios the information cannot reach all the nodes, i.e., when the network is partitioned. Figure 4.13 shows such an example for a target area size of 1000 m. Again, the figure represents a single sample for illustration and explanation.

In Figure 4.13 the GeoBC message reaches only 75% of the vehicles inside the target area due to a network partition. Even when connectivity resumes shortly after the message transmission, the additional vehicles are not informed because the message is only issued once. With the fluctuation of vehicles in the target area, the information decreases continuously as the vehicles leave the target area. After 21 s, all informed vehicles have left the target area, and the information distribution ratio decreases to zero.

The average duration for which GeoBC is able to keep all vehicles inside a geographical target area of 1000 m informed is 4 s.

Figure 4.14 shows the CDF of the GeoBC message distribution in a target area

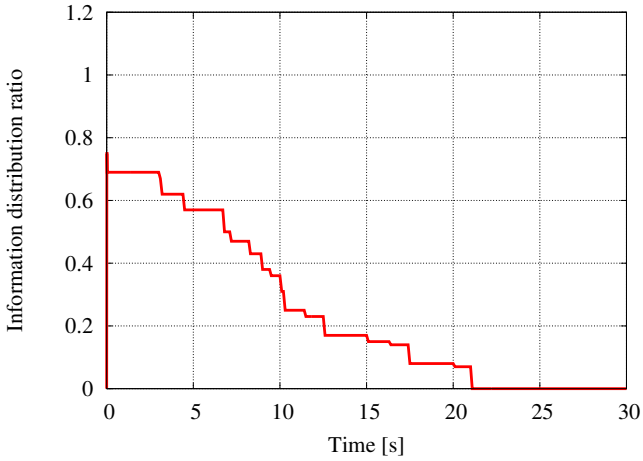

Figure 4.13: GeoBC information distribution ratio sample for target area length 1000 m.

of 1000 m. The CDF includes curves for the duration when 75% and 50% of the vehicles in the target area are informed since the information of *all* (i.e. 100%) vehicles is improbable due to network partitions, as explained above.

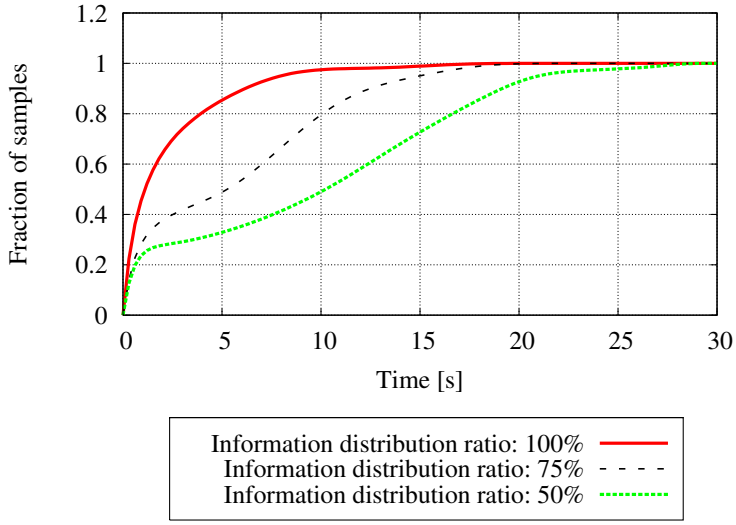

Figure 4.14: CDF of the loss of GeoBC information distribution to different percentages of nodes for different target area lengths above radio range.

The probability that the GeoBC algorithm is able to inform all nodes (i.e., 100% curve) in the target area of 1000 m decreases exponentially. After 5 s, only 14% of the samples can keep all vehicles informed. The fraction of samples de-

creases to 3% after 10 s and at maximum 18 s the probability that all vehicles are informed is zero.

The curve for an information distribution ratio of 75% shows that half of the samples can keep three-quarters of the vehicles informed for up to 5 s. However, after 15 s only in 2.5% of the samples 75% of the vehicles are still informed, and the information is lost in all samples after 18 s, as well.

Finally, half of the vehicles are informed up to 10 s in 50% of the samples. After 20 s, the fraction of samples with 50% of the vehicles informed is below 10%, and half of the vehicles are not informed for up to 27 s.

Summarizing, for target area sizes greater than the radio transmission range, multi-hop forwarding of messages is required to inform all vehicles in the target area. However, not all vehicles inside the target area might be reached in case the network is partitioned. The CDF shows that the fraction of samples that keep all nodes informed decreases exponentially. For a target area size of 1000 m, all vehicles are informed for 4 s on the average. Furthermore, the CDF evaluates the information distribution ratio of 75% and 50% of the vehicles. The results show that not even half of the vehicles in the target area can be informed for a lifetime of 30 s of the event.

Redundant Packet Repetition Ratio

This section presents the results of total versus redundant message transmissions in GeoBC. A message transmission is considered redundant when the message does not reach and inform additional neighbors, according to the metric definition in Section 4.6.2. Again, the results are presented according to target area sizes, i.e., smaller and greater than the radio transmission range.

Target Area Size Smaller Than Radio Range

This section compares the total number of packets and the redundant number of packets for target area sizes smaller than the radio transmission range.

Figure 4.15 shows the average number of total versus the average number of redundant GeoBC packets for a target area size of 50 m. GeoBC transmits on the average 1.2 packets whereof 0.8 packets are redundant. Note that for a target area size below radio transmission range it requires only a single packet to inform all vehicles in the target area. The average number of required packets is even below one because no message should be transmitted when the target area is empty. Thus, for a target area size of 50 m, the average required number of packets is 0.4. Furthermore, all GeoBC packets are transmitted within the first evaluation time interval. Consequently, vehicles entering after that distribution period are not informed.

Figure 4.16 illustrates the number of total and redundant GeoBC packets separately in order to show the standard deviation.

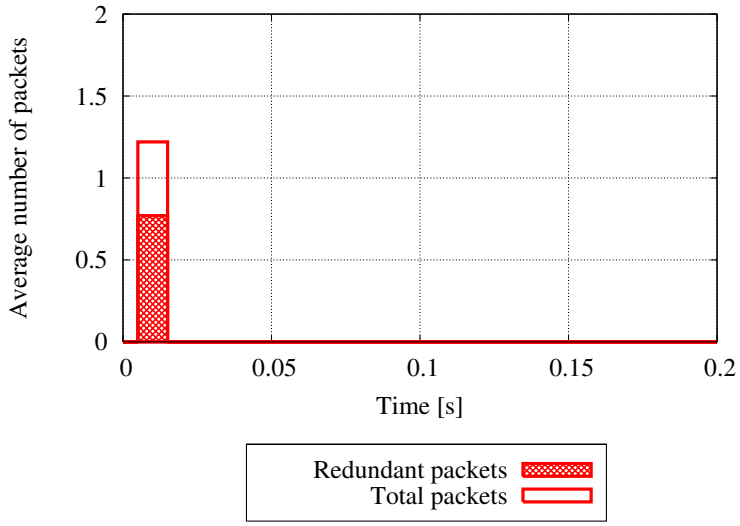

Figure 4.15: GeoBC average number of total versus redundant packets for target area length 50 m.

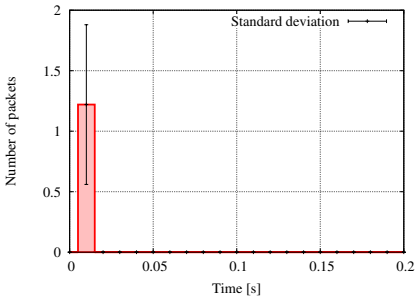

(a) Average number and standard deviation of total packets for target area length 50 m.

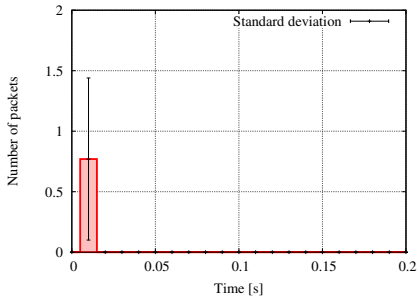

(b) Average number and standard deviation of redundant packets for target area length 50 m.

Figure 4.16: Average number and standard deviation for total and redundant packets for different target area sizes below radio range.

Table 4.1 summarizes the average number of total and redundant packets and provides the respective ratios. With increasing target area size (i.e., as long as the target area size remains smaller than the radio range), the ratio of redundant over total packets increases because there are more vehicles in the target area that rebroadcast a packet. Note, that the ratio for 50 m and 100 m is similar, due to the average small number of vehicles in the target area.

TA	Avg. redundant	Avg. total	Samples	Ratio
50 m	0.77	1.22	60	0.630
100 m	1.23	1.97	60	0.627
250 m	2.97	3.97	60	0.748

Table 4.1: GeoBC: Redundant packet repetition ratio for different sizes of target areas below radio range.

Summarizing, for target area sizes below the radio transmission range, the initial GeoBC message by the originator typically reaches all vehicles in the target area. The average number of required messages is even below one because no message is required when the target area is empty. In GeoBC, each vehicle rebroadcasts each message once, resulting in redundant retransmissions. The ratio of redundant over total message increase from 63% to 74,8% when increasing the target area size from 50 m to 250 m, respectively. All retransmissions occur upon the initial distribution of the message.

Target Area Size Greater Than Radio Range

This section compares the total number of GeoBC packets and the redundant number of packets for target area sizes greater than the radio transmission range.

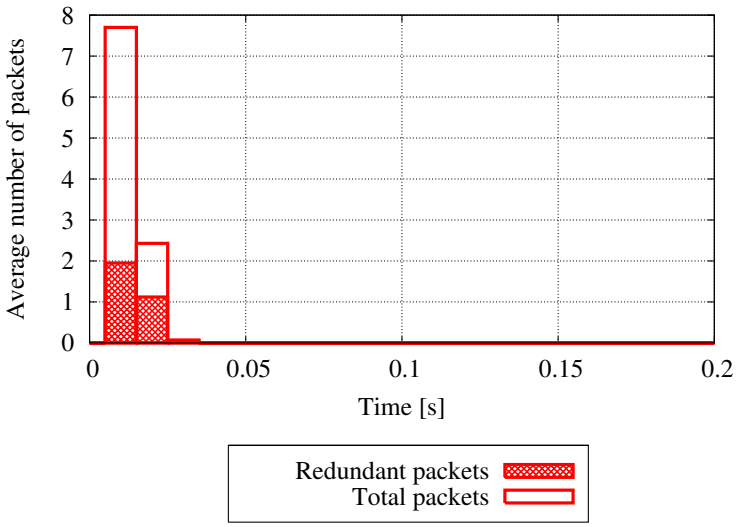

Figure 4.17: GeoBC average number of total versus redundant packets for target area length 1000 m.

Figure 4.17 compares the total and redundant number of GeoBC packets for a target area size of 1000 m. The total and required number of packets increase disproportionally high because each rebroadcast along a straight road covers an ad-

ditional section of the highway, according to the distance to the predecessor. Consequently, several rebroadcasts reach new, additional neighbors whereas a more advanced scheme could select a single, optimized forwarder that maximizes the covered road segment and reduce the number of required messages. Furthermore, the messages are distributed to several intervals, due to the multi-hop relaying of messages in target areas that are greater than the radio transmission range.

Again, Figure 4.16 illustrates the number of total and redundant GeoBC packets separately in order to show the standard deviation.

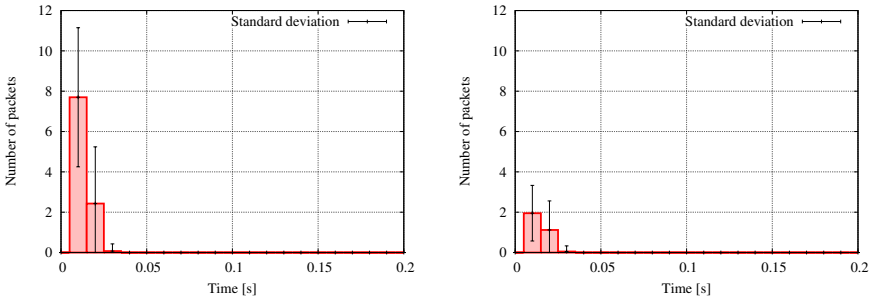

(a) Average number and standard deviation of total packets for target area length of 1000 m.

(b) Average number and standard deviation of redundant packets for target area length of 1000 m.

Figure 4.18: Average number and standard deviation for total and redundant packets for different target area sizes above radio range.

Table 4.2 summarizes the average number of total and redundant messages and calculates the respective ratios.

TA	Avg. redundant	Avg. packets	Samples	Ratio
500 m	2.60	6.85	60	0.380
1000 m	3.12	10.20	60	0.306
2000 m	3.50	13.28	60	0.263

Table 4.2: GeoBC: Redundant packet repetition ratio for different sizes of target area above radio range.

For increasing target area sizes the ratio decreases since more vehicles are present in greater target areas. This increases the probability that a rebroadcast reaches an uninformed vehicle. As stated before, the radio transmission range of each rebroadcast covers an additional segment of the road, according to the distance to the predecessor. Therefore, even each of two rebroadcasts of vehicles in only a couple of meters distance might both reach uninformed nodes.

Summarizing, this section evaluates the total and redundant messages of GeoBC for target areas greater than the radio transmission range. The ratio of redundant over total GeoBC messages decreases because the probability that a rebroadcast reaches uniformed neighbors is high along a road segment. Consequently, the required number of GeoBC messages increases above average which can be avoided by a more advanced scheme that selects a forwarder covering a maximal additional area.

4.6.3.3 TERGF Simulation Results

This section evaluates the TERGF algorithm according to the same methodology used for the GeoBC evaluation in the previous Section.

Information Distribution Ratio

Once more, the information distribution ratio evaluates the number of informed vehicles over all vehicles in the target area, as defined in Section 4.6.2. This evaluation uses the TERGF algorithm for information distribution. The results distinguish between target area sizes greater and smaller than the radio transmission range and show a sample of the information distribution ratio over time, as well as the CDF for both scenarios.

Target Area Size Smaller Than Radio Range

For target area sizes smaller than the radio range, the initial message reaches all vehicles inside the target area via single-hop broadcast, similar to GeoBC. In contrast, an informed vehicle that uses the TERGF algorithm generates a new message when a new, not-informed vehicle enters the transmission range. Consequently, the information is kept inside the target area as long as any informed node remains inside the target area. Only in case all informed nodes leave the target area (i.e., the target area is temporally empty), the information is lost.

Figure 4.19 illustrates a typical sample for the information distribution via TERGF in a target area of 50 m. Upon the initial message transmission of the originator at time zero, all vehicles receive the information, and the information distribution ratio is one. At time 0.8 s, a peak drop occurs in the curve because a new node enters the target area. The information distribution ratio drops for the time required to detect the new vehicle (i.e., until the vehicle announces its presence by a beacon) and the generation and transmission of the respective message. When the vehicle that enters the target area receives the safety message, the information distribution ration goes back to one. At time 2.1 s, all vehicles leave the target area and, thus, the information distribution ratio drops to zero. Note that in TERGF the information is only lost when all vehicles leave the target area, i.e., when the target area is temporally empty. Once the information ratio reaches zero,

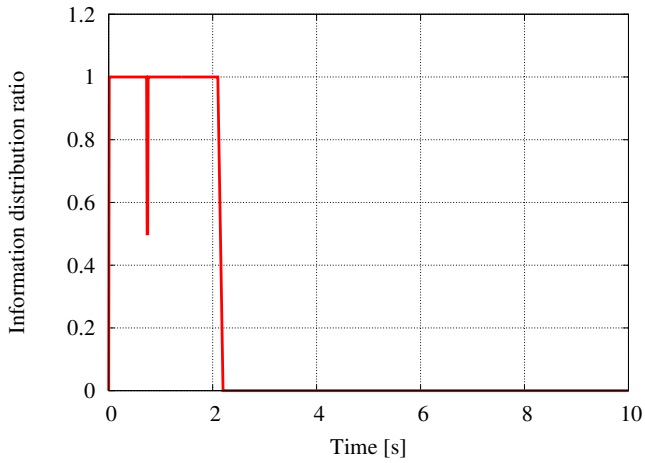

Figure 4.19: TERGF information distribution ratio sample for target area length 50 m.

it remains there because there is no informed node left that could generate a new message for vehicles entering the target area.

In TERGF, the average interval in which all vehicles in a geographical target area of 50 m remain informed is 2.24 s.

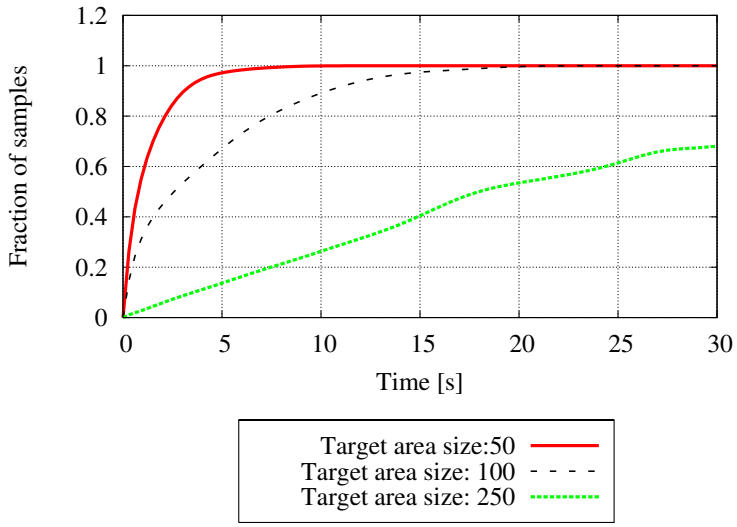

Figure 4.20: CDF of the loss of TERGF information distribution to all nodes for different target area lengths below radio range.

Figure 4.20 shows the CDF of the information distribution ratio to all vehicles (i.e., probability to keep all nodes in the target area informed) for different target area sizes below radio transmission range.

For a target area size of 50 m, only 2% of the samples keep all vehicles informed for more than 5 s. After 8 s, the information is lost in all samples.

In a target area of 100 m, 35% of the samples maintain a distribution ration of one for longer than 5 s. After 10 s, 10% of the samples keep the information and after 20 s the information is lost out of the target area in all samples.

For a target area size of 250 m, 83% of the samples maintain all vehicle informed for more than 5 s and 32% of the samples keep the information for the complete lifetime of the event, i.e., 30 s.

Summarizing, with increasing target area sizes up to the radio transmission range, the probability to keep all nodes informed increases. The TERGF algorithm keeps all nodes informed, including nodes entering the target area over time, until all nodes leave the target area.

Target Area Size Greater Than Radio Range

This section evaluates the information distribution ratio of the TERGF algorithm for target area sizes greater than the radio transmission range.

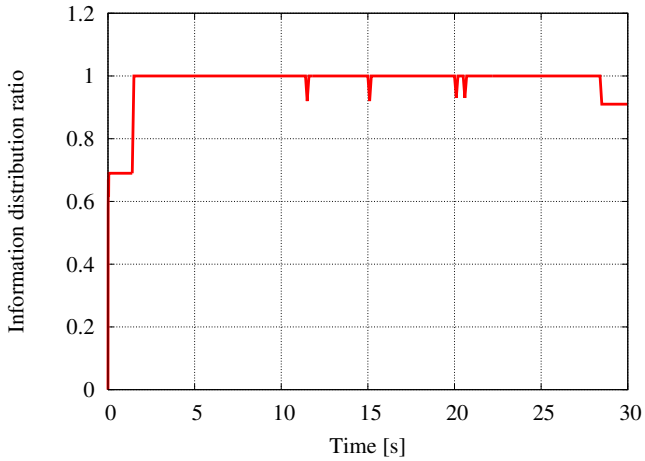

Figure 4.21: TERGF information distribution ratio sample for target area length 1000 m.

Figure 4.21 shows a typical sample of the information distribution ratio for a target area size of 1000 m. The initial safety message by the originator reaches only 75% of the vehicles in the target area, due to a network partition. At time 1.5 s, full network connectivity resumes in the target area because of vehicular movement. When connectivity resumes, the TERGF algorithm on the informed vehicles generates a new message and informs the previously not informed vehicles, resulting in a information distribution ratio of 100%. The graph shows four drops when new vehicles enter the area. Similar to the TERGF results in small target areas, these drops indicate the information of vehicles entering the target area. At time

28.5, further vehicles enter the target area, which cannot be reached, again, due to a network partition. However, TERGF keeps the information in the target area over the lifetime of the safety event.

In TERGF, the average time that all vehicles in a geographical target area size of 1000 m remain informed is 18.46 s.

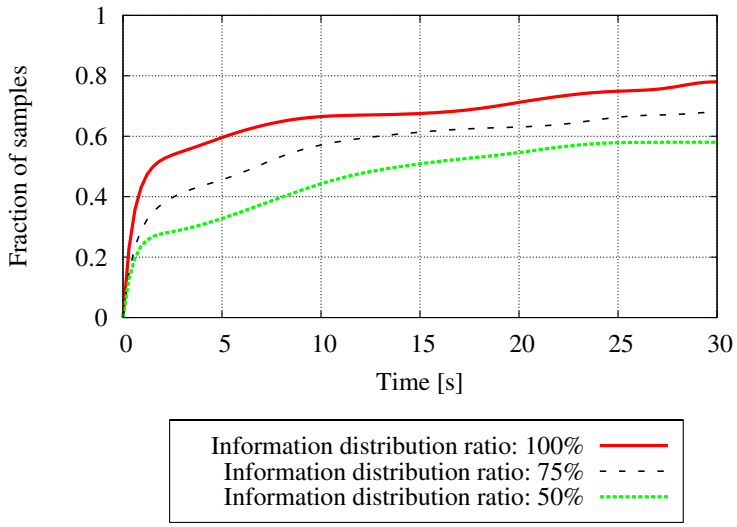

Figure 4.22: CDF of the loss of TERGF information distribution to different percentages of nodes for different target area lengths above radio range.

Figure 4.22 shows the CDF of the information distribution in TERGF for target area sizes greater than the radio transmission range. The figure shows the CDF of information distribution ratios of 100%, 75% and 50%, due to the temporal network partitions in large target areas. TERGF cannot inform all vehicles inside the target area without full connectivity. Consequently, the results include the evaluation of the information of three-quarters and half of the nodes.

21% of the samples keep all nodes informed over the 30 s lifetime of the safety event. In comparison, in 31% of the samples 75% of the nodes remain informed for more than 30 s and 41% of the samples keep half of the vehicles informed over the lifetime.

Summarizing, TERGF provides reliability over time with a high probability since the CDF does not reach one, i.e., the information is not lost out of the target area for the 30 s lifetime of the event.

Redundant Packet Repetition Ratio

This section presents the ratio of total versus the redundant packet transmissions, using the TERGF algorithm. The simulations accumulate all packet transmissions and consider a transmission redundant when no additional neighbor is reached, as defined in Section 4.6.2. Once more, the results distinguish target area sizes smaller and greater than the radio transmission range.

Target Area Size Smaller Than Radio Range

This section compares the total number of packets and the redundant number of packets in TERGF for target area lengths greater than the radio transmission range.

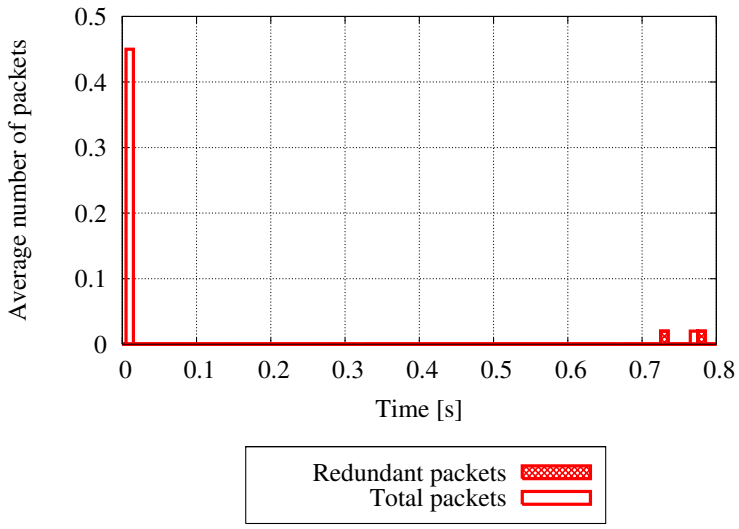

Figure 4.23: TERGF average number of total versus redundant packets for target area length 50 m.

Figure 4.23 shows the average number of total and redundant packets in TERGF for a target area size of 50 m over time. The average number of total packets upon the initial transmission is 0.45 without redundant retransmissions. The average number of total packets is below one because TERGF does not transmit packets when there is no other node in the target area. Furthermore, TERGF avoids redundant packets since the initial packet reaches all vehicles inside the target area and consequently, the vehicles in the target area do not rebroadcast the packet.

Between 0.7 s and 0.8 s, Figure 4.23 shows rebroadcasts when new vehicles enter the target area. These rebroadcasts might include redundant retransmissions. Though each vehicle initiates a backoff-timer upon detection of a new neighbor, several informed vehicles may generate a message to inform the newcomer in parallel when the differences between timer offsets are too small. However, the ratio of the average number of these redundant packets is with 0.03 extremely small.

The standard deviation, which is approximately 0.5 for the total packets upon the initial transmission and 0.02 for total and redundant rebroadcast over time, is omitted in this Figure for better readability. The illustration of the standard deviation is included in Appendix B.

Table 4.3 summarizes the average number of total and redundant packets in TERGF for target area sizes smaller than the radio transmission range and calculates the respective ratio. The number of redundant packets and the ratio is extremely small because redundant transmissions only occur when the difference between timer offsets to inform vehicles entering the target area is too small, and thus two nodes simultaneously generate and transmit a packet to inform the newcomer. With increasing target area size, the ratio decreases since the number of total and required packets increases whereas the average number of redundant packets on the boundary of the target area remains stable.

TA	Avg. redundant	Avg. total	Samples	Ratio
50 m	0.03	0.50	60	0.067
100 m	0.03	1.12	60	0.030
250 m	0.03	2.12	60	0.016

Table 4.3: TERGF: Redundant packet repetition ratio for different sizes of target areas below radio range.

Summarizing, TERGF reduces the average number of total and required packets, as well as the average number of redundant packets, while keeping the information in the target area over time. Besides the initial distribution of safety information, TERGF nodes generate and retransmit information to nodes entering the target area over time.

Target Area Size Greater Than Radio Range

Finally, this section compares the total number of packets and the redundant number of packets in TERGF for target area sizes greater than the radio transmission range.

Figure 4.24 shows the average number of total packets in TERGF for a target area size of 1000 m over time. Redundant packets are not shown in the figure because there are no redundant packets in the simulation results of TERGF for large target area sizes. After the initial distribution of a safety message by the originator, the average number of total (i.e., required) packets is 2.8. Continuing along the time axis, small amplitudes of total packets occur when new vehicles enter the target area and the informed nodes generate and transmit new safety packets. These packets are required to keep all nodes informed, especially including newcomers.

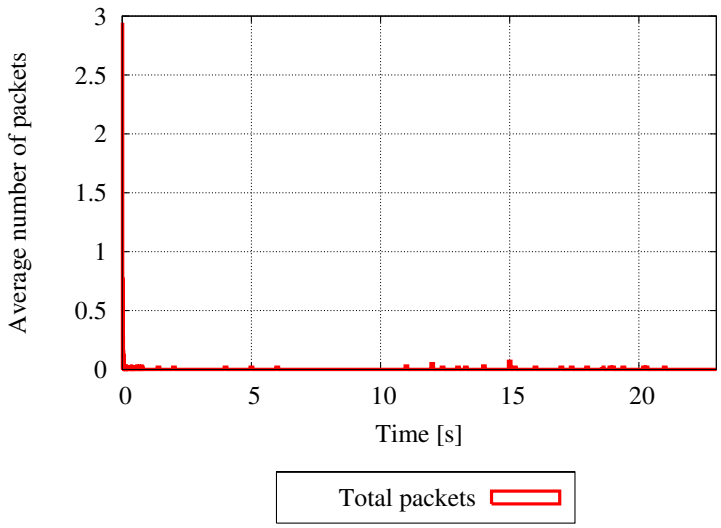

Figure 4.24: TERGF average number of total versus redundant packets for target area length 1000 m.

Table 4.4 shows the average number of total packets in TERGF for different target area sizes greater than the radio transmission range. It reveals that the number of redundant packets is indeed zero for these scenarios. Consequently, the ratio of total over redundant packets is zero.

TA	Avg. redundant	Avg. packets	Samples	Ratio
500 m	0.000	4.48	60	0.000
1000 m	0.000	5.95	60	0.000
2000 m	0.000	8.18	60	0.000

Table 4.4: TERGF: Redundant packet repetition ratio for different sizes of target areas above radio range.

Summarizing, TERGF achieves an significant increase in the distribution of safety information in target area sizes greater than the radio transmission range efficiently, i.e., without redundant retransmissions.

4.6.3.4 Comparison

This section explicitly compares the simulation results of GeoBC and TERGF. The comparison considers the metrics *information distribution ratio* and *total versus redundant packet transmissions* as evaluated separately in the previous sections.

Information Distribution Ratio

The comparison of the GeoBC and TERGF information distribution ratio distinguishes once more between target area sizes smaller and greater than the radio transmission range.

Target Area Size Smaller Than Radio Range

Figure 4.25 compares the GeoBC and TERGF information distribution ratio of a single sample for the target area size of 50 m.

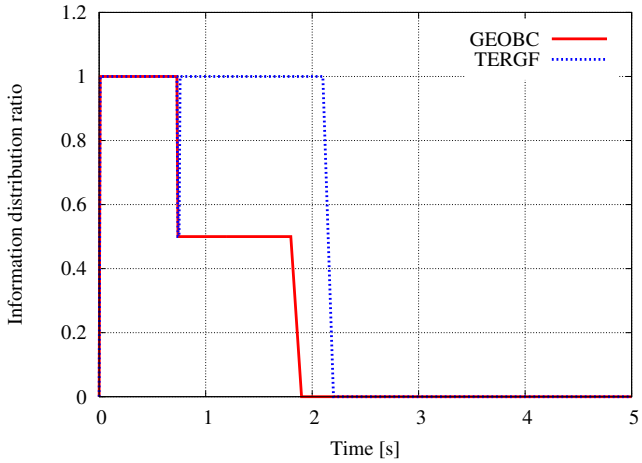

Figure 4.25: Comparison of an information distribution ratio of one sample for target area size of 50 m.

In GeoBC, the information distribution ratio decreases when further vehicles enter the target area, as indicated by the decrease of the curve to 50% after 0.8 s. At 1.9 s, the last vehicle informed by the initial message leaves the area and the information distribution ratio drops to zero.

In contrast, TERGF keeps all vehicles informed as long as a vehicle remains in the target area. The vehicle entering the target area at 0.8 s results only in a short drop. This drop illustrates the time required to detect the new vehicle, generate and transmit the safety message. The information distribution ratio in TERGF resumes to one after the newcomer is informed and remains at this level until all vehicles leave the area at time 2.1 s.

For a target area size of 50 m, TERGF improves the average time that all vehicles in a target area remain informed by approximately 80%, compared to GeoBC. Table 4.5 shows the absolute average values for different target area sizes below radio transmission range.

Figure 4.26 compares the CDF of the GeoBC and TERGF information distri-

TA	GEOBC	TERGF
50 m	1.48 s	2.24 s
100 m	1.94 s	5.69 s
250 m	2.65 s	17.44 s

Table 4.5: Average time that all vehicles remain informed for different target area sizes below radio transmission range.

bution ratio, considering the durations that all vehicles remain informed for a target area size of 250 m.

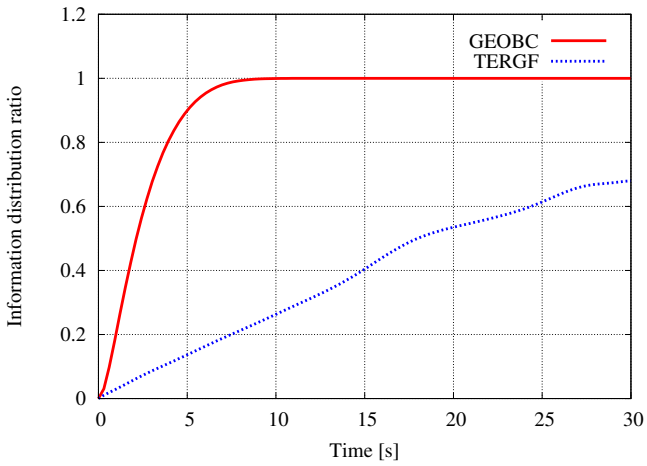

Figure 4.26: Comparison of CDF of a 100%information distribution for target area size of 250 m.

In GeoBC, in 90% of the samples the information is lost after 5 s and it is lost in all samples after 8 s. In contrast, in TERGF 18% of the samples keep the information for up to 5 s, and 68% of the samples keep all nodes informed up to the 30 s lifetime of the event. Consequently, TERGF is able to keep all vehicles informed longer than 30 s lifetime with a probability of 32%.

Summarizing, TERGF improves the information distribution ratio in target area sizes smaller than the radio transmission range significantly. TERGF informs vehicles entering the target area and it keeps the information, as long as an informed nodes remains in the area.

Target Area Size Greater Than Radio Range

Figure 4.27 compares the GeoBC and TERGF information distribution ratio of a single sample for the target area size of 1000 m.

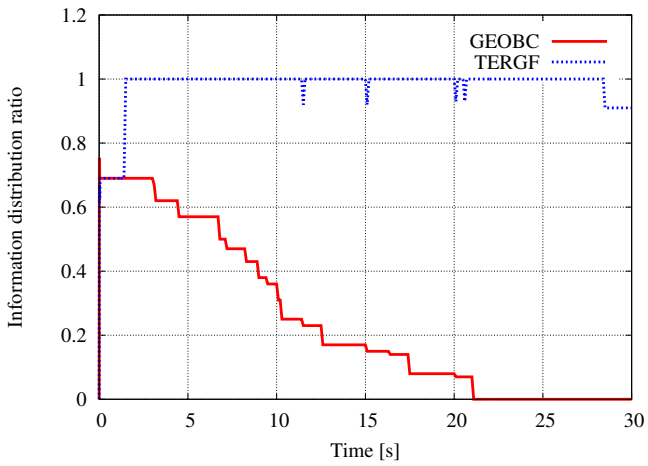

Figure 4.27: Comparison of the information distribution ratio of one sample for target area size of 1000 m.

Upon the initial transmission of the safety message, the network is partitioned which is typical for large target area sizes. Consequently, both algorithms inform 75% of the vehicles with the initial safety message. In GeoBC, the information ratio decreases linearly with vehicles entering the target area. After 21 s, the information distributed with GeoBC is lost out of the target area. In contrast, the TERGF algorithm informs all vehicles as soon as the network partition is over. TERGF keeps all vehicles informed over time. When further vehicles enter the target area, the TERGF algorithm shows short drops, representing the time to detect the vehicle, generate and transmit the respective message. Only when further vehicles enter the target area and the network is partitioned, as at time 28.5 s, TERGF can temporally not inform all vehicles in the target area.

Table 4.6 shows the average durations to keep all nodes informed in target area sizes greater than the radio transmission range. For the target area size of 1000 m, as illustrated in the sample above, the average duration to keep all nodes informed improves by a factor 4.5 when using TERGF.

TA	**GEOBC**	**TERGF**
50 m	3.01 s	17.99 s
100 m	4.00 s	18.46 s
250 m	3.58 s	22.81 s

Table 4.6: Average time that all vehicles remain informed for different target area sizes above radio transmission range.

Figure 4.28 compares the CDF of information ratio to all vehicles in a target area of 1000 m.

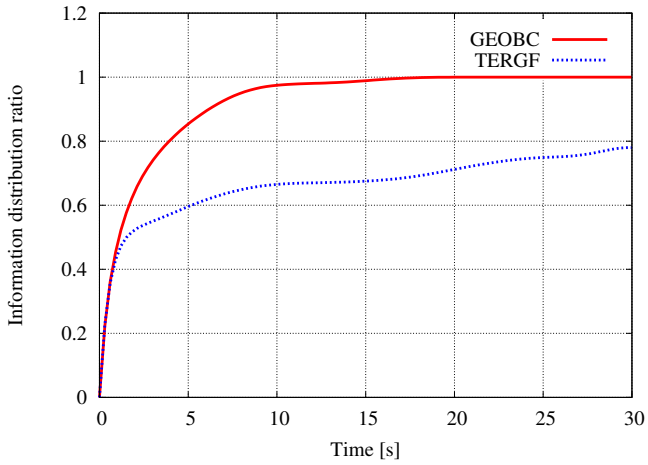

Figure 4.28: Comparison of the CDF of information distribution ratio to all vehicles for a target area size of 1000 m.

In GeoBC, only 15% of the samples keep all nodes informed for more than 5 s, and after 15 s, the information is lost in all the samples. In comparison, TERGF keeps all vehicles informed for more than 5 s in 40% of the cases, and after the lifetime of 30 s still 21% of the samples keep all vehicles informed.

Summarizing, TERGF increases the probability of information distribution significantly also for target area sizes greater than the radio transmission range. TERGF keeps the information in the target area as long as vehicles remain in the target area and these vehicles are fully connected. TERGF informs vehicles entering the target area, as well as when a network partition is over.

Redundant Packet Repetition Ratio

This section compares the ratio of total versus redundant packet transmissions for GeoBC and TERGF, classified according to target area sizes smaller and greater than the radio transmission range.

Target Area Size Smaller Than Radio Range

Table 4.7 compares the redundant packet repetition ratio of GeoBC and TERGF for different target area sizes below radio range.

In comparison to GeoBC, TERGF decreases the ratio of total versus redundant packets by approximately 90% on the average. Furthermore, TERGF reduces the

TA	GEOBC	TERGF
50 m	0.630	0.067
100 m	0.627	0.030
250 m	0.748	0.016

Table 4.7: Comparison of a redundant packet repetition ratio for different sizes of target areas below radio range.

number of total packets, e.g., for a target area of 50 m GeoBC needs 1.2 packets in average whereas TERGF requires 0.45 packets on the average.

Target Area Size Greater Than Radio Range

Table 4.8 shows the redundant packet repetition ratio of GeoBC and TERGF for different target area sizes above radio range.

TA	GEOBC	TERGF
500 m	0.380	0.000
1000 m	0.306	0.000
2000 m	0.263	0.000

Table 4.8: Comparison of a redundant packet repetition ratio for different sizes of target areas below radio range.

TERGF provides efficient distribution of information in the target area, as no redundant packet transmission is observed. Though rare redundant retransmissions may occur when several vehicles detect the same vehicle entering the target area, this ratio is negligibly small, as shown by the simulation results. Also, for large target area sizes, TERGF reduces the number of total packets, e.g., for a target area of 1000 m GeoBC needs 7.7 packets on the average whereas TERGF requires only 2.8 packets on the average.

4.7 Summary and Conclusions

One of the most important applications in vehicular ad hoc networks is traffic and road safety, such as extended brake lights or emergency vehicle approaching. Traffic safety applications demand for efficient and reliable distribution of safety information since this information affects human life and health.

Typically, safety information is beneficial and valid within restricted geographical boundaries. Presently, there are schemes to address all vehicles in a geographical target area, termed GeoBroadcast (GeoBC). However, GeoBC does not provide any means of reliability, and GeoBC causes a high network load due to retransmission of each packet by every vehicle. Consequently, more advanced algorithms

for the reliable and efficient distribution of safety information in a spatial area is required.

Furthermore, the high degree of mobility in VANETs causes vehicles to frequently enter and leave the target area. This high topological change rate (TCR) further modifies the definition of reliability for VANETs: According to the lifetime of a safety event, reliability in VANETs includes the distribution of safety information to vehicles that enter the target area after the initial transmission of the safety message.

Consequently, VANETs demand for efficient and reliable distribution of safety information in a geographical target area over time, i.e., during the lifetime of a safety event.

This chapter has proposed a time-extended reliable geographical flooding (TERGF) algorithm. TERGF aims at efficient and reliable information distribution in a geographical target area over time. TERGF introduces a information management component and information structures since TERGF focuses on distribution of information rather than on traditional packet-based reliability. TERGF avoids redundant retransmissions by adding a list of single-hop neighbors that are located inside the target area to the geographical addressing in the message header. Furthermore, TERGF provides reliability in a single-hop scope by retransmissions, based on passive acknowledgments. A sender maintains a confirmation list of its neighbors. When the sender *overhears* the forwarding of the information by a successor, it marks the respective node as confirmed. A timer schedules retransmissions of unacknowledged packets in absence of acknowledgments. This passive acknowledgment scheme scales since it is applied in the single-hop scope of a sender or forwarder. Finally, each node maintains a list of single-hop neighbors (or accesses the neighbor list of the routing layer) in order to detect and inform vehicles entering the target area, i.e., to provide reliability over time.

A simulative study with the ns-2 network simulator evaluates the performance of the TERGF algorithm and compares it to GeoBC. The simulative study evaluates the following metrics:

(i) The topological change rate (TCR) which describes the dynamic of vehicles by accumulating the vehicles entering and leaving the target area over the total number of vehicles remaining in the target area.

(ii) The information distribution ratio which denotes the number of informed vehicles versus the total number of vehicles in the target area over time.

(iii) The total versus the redundant number of packets defines the redundant packet repetition ratio.

The simulations are based on validated, realistic highway movement patterns that are typical for weekday traffic on German highways. The simulations include different traffic densities, such as different number of vehicles per kilometer or different number of lanes per direction. Each simulation run selects an originator in the center of the overall area, such that the target area determined by the originator

is completely located inside the overall simulation area. The geographical target area might be in front or behind the originator, representing different safety applications. The originator continues driving, which represents a *worst case* scenario because the information can leave the target area completely (i.e., all nodes might temporally leave the target area).

The following paragraphs summarize the main simulation results.

The TCR simulation results are independent of the distribution algorithm, as the TCR corresponds to the movement pattern. In small target areas, the TCR is high whereas in large target areas the TCR is low. These results relate to the total number of vehicles which increases with the target area size. However, the TCR results show a frequent fluctuation of vehicles in the target area which proves the demand for a reliable distribution scheme over time that efficiently copes with this high dynamic.

The information distribution simulation results show that TERGF significantly improves the duration during which vehicles in the target area remain informed compared to GeoBC. As examples, in a small target area of 50 m, the average duration to keep all vehicles inside the target area informed is 1.24 s and 2.24 s for GeoBC and TERGF, respectively. For a large target area of 1000 m, the average duration increases from 4 s to 18.46 s for GeoBC and TERGF respectively. While GeoBC distributes a message once, TERGF keeps all connected nodes in a target area informed, as long as there are any vehicles in the target area. Note that the simulation results are based on a *worst case* scenario, such that the originator continuous driving. Assuming that the originator stops, e.g., due to a breakdown, there is always a vehicle in the target area during the lifetime, and the TERGF results would further improve.

The simulative study evaluates the efficiency of both algorithms by accumulating the number of total and redundant packet transmissions. In GeoBC, a huge amount of total and redundant packets occurs upon the initial transmission since each vehicle in the target area rebroadcasts the message once. There are no transmissions over time since GeoBC does not retransmit messages to inform vehicles entering the target area. In contrast, TERGF reduces the total and redundant packet transmissions by 90% to 100% upon the initial information distribution since each vehicle in the target area rebroadcast a packet only if it can reach and inform additional vehicles in the target area. TERGF generates packets over time. These rebroadcasts are required to inform vehicles entering the target area. In this situation, redundant packets may occur when several vehicles detect the same newcomer and the packet generation offset cannot avoid both transmissions. However, these redundant transmissions are negligibly small (i.e., below 1%).

Summarizing, the simulation results show that the TERGF algorithm provides efficient and reliable distribution of information in a geographically scoped target area over time.

An initial implementation of the algorithm has already been done and is integrated in the NoW vehicular test network. This implementation will be extended in the future work, e.g., by implementing the enhancements of TERGF. TERGF will be tested and further evaluated by *real-world* measurements in the framework of the NoW project. A potential further development may integrate receiver centric approaches where the receivers take the forwarding decision, taking into account the effects of different beacon intervals.

Field measurements in the protocol development of the NoW project, such as [96], reveal a significant affect of radio characteristics on network and protocol performance. Even in static and line-of-sight scenarios, the omnipresent environmental motion causes e.g., topology changes due to signal strength fluctuations.

The following Chapter 5 evaluates and quantifies the effects of signal strength fluctuations. First, the study quantifies IEEE 802.11b wireless LAN signal strength fluctuations in field measurements. Second, the chapter develops a simple, but realistic signal strength fluctuation model, as adapted from the measurement results. Third, a simulative study uses the model to quantify the impact of signal strength fluctuations on the ad hoc network performance and compares it to the impact of mobility.

Chapter 5

The Impact of Radio Fluctuations on Ad Hoc Network Performance

5.1 Introduction

The performance of a wireless ad hoc network primarily depends on its ability to adapt to changes in the network topology. Node mobility is considered as the main reason for topological changes in an ad hoc network. However, radio propagation characteristics, such as signal strength fluctuations, significantly affect the performance of an ad hoc network. Fluctuations in signal strength, which lead to variations in the wireless transmission range, are caused by environmental factors, e.g., obstacles or weather conditions. These signal strength fluctuations result in network topology changes even when both transmitter and receiver are stationary. Omnipresent environmental motion introduces temporal fading, as it can be easily observed when performing field trials, e.g., as indicated in [96].

Node mobility is usually taken into account in ad hoc network performance evaluation and protocol design by employing a mobility model or realistic movement patterns. However, most of today's ad hoc network simulation models consider only idealized radio propagation scenarios with constant wireless transmission ranges. As a consequence, such a simulative analysis of ad hoc network performance only reflects the protocol efficiency with respect to node mobility, but neglect the impact of radio propagation characteristics.

We argue that the design of ad hoc network protocols should consider temporal radio signal strength fluctuations in order to avoid performance degradation when transferred to the *real world*, as shown by the following example:

Assume that a routing protocol is employed that makes use of *hello messages* or beacons by which a node can learn about its direct neighbors. With strong temporal transmission range fluctuations, it might happen that a sender beyond the wireless transmission receives a hello message. However, due to fluctuations, the link cannot be used for following data transmissions, and the node's view on the

network topology is inaccurate. In turn, this inaccurate view leads to wrong routing decisions and to performance degradation.

Temporal fading can be studied at various scales: Radio engineers might study fast fading at the μs-level, network-planning engineers might study slow fading or *cell breathing* over hours or days. The ad hoc network protocol design requires studying temporal fading at the scale of the size of packets and their inter-arrival times.

This chapter evaluates and quantifies by means of simulations the impact of signal strength fluctuations on the ad hoc network performance. The study focuses on the metrics *topological change rate* (TCR) and *link stability* which directly influence the performance of the ad hoc network.

As a pre-requisite, the study measures the temporal signal strength fluctuations during an IEEE 802.11b wireless LAN communication. The outdoor measurements are conducted in a static line-of-sight environment, representing a *best case* scenario. Based on the measurement results, the evaluation derives a simple, but realistic signal power fluctuation model. Based on this model, a simulative study evaluates the impact of pure node mobility, pure signal strength fluctuations and the combination of mobility *and* radio fluctuations on the metrics TCR and link stability.

5.2 Related Work

The complexity of wireless radio propagation in *real-world* environments challenges engineering for concrete applications [59], as well as modeling for research purposes [97, 104]. [111, 128]. For example, the MobiHoc tutorial in 2003 [32] expresses the need to incorporate advanced radio propagation models into simulations of higher layer network protocols. However, so far only a small number of simulation studies have addressed time-varying transmission ranges for studying protocol performance in ad hoc networks, as addressed in this chapter.

A notable exception is given in [11] where the authors study from an algorithmic point of view a robust location-aware ad hoc routing protocol that tolerates up to 40% of variation in the wireless transmission range. Still, most studies focus on the unit-disc graph model [29], which assumes static links in a connected graph of the network topology. In [111, 128] it is shown that non-deterministic radio propagation models can have an impact on the performance of packet delivery rate or end-to-end delay for specific protocols.

[11, 111, 128] study the effects of non-deterministic radio propagation models on the performance of specific network protocols, such as position-based routing. In contrast to the related work, this chapter does not focus on specific protocols, but investigates the impact of radio fluctuations on protocol-independent metrics that are closely related to the ad hoc network performance. Both, the related evaluations

and the study of this chapter show that modeling of radio fluctuations is at least as important as modeling of mobility when studying mobile ad hoc networks.

Many network simulation tools provide radio fluctuations models, but unfortunately, this configuration is rarely used in ad hoc network simulations. The widely used network simulator ns-2 [132] includes a shadowing model in order to model fading effects in a log-normal scale. The wireless network simulator WIPPET [73] includes modeling of radio propagation, mobility and multiple protocol layers. The deployed radio propagation model considers fluctuations in space and time due to environmental motion, and it includes activity of environmental objects.

The evaluation in this chapter builds on and extends the work in [107]. The work in the paper analyzes the impact of node mobility on the topological change rate, based on the random waypoint mobility model.

5.3 IEEE 802.11b Radio Fluctuation Measurements

In order to quantify IEEE 802.11b signal strength fluctuations, the wireless LAN driver measures the radio signal strength of correctly received packets in an ongoing data communication, as explained in detail in the remainder of this section.

The goal is to obtain a simple, but reasonable wireless radio fluctuation model, such that the simulation results are not overestimating the impact of radio fluctuations. Therefore, the experimental environment considers a *best case* scenario, i.e., one with a low signal strength variance. Therefore, sender and receiver are statically located in line-of-sight, without obstacles around.

This section describes the scenario, measurement setup and presents the results, which provide the basis for the radio model in the following simulative study.

5.3.1 Measurement Scenario and Setup

The scenario consists of a transmitting and a receiving node that are equipped with IEEE 802.11b wireless LAN cards. The nodes are placed statically in a distance of 250 m in an open field without obstacles closer than 150 m to the direct communication path. Thus, the nodes have a clear line-of-sight during the whole experiment, i.e., the direct communication path strongly dominates over any other path, and the lowest level of fluctuation is expected to occur. Figure 5.1 illustrates the measurement scenario.

The sender generates 18.000 UDP packets with an inter-packet delay of 20 ms. On the successful reception of a packet, the receiver requests the actual signal strength for the current packet out of the Linux interface driver, using the socket interfaces, as provided by the driver. Lost packets are not considered in the measurements (the share of lost packets is negligible, i.e., below 1%).

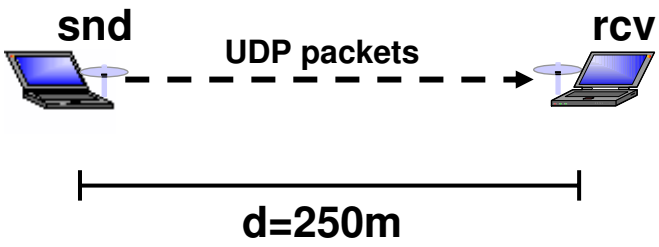

Figure 5.1: IEEE 802.11b measurement scenario.

5.3.2 Measurement Results

Figure 5.2 presents a typical example of the resulting signal strength fluctuations for received packets in the static scenario, as described above.

Figure 5.2(a) shows the signal strength samples in a dBm scale over time. Even in the static line-of-sight scenario, the signal strength fluctuates between -80dBm and -86dBm. Note that this fluctuation in the dBm domain represents almost a doubling of the received signal strength when translated to the decimal scale.

Figure 5.2(b) represents the histogram of the received signal strength. The radio fluctuation model is based on the distribution of the number of occurrences in the measured histogram, as explained in the following section.

5.4 Radio Fluctuation Model

The radio fluctuation model assumes that the distribution of received signal strength represents a normal distribution in the dBm domain, following the basic shape of the measurement result graph. A *least squares fit* algorithm determines the corresponding parameters. Furthermore, the distribution of the signal strength variances between consecutive samples provides the temporal dependency of the measured fluctuations. This temporal dependency is modeled as a first-order Gauss-Markov process, as given in Equation (5.1). The α-value in the formula reflects the temporal dependency. It is determined from the measurements via numerical approximation to be $\alpha = 0.5$.

$$\begin{aligned} r_0 &\approx N(0,1) \\ r_{i+1} &= \alpha * r_i + \sqrt{(1-\alpha^2)} * e \\ &\quad \text{with} \quad e \approx N(0,1) \end{aligned} \tag{5.1}$$

Note that the temporal dependency between subsequent samples is of critical importance since the TCR and the link stability are highly depending on it, as shown in Section 5.5.

In the model, the *sender* transmits with varying power in order to calculate the direct impact of signal strength fluctuations on the radio transmission range.

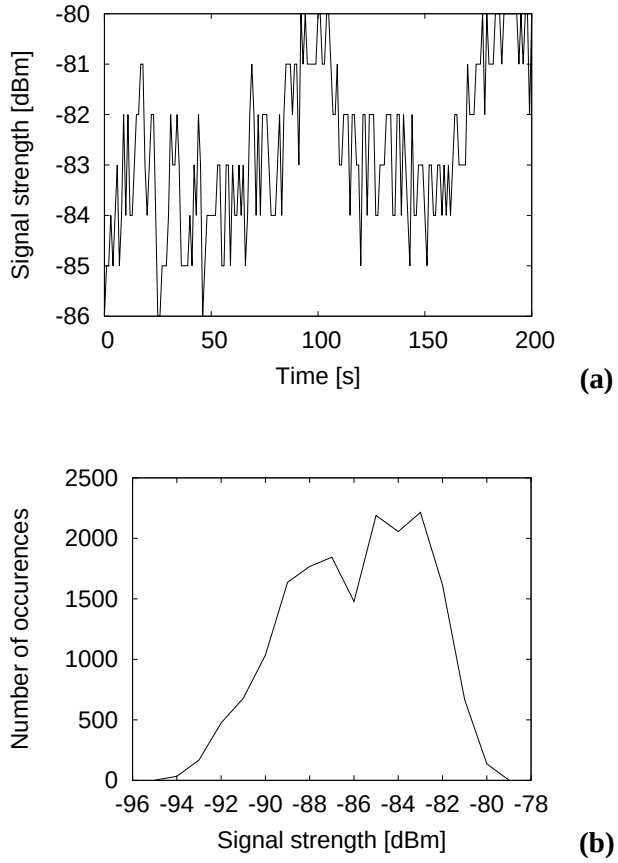

Figure 5.2: IEEE 802.11b WLAN signal strength field measurements
(a) Measured samples over time (b) Histogram of measured signal strengths.

The model employs a coverage area, as in the classical unit-disk graph model [29], which is determined by a circle around the sender. In contrast to the static unit-disc graph approach, the radius of the circle changes over time. The radio cell *breathes*. Note that the approach of fluctuating sending power is valid for this model because with a static sender and only one receiver spatial dependencies are not considered.

In order to obtain the appropriate signal amplitudes at the sender, as well as the conversion between signal strength and radio transmission range, the model uses the free space propagation model, as shown in Equation (5.2). Although there are more accurate propagation models, this simple and optimistic model is appropriate for the evaluated scenario.

$$S_r = S_t + G_t + G_r + 20log\left(\frac{\lambda}{4\pi}\right) - 20log(d) \tag{5.2}$$

S_r and S_t represent the receiving and transmission power in dBW, respectively. The antenna gains G_r and G_t (in dBi) are set to one. The wavelength λ = 0.125 m derives from the deployed 802.11b frequency and the distance is set to 250 m. As a result, the transmission power of the model follows a normal distribution with $N[1, 4.085^2]$.

Finally, the model restricts the maximum signal strength variance to three times the standard deviation, which assures that 99.73% of the sampling values are located within the minimum/maximum borders, according to the 3σ-rule [56]. This minimum and maximum approximately correspond to the borders obtained from the field trial results.

Summarizing, Figure 5.3 illustrates the signal strength fluctuation model. Figure 5.3(a) presents samples over time whereas Figure 5.3(b) shows the histogram of the fluctuations in the model.

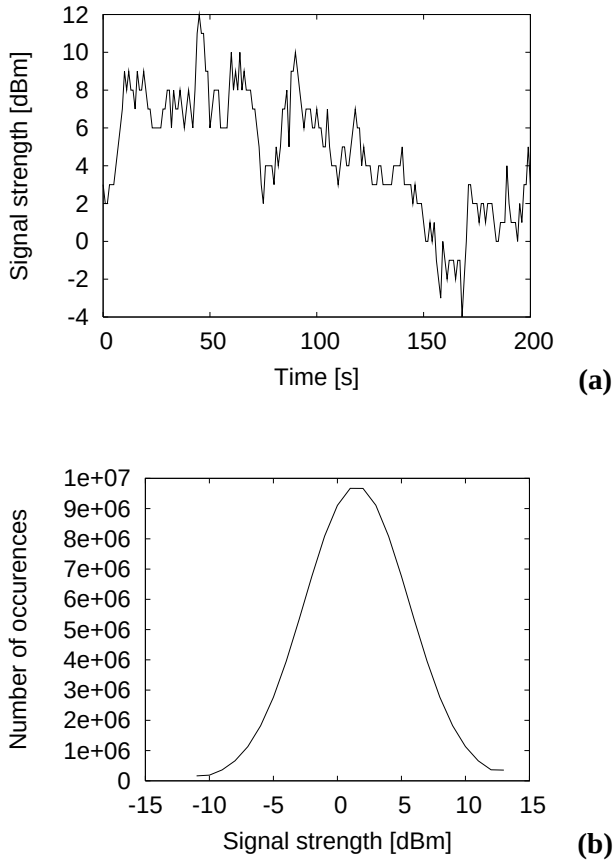

Figure 5.3: Signal strength fluctuation model
(a) Modeled samples over time (b) Histogram of the modeled signal strengths.

Signal power, reception thresholds and radio propagation laws determine the maximum distance at which a data packet can be properly received. The minimal and maximal values of the signal strength deviation define a minimum radius r' and a maximum radius r'', as depicted in Figure 5.4. This defines the border of a wireless *cell* in the model. According to the fluctuation model above, as derived from the measurements, the signal strength variations cause transmission radius changes between r' = 50 m and r'' = 790 m. The average transmission radius is r_{av} = 200 m. Note that the range distribution is not symmetric due to the logarithmic scale in the signal strength fluctuations.

The model assumes the following characteristics in the spatial domain:

- Devices within a distance d to the node less than the minimum transmission range r' (i.e., $d \leq r'$) are able to communicate perfectly (i.e., without packet losses).

- Devices within a distance d to the node greater than the maximum transmission range r'' (i.e., $d > r''$) are not able to communicate at all.

- Between the minimum r' and the maximum r'' the radius varies over time according to the fluctuation model. Therefore, devices within a distance d between r' and r'' (i.e., $r' < d \leq r''$) may or may not be able to communicate depending on the temporal fluctuation $r(t)$.

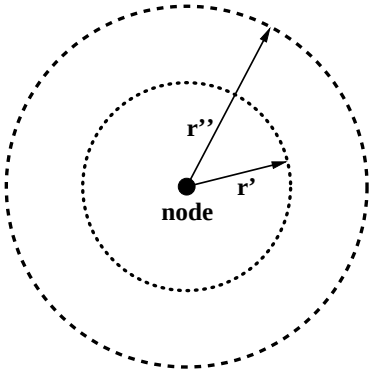

Figure 5.4: Transmission range variations.

The model assumes equal transmission ranges for transmitter and receiver at a time, according to the assumption that the radio channel between two nodes is reciprocal. Unidirectional communication capabilities in the model and in the line-of-sight scenario result in bi-directional channels when we neglect interferences. When a node is in communication range of its peer, the reverse path exists as well due to the same sending power of the interfaces and equal environmental conditions.

5.5 Simulative Evaluation

This section defines the performance metrics TCR and link stability as protocol independent measures of the ad hoc network performance. Furthermore, it describes the simulation scenario and presents the simulation results.

The simulation evaluates the effects of pure mobility, pure signal strength fluctuations and compares the impact of both parameters in a scenario that considers the combination of mobility *and* radio fluctuations.

5.5.1 Performance Metrics

This study focuses on the metrics: *Topological Change Rate (TCR)* and *link stability* because both metrics correlate with the ad hoc network performance. Both metrics are used to characterize the network dynamics in [17, 32].

Since both metrics are built on the notion of a *link*, the concept of a *link* must be translated to the wireless networking environment: While in wired networking a link usually is either *on* or *off*, a *wireless link* can differ in its quality to transmit a signal. Therefore, the wireless environment requires a definition when a link is to be considered *on* or *off*.

In the simulations, the communication capability is checked every 20 ms. A link state change at every check (i.e., with every received or lost packet) results in an unreasonable *ping-pong* behavior. In order to avoid these frequent link state changes, the simulation applies the following link state management scheme:

When the link state is *connected*, it requires 25 consecutive checks with the result *not connected* before changing the state to *off*. This represents 500 ms without being able to communicate. In the other case when changing from *not connected* to *connected*, it requires three successfully received packets within the 500 ms period in order to avoid an inaccurate network view due to a single received packet. This hysteresis avoids fluctuations due to single or small number of lost packets, as experienced in the measurements in static scenarios.

The TCR is the *natural* mobility and radio fluctuation metric because it reflects per definition the number of link changes per time unit, as observed by a single node.

Link stability represents the average link duration of a single link, again, as observed by a single node.

5.5.2 Simulation Scenario and Environment

The following evaluation uses our own discrete event simulation tool, implemented in C. The overall simulation area is circular with the radius R = 1000 m and considers two nodes. Each node is equipped with a wireless radio interface. A fixed

observer (Obs) is located in the center, while a mobile node (MN) is either positioned statically at various distances to the center or moving through the area based on the Random Waypoint (RWP) movement pattern [14].

With the RWP model, a node selects the next *waypoint*, i.e., the next destination, by sampling from a uniform distribution over the system area. The node moves with constant speed to this *waypoint*. It might pause there before iterating the scheme. The basic principle of this scheme is illustrated in Figure 5.5.

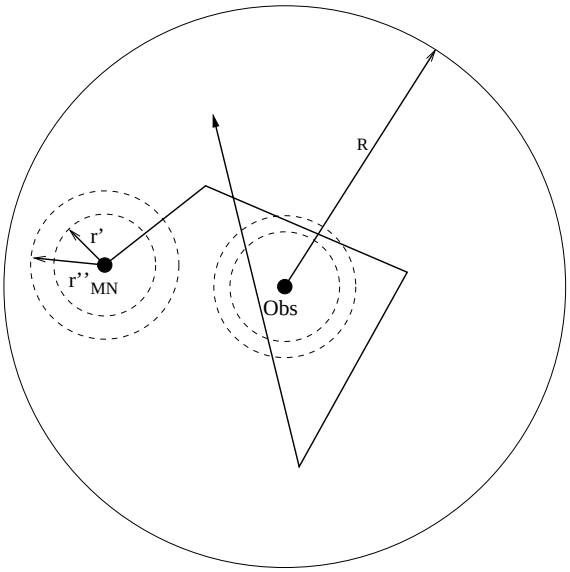

Figure 5.5: RWP movement pattern.

Recent studies of RWP have shown some unexpected behavior, e.g., with respect to the spatial node distribution [13, 15] and the average velocity [140]. However, these results should not be interpreted in a way that RWP as a mobility model is invalid, but show the importance of proper use. In the simulative evaluation of this chapter, the mobile node travels with a constant speed and does not pause at the waypoints.

As mentioned before, the simulation checks every 20 ms whether connectivity between the MN and Obs exists. Therefore, the signal strength is converted to the maximum distance at which a packet can still be properly received via the free space propagation model. This maximum distance is compared to the actual distance between MN and Obs.

In order to evaluate the effects of pure mobility, the signal strength and, thus, the radio transmission range, is kept static while the MN moves according to the RWP model shown in Figure 5.5. For the evaluation of pure signal strength fluctuations, MN and Obs are statically placed at certain distances, as shown in Fig-

ure 5.6. The combination of both parameters requires the combination of signal strength fluctuation and mobility, as shown in Figure 5.5.

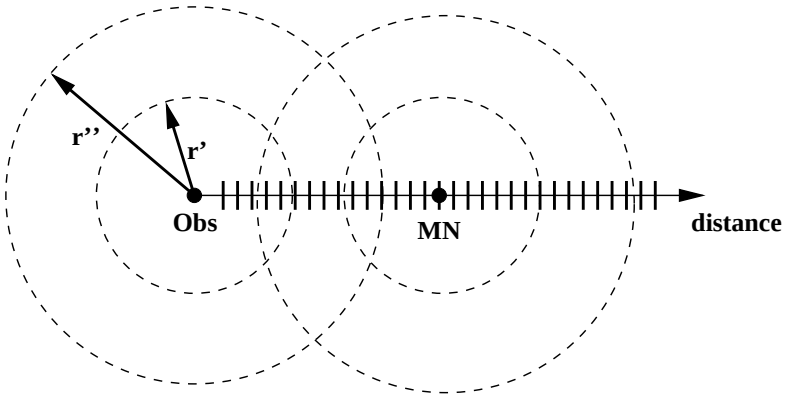

Figure 5.6: Static simulation scenario.

5.5.3 Simulation Results

This section presents the simulation results, which evaluate and compare the impact of pure mobility, pure signal strength and the combination of both on the network performance metrics TCR and link stability.

5.5.3.1 Topological Change Rate Analysis

The TCR reflects the number of link changes per time unit. In order to compare the effects of mobility and radio fluctuations, this section structures the results according to the methodology above.

Impact of Pure Mobility on the TCR. The simulation results depicted in Figure 5.7 illustrate the impact of pure node mobility on the performance metric TCR. Within this scenario, the observer is located statically in the center, while the MN moves in the simulation area according to the RWP mobility pattern. The wireless transmission radius r (i.e., the sending signal power) remains constant for each simulation run. The ratio r/R shows the effects of mobility for different constant radio coverage areas in relation to the overall simulation area. The velocity in this simulation is 1 m/s reflecting a scenario of pedestrian motion. We have chosen the scenario with slow moving nodes with respect to the previous measurements. Such a scenario allows a more accurate modeling since side effects due to high mobility can be neglected. However, the model is extensionable to highly dynamic environments, such as vehicular networks. In this scenario, the maximum TCR of 0.001/min is reached for the ratio r/R = 0.65.

The TCR depends linearly on the speed, i.e., different velocities result in modified amplitude. The results, originally presented in [107], show the impact of

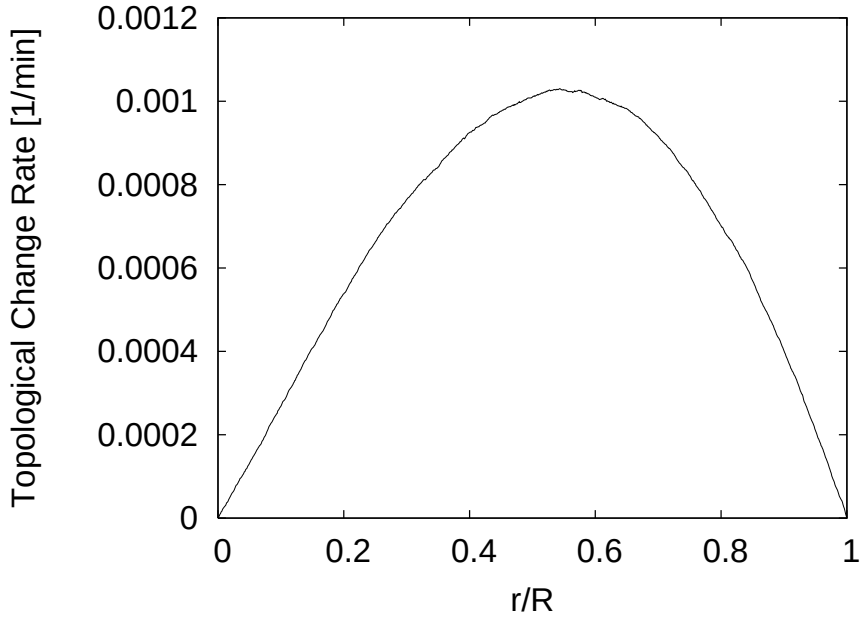

Figure 5.7: Impact of pure mobility on the TCR for constant transmission radius and velocity 1 m/s.

node mobility on the TCR and provide the basis of comparison for the following evaluation of mobility and radio fluctuations.

Impact of Pure Signal Strength Fluctuation on the TCR. In order to evaluate the impact of pure signal strength fluctuations, the MN is positioned statically for each simulation run in one meter steps along the x-axis for the distances $0 < d \leq 800$. There is a significant impact on the TCR for certain distances close to the edge of the communication range where the ability of communication is highly affected by fluctuations. Figure 5.8 illustrates the TCR over distance for pure signal strength fluctuations.

The consideration of the deployed link state management scheme causes the maximum TCR to occur at a distance of 360 m when a node receives 10.05% of the packets. At this point, the small percentage of received packets along the whole simulation is distributed in the way that causes the link management scheme to detect the highest number of link changes; the maximum TCR reaches 0.78/min.

Impact of Mobility and Signal Strength Fluctuations on the TCR. The previous sections have evaluated node mobility and signal strength fluctuations separately to show the impact of both parameters on the link individually. A direct comparison of the individual impact was not possible due to the different settings in the scenarios. In order to provide a comparative evaluation of the TCR metric,

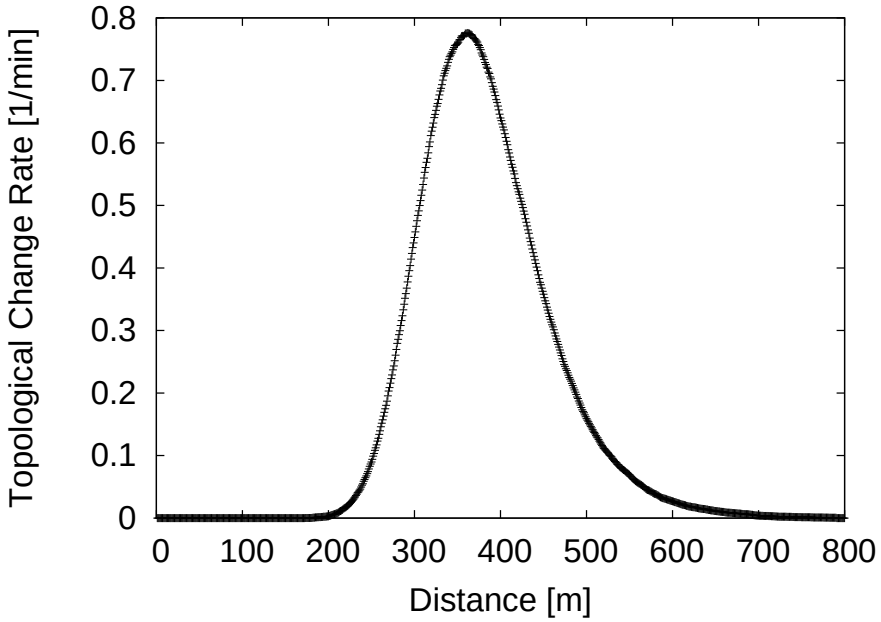

Figure 5.8: Impact of pure signal strength fluctuations on TCR for static nodes at various distances.

the following scenario combines node mobility and signal power fluctuations. The observer node is positioned statically in the center while the mobile node moves according to the random waypoint mobility model.

Figure 5.9 presents the TCR over varying signal power deviation for combined mobility and signal strength fluctuations for two different velocities. A standard deviation of σ = 0 reflects the TCR for constant transmission power. With increasing standard transmission power deviation, the TCR increments exponentially due to the growing influence of signal strength fluctuation. The speed determines the gradient for small standard deviation but the effect of velocity is insignificant when the deviation increases, as indicated by v = 1 m/s and v = 10 m/s. Again the scenarios of slow moving nodes are chosen on purpose in order to avoid side effects of high dynamic networks. Still the model can be extended to such scenarios.

Figure 5.10 examines the same scenario from a different perspective. It shows the TCR over velocity for the standard deviation σ=4.085, as obtained from the field trial. The figure includes the results of constant signal strength for comparison, which reflects the effects of pure mobility with constant transmission radius. The latter graph visualizes the linear dependency of TCR and velocity for constant signal power. The difference between the graphs indicates the TCR increase caused by wireless signal strength fluctuations. The slope is almost zero, which illustrates the low impact of the velocity when considering fluctuations. Note that when fluctuations and movement are considered, the crucial factor is the overall time in which the MN is located in the *fluctuation area* (i.e., between r' and r'').

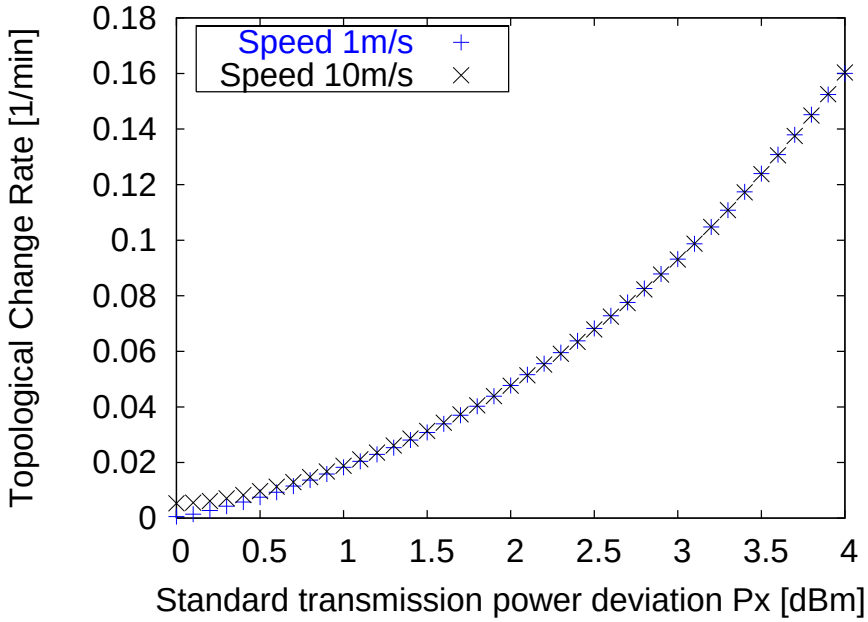

Figure 5.9: Impact of mobility and fluctuation on the TCR over standard transmission power deviation.

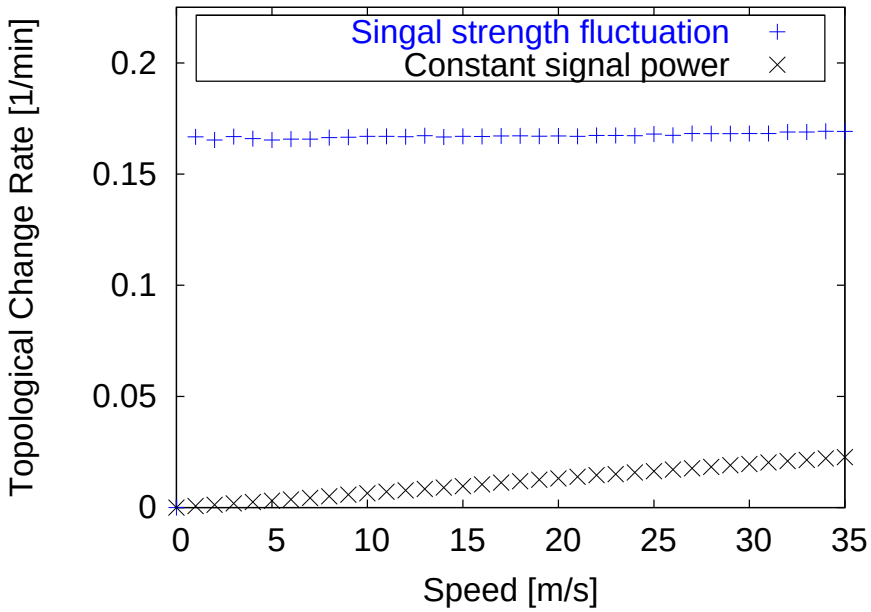

Figure 5.10: Impact of mobility and fluctuations on the TCR over speed for a standard transmission power deviation 4.085 dBm.

This time is approximately the same for slow and fast nodes since fast nodes visit the area for a shorter time but more frequently, while slow moving nodes visit the areas for a longer time but less frequently.

5.5.3.2 Link Stability Analysis

This section presents the impact of pure node mobility, pure transmission radius fluctuations and the combination of both on the network performance metric link stability in the scenario as described above.

The following evaluation focuses on the average link stability since the accumulated overall link duration is equal for slow and fast nodes. As mentioned at the end of the previous section, slow nodes enter less often the link state *connected*, but the duration per state is longer compared to fast moving nodes, which show the opposite behavior. This results in equal overall link duration.

Impact of Pure Mobility on Link Stability. Figure 5.11 illustrates the effects of pure mobility on the link stability metric for a constant transmission range (i.e., constant signal strength) over the ratio r/R, i.e., the wireless transmission radius in relation to the overall area. As in the TCR evaluation of pure mobility effects, the observer is located statically in the center while the MN moves around in the simulation area according to the RWP mobility model. Both nodes send with constant transmission power, i.e., constant transmission radius for each simulation run.

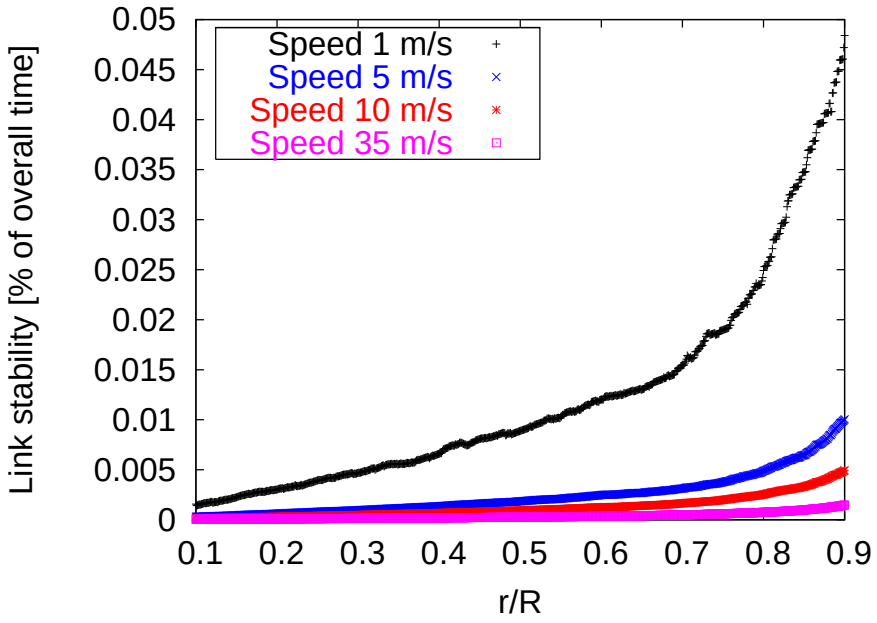

Figure 5.11: Impact of pure mobility on link stability for constant transmission radius.

The increase of the r/R ratio in Figure 5.11 improves the link stability because a larger radio coverage area enhances the probability to be in the link state *connected*, which affects the link stability positively. Naturally, the slope is more significant for slow moving nodes, as shown in the curve for the speed 1 m/s because slow nodes remain within communication distance for a longer time interval.

Impact of Pure Signal Strength Fluctuation on Link Stability. The evaluation of the impact of pure signal strength fluctuation on link stability requires, once more, static positioning of the communication peers in different distances.

Figure 5.12 shows the performance of link stability over increasing distances between the two static nodes. While for a constant transmission radius, the nodes would be able to communicate up to the distance equal to the transmission radius and there would be no communication beyond this point, the link stability for varying signal strength decreases as a function of the distance.

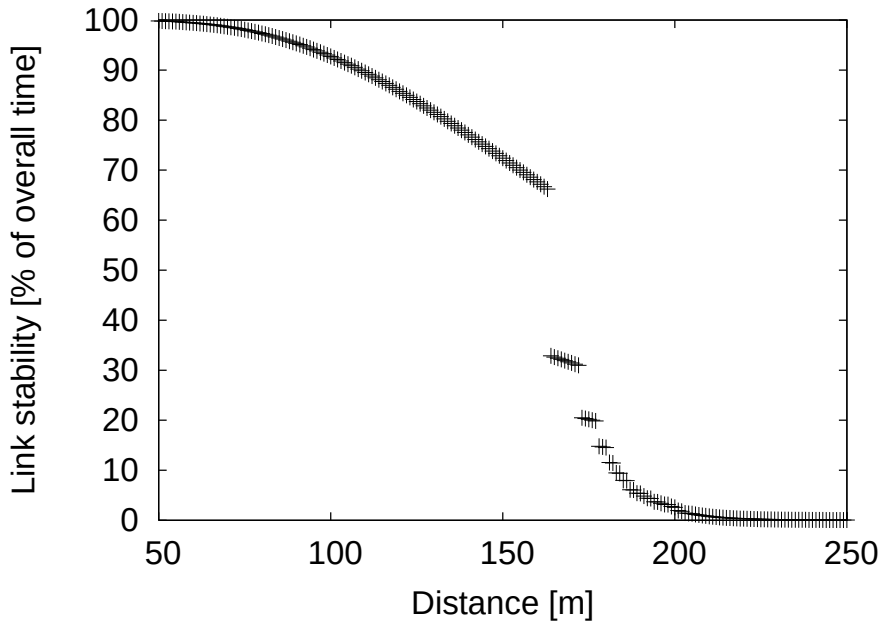

Figure 5.12: Impact of pure fluctuations on link stability for pure signal strength fluctuations with static nodes.

Up to a distance of approximately 160 m, there is only one long *connected* state with exponentially decreasing link duration. When reaching the distance d = 163 m, several *connected* link states occur during a simulation run, which results in the drop of the average link stability below 50%, as visible in the Figure. Finally, the graph finds back to the exponential decrease beyond d = 170 m when a difference in the number of connected states is not so significant.

Impact of Mobility and Signal Strength Fluctuations on Link Stability. The following comparative evaluation of node mobility and signal strength fluctuations and their impact on the link stability metric assumes, once more, the static observer in the center, while the MN moves within the simulation area according to the RWP movement pattern.

The combined impact of mobility and fluctuations on the link stability metric over different standard deviations is shown in Figure 5.13 for two different velocities. In contrast to the TCR evaluation, the link stability decreases exponentially when increasing standard transmission power deviation, representing the impact of fluctuations. Again, the low impact of mobility can be observed when increasing the deviation.

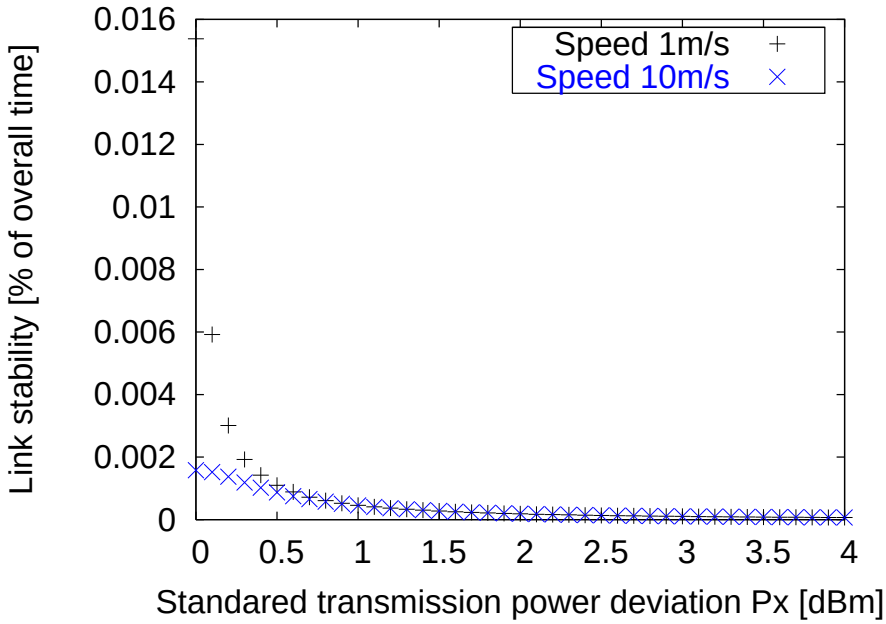

Figure 5.13: Link stability over standard signal power deviation for 1 m/s velocity.

Finally, Figure 5.14 illustrates link stability over speed for standard deviation σ=4.085. Similar to the TCR, link stability is not affected by velocity. Again, the reason for the independence is that the accumulated period of time in which the MN moves through the *fluctuation area* is equal for slow and fast nodes, as explained before.

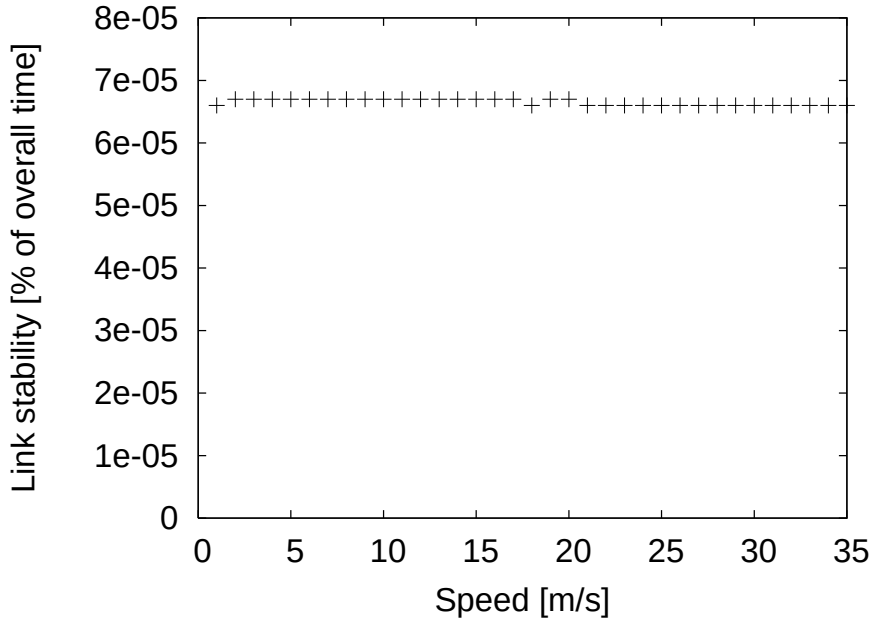

Figure 5.14: Link stability over velocity for standard signal strength deviation 4.085 dBm.

5.6 Summary and Conclusions

Radio fluctuations affect the performance of ad hoc networks, in addition to node mobility. Our evaluation quantifies the amount of radio fluctuations in a static line-of-sight scenario via IEEE 802.11b field trial measurements. The measurement results prove that even in this *best case* scenario, significant fluctuations occur due to omnipresent environmental effects.

Based on the measurement results, the analysis derives a simple, but realistic model for radio fluctuations. A normal distribution in the dBm domain models the distribution of radio fluctuations and a Gauss-Markov process models the temporal dependency of consecutive samples. The model adapts to *real world* conditions through *least square* parameter fitting, according to the results of the field measurements.

Since the concept of a *link* does not readily take over for the wireless domain, the simulation model uses a heuristic to check whether a link is available or not. This heuristic avoids *ping-pong* effects in the link state by requiring a pre-defined number of received or lost packets to change the link state in the simulation.

As the main contribution, the simulations evaluate the ad hoc network performance metrics *topological change rate* (TCR) and *link stability*. The TCR reflects the number of link changes over time, and link stability represents the average link

duration. Both metrics directly impact the ad hoc network performance.

The simulations evaluate and compare the effects of pure node mobility and the effects of pure signal strength fluctuations of a wireless channel, as well as the combination of node mobility *and* signal strength fluctuations.

The simulation results indicate that signal strength fluctuations have a significant impact on the ad hoc network performance in addition to node mobility. Therefore, temporal signal power fluctuations due to radio propagation laws have to be taken into account in ad hoc simulation models and in the design of robust ad hoc network protocols. As an example, a robust network protocol may not use a potential forwarding node upon reception of a single beacon when this forwarder is located close to the radio transmission border.

As future work, further radio effects will be evaluated in different environments. As with mobility, there is a quest for a model that is simple enough to be computed but realistic enough to trust the results. Furthermore, the effects of signal strength fluctuations and radio characteristics on ad hoc routing and transport layer will be analyzed in order to provide guidelines for future *cross-layer* protocol design.

Chapter 6

Conclusions and Outlook

The thesis has studied different aspects of transport and network layer performance in single-hop and multi-hop wireless networks, such as:

- Handover latency measurements, TCP and UDP performance in presence of Mobile IPv6 and Fast Mobile IPv6 handovers.
- Design and evaluation of a point-to-point transport protocol that is tailored to highway scenarios in vehicular ad hoc networks.
- Design and evaluation of an algorithm for the reliable and efficient distribution of safety information in a geographically limited target area over time (i.e., during the lifetime of a safety event).
- Evaluation of signal strength fluctuations on the impact of ad hoc network performance.

Chapter 2 evaluated handover performance in an integrated, mobility-enabled IPv6 environment, including Quality of Service (QoS), Authentication, Authorization, Accounting and Charging (AAAC).

First, the study provided results, as obtained from measurements in an integrated test network, on standard and fast handovers in Mobile IPv6. The fast handover implementation provides interruptionless Ethernet-WLAN handover and in average 0.23 ms handover latency for WLAN-WLAN intra-technology handover, assuming a pre-established security association between the access routers and ideal QoS and AAAC attendants. Furthermore, the fast handover latency is independent of round trip time between MN and HA or the router advertisement interval. In comparison, both network characteristics significantly affect the handover latency of standard Mobile IPv6 handover, which in contrast to fast handover is between 500 ms and 1500 ms.

Based on the handover latency results, the second part of the study measures UDP and TCP performance in presence of fast handovers in detail. The results show that neither UDP nor TCP performance is affected by fast handovers. Due

to the extremely short handover latency, a single packet is lost in the worst case. This single packet loss does not affect UDP performance, e.g., in real-time UDP communications, the user cannot notice a single lost packet. The TCP version NewReno, as used in the measurements, quickly repeats the lost packet via the Fast Retransmission scheme without modifying its throughput or the congestion window (i.e., TCP does not invoke a slow-start). Consequently, the fast handover does not affect the TCP performance.

The transport performance evaluation results of this chapter prove that a fast handover scheme is indispensable for the successful deployment of next generation All-IP networks. Furthermore, the results show that the traditional transport protocols UDP and TCP are able to provide uninterrupted service without performance degradation in the presence of fast handovers.

The future work continues the architecture design to integrate network controlled handover initiation and enhance the security level. Network providers have a strong interest to increase their control, e.g., in order to support load balancing or for business models that offer different levels of QoS at different prices.

In the following chapters we have focused on transport issues for wireless multi-hop communication in vehicular ad hoc networks (VANETs).

Chapter 3 designed and evaluated a vehicular transport protocol (VTP) that is optimized for the unique characteristic of this highly dynamic environment.

Prior to the transport protocol design, the chapter evaluated through simulations the network path characteristics of VANETs in a highway scenario, such as expected communication and disruption duration for specific source-destination distances, packet loss characteristics, reordering, round trip time (RTT) and RTT jitter. Based on these results, the chapter then designed a vehicular transport protocol (VTP). The key features of VTP are:

- The VTP sender uses a rate-based transmission scheme.
- VTP decouples error and congestion control.
- VTP uses explicit signaling of available bandwidth from intermediate nodes for congestion control.
- VTP provides reliability via retransmissions of lost packets. The VTP receiver reports received and missing packets via selective acknowledgments (SACKs). The receiver transmits SACKs in periodic intervals, depending on the source-destination distance.
- The VTP sender uses statistical knowledge to predict the expected communication behavior of a connection, e.g., in the absence of acknowledgments.

A simulative study evaluated the performance of VTP, such as throughput and fairness, in static and mobile multi-hop wireless scenarios and compared it to TCP

as the reference transport protocol. We found that VTP outperforms TCP, particularly with respect to throughput and fairness. The results show that VTP maintains a steady transmission rate and quickly adapts to disruption or congestion. The transmission rate is adjusted according to feedback (or absence of feedback) of intermediate nodes and includes statistical knowledge (e.g., expected communication duration for certain distances) in the transmission rate calculation. The selective acknowledgment scheme allow efficient retransmission, and thus, provide reliability.

As a main result, VTP provides reliable end-to-end connections and outperforms the varying throughput and unfairness of TCP by maintaining a steady throughput above the average throughput of TCP.

The future work will adjust and evaluate VTP in city scenarios. This includes a simulative evaluation and, furthermore, the NoW project intends to implement VTP and perform *real-world* measurements. Finally, interoperability with TCP (e.g., translation proxies at road side access points) is required in order to provide connectivity to the Internet or any arbitrary fixed, local infrastructure networks.

Beyond the point-to-point transport requirements, as evaluated above, VANETs demand for reliable and efficient distribution of information to multiple receivers, e.g., for traffic safety applications.

Chapter 4 designed and evaluated an efficient, time-extended reliable flooding algorithm for geographical target areas (TERGF). This algorithm provides efficient and reliable distribution of information in a geographical area over time. In particular, it informs vehicles that enter the target area after the initial distribution of the message. It combines the following mechanisms:

- Reliable distribution of information, including content-related interpretation and aggregation of data (i.e., in contrast to traditional packet-based reliability).
- Efficient broadcasting of information based on the combination of GeoCast (i.e., geographical addressing) and self-pruning (i.e., explicit addressing of single-hop neighbors in the target area in the packet header).
- The single-hop broadcasting is extended by acknowledgment-based reliability, using a passive acknowledgment scheme to detect single-hop losses.
- Redistribution of safety information in case of single-hop losses.
- Redistribution of safety information when vehicles enter the target area, for the lifetime of the safety event.

The simulative evaluation of TERGF shows a significant increase of the *information distribution ratio* whereas the number of *redundant packet repetitions*

decreases, compared to the standard GeoCast packet distribution scheme. The information is kept longer in the target area since the information is rebroadcast when further vehicles become reachable, i.e. reliability over time. In TERGF, the information is only lost when all informed vehicles leave the target area before new vehicles enter. Efficiency is increased because information is only rebroadcast when additional (i.e., new) neighbors are reachable.

An initial implementation of TERGF is already available on the nodes of the NoW vehicular test network. The future work will enhance and extended this implementation, e.g., the extensions still have to be included. An evaluation via *real-world* measurements in the NoW test network is also part of the future work. An interesting approach might also consider to think about performing the forwarding decision on the receivers and compare such an enhancement to standard TERGF.

Since field measurements in protocol developments for VANETs have shown a significant impact of radio characteristics on the ad hoc network performance, the final study of this thesis evaluated the impact of signal strength fluctuations on the ad hoc network performance.

Chapter 5 evaluated the impact of radio fluctuations on the topological change rate (TCR) and the link stability and compared it to the impact of mobility. Both metrics directly influence the performance of ad hoc networks. Our field measurements quantify signal strength fluctuations in a static, line-of-sight environment, which represent a *best case* scenario. Based on the measurement results, the study derives a simple, but realistic model for signal strength fluctuations. The simulative study that uses this signal strength fluctuation model concludes that the impact of fluctuations is as high as the impact of mobility and in some specific scenarios even higher.

Consequently, the chapter concludes that in future works temporal radio fluctuations have to be taken into account in ad hoc simulation models and in the design of robust ad hoc network protocols, in addition to node mobility.

Though the different chapters of this thesis study different aspects of transport and performance issues of ad hoc networks, an inherent part of the future work is the convergence of the different research disciplines. The radio fluctuation results impact the further design and development of all network and transport protocols. The integration of vehicular transport protocols and TCP in an integrated future All-IP mobility environment will facilitate access to fixed networks out of vehicles, e.g., for traffic related information or passenger entertainment. *Putting the pieces together* in the future work will significantly enhance performance in future mobile networks in a variety of scenarios, even in high dynamic environments such as vehicular networks.

Bibliography

[1] G.-S. Ahn, A.T. Campbell, A Veres, and L.-H. Sun. SWAN: Service Differentiation in Stateless Wireless Ad Hoc Networks. In *IEEE Conference on Computer Communications (INFOCOM)*, New York, NY, USA, June 2002.

[2] T. Henderson amd R. Katz. Transport Protocols for Internet compatible Satellite Networks. *IEEE JSAC*, 17(2), Februrary 1999.

[3] V. Anantharaman, S.-J. Park, K. Sundaresan, and R. Sivakumar. TCP Performance over Mobile Ad Hoc Networks: A Quantitative Study. *Wireless Communications and Mobile Computing*, 4(2), March 2004.

[4] L. Andrew, S. Hanly, and R. Mukhar. CLAMP: Differentiated Capacity Allocation in Access Networks. In *Proc. IEEE International Performance Comparison and Communication Conference*, Phoenix, AZ, USA, April 2003.

[5] A. Bakre and B. Dadinath. I-TCP: Indirect TCP for mobile hosts. In *Proc. 15th International Conference on Distributed Computing Systems (ICDCS)*, Vancouver, BC, Canada, May 1995.

[6] A.V. Bakre and B.R. Badrinath. Implementation and Performance Evaluation of Indirect TCP. *IEEE Transactions on Computers*, 46(3), March 1997.

[7] B. Bakshi, P. Krishna, N.H. Vaidya, and D.K. Pradhan. Improving performance of TCP over wireless networks. In *Proc. 17th International Conference on Distributed Computing Systems*, 1997.

[8] H. Balakrishnan, S. Seshan, E. Amir, and R.H. Katz. Improving TCP/IP Performance over Wireless Networks. In *Proc. 1st ACM Conf. on Mobile Computing and Networking*, Berkely, CA, November 1995.

[9] H. Balakrishnan, S. Seshan, E. Amir, and R.H. Katz. Improving TCP/IP Performance over Wireless Networks. In *Proc. International Conference on Mobile Computing and Networking (MobiCom)*, Berkeley, CA, USA, November 1995.

[10] H. Balakrishnan, S. Seshan, and R.H. Katz. Improving Reliabile Transport and Handoff Performance in Cellular Wireless Networks. *ACM Wireless Networks*, 1(4), December 1995.

[11] L. Barriere, P. Fraigniaud, and L. Narayanan. Robust Position-Based Routing in Wireless Ad Hoc Networks with Unstable Transmission Ranges. In *Proc. 5th International Workshop on Discrete Algorithms and Methods for Mobile Computing and Communications*, pages 19–27, Rome, Italy, July 2001.

[12] D. Bertsekas and R. Gallager. *Data Networks (2nd Ed.)*. Prentice Hall, 1992. Chapter 6.

[13] C. Bettstetter. Mobility in wireless networks: Categorization, smooth movement and border effects. *ACM MC2R*, 5(3):55–67, 2001.

[14] C. Bettstetter, H. Hartenstein, and X. Perez-Costa. Stochastic properties of the random waypoint mobility model. *ACM/Kluwer WINET*, 10(5), September 2004.

[15] C. Bettstetter, G. Resta, and P. Santi. The node distribution of the random waypoint mobility model for wireless ad hoc networks. *IEEE Transactions Mobile Computing*, 2(3):257–269, 2003.

[16] J. Boleng, T. Camp, and V. Tolety. Mesh-Based GeoCast Routing Protocols in an Ad Hoc Network. In *Proc. 15th International Parallel and Distributed Processing Symposium (IPDPS)*, San Francisco, CA, USA, April 2001.

[17] J. Boleng, W. Navadi, and T. Camp. Metrics to enable adaptive protocols for mobile ad hoc networks. In *Proc. International Conference on Wireless Networks (ICWN)*, pages 293–298, Las Vegas, Nevada, USA, June 2002.

[18] R. Boppana and S. Konduru. An adaptive Distance Vector Routing Algorithm for Mobile, Ad Hoc Networks. In *Proc. 20th Annual Joint Conference of the IEEE Computer and Communications Societies (INFOCOM)*, Anchorage, Alaska, USA, April 2001.

[19] P. Bose, P. Morin, I. Stojmenovic, and J. Urrutia. Routing with guaranteed delivery in ad hoc wireless networks. *ACM Wireless Networks*, November 2001.

[20] K. Brown and S. Singh. M-TCP: TCP for Mobile Cellular Networks. *ACM SIGCOMM Computer Communications Review*, 27(5), October 1997.

[21] C. Janneteau. IST project WINE GLASS (IST-1999-10699), Deliverable D09: Specification of Wireless Internet Testbed and Research Results. http://wineglass.tilab.com, September 2001.

[22] R. Caceres and L. Iftode. Improving the Performance of Reliable Transport Protocols in Mobile Computing Environments. *IEEE Journal of Selected Areas in Communications (JSAC)*, 13(5):850–857, June 1995.

[23] T. Camp and Y. Liu. An Adaptive Mesh-Based Protocol for GeoCast Routing. *Journal of Parallel and Distributed Computing: Special Issue on Routing and Mobile Wireless Ad Hoc Networks*, 62(2), 2003.

[24] S. Capkun, M. Hamdi, and J. Hubaux. GPS-Free Positioning in Mobile Ad Hoc Networks. In *Proc. 34th Annual Hawaii International Conference on System Sciences (HICSS)*, Maui, Hawaii, USA, January 2001.

[25] K. Chandran, S. Raghunathan, S. Venkatesan, and R. Prakash. A Feedback-Based Scheme for Improving TCP Performance in Ad Hoc Wireless Networks. In *Proc. 18th International Conference on Distributed Computing Systems (ICDCS)*, Amsterdam, Netherlands, May 1998.

[26] C.-Y. Chang, C.-T. Chang, and S.-T. Tu. Obstacle-Free Geocasting Protocols for Single/Multi-Destination Short Message Services in Ad Hoc Networks. *Wireless Networks*, 9(2), 2003.

[27] K. Chen and K. Nahrstedt. EXACT: An Explicit Rate-based Flow Control Framework in MANET. *Technical Report UIUCDCS-R-2002-2286/UILU-ENG-2002-1730*, July 2002.

[28] K. Chen, K. Nahrstedt, and N. Vaidya. The Utility of Explicit Rate-Based Flow Control in Mobile Ad Hoc Networks. *Proc. IEEE Wireless Communications and Networking Conference (WCNC)*, March 2004.

[29] B.N. Clark, C.J. Colburn, and D.S. Johnson. Unit disk graphs, Discrete Mathematics, Volume 86. *Elsevier Science Publishers*, pages 165–177, 1991.

[30] T. Clausen and P. Jacquet. Optimized Link State Routing Protocol (OLSR). In *Internet Engineering Task Force (IETF), RFC 3626*, October 2003.

[31] A. Colvin. CSMA with collision avoidance. *Computer Communications*, 6(5), March 1983.

[32] C. Constantinou. Radiowave Channel Modelling for Radio Networks. In *Tutorial 4th ACM International Symposium on Mobile Ad Hoc Networking and Computing (MobiHoc)*, Annapolis, MD, USA, June 2003.

[33] D. Johnson and C. Perkins and J. Arkko. Mobility Support in IPv6, May 2003. draft-ietf-mobileip-ipv6-22.txt, IETF Internet Draft, work in Progress.

[34] S.R. Das, R. Castaneda, and J. Yan. Simulation Based Performance Evaluation of Mobile Ad Hoc Network Routing Protocols. *ACM/Baltzer Mobile Networks and Applications (MONET) Journal*, July 2001.

[35] S. Deering. Multicast Routing in a Datagram Internetwork. *PhD Thesis, Stanford University, Palo Alto, CA, USA*, December 1991.

[36] R. Durst, G. Miller, and E. Travis. TCP Extensions for Space Communications. In *Proc. International Conference on Mobile Computing and Networking (MobiCom)*, Rye, NY, USA, November 1996.

[37] T. Dyer and R. Boppana. A Comparison of TCP Performance over Three Routing Protocols for Mobile Ad Hoc Networks. In *Proc. 4th ACM International Symposium on Mobile Ad Hoc Networking and Computing (MobiHoc)*, Long Beach, CA, USA, October 2001.

[38] K. Egevang and P. Francis. The IP Network Address Translator (NAT), May 1994. Internet Engineering Task Force (IETF), RFC 1631.

[39] H. Elaarag. Improving TCP performance over mobile networks. *ACM Computing Surveys*, 34(3):357–374, 2002.

[40] S. ElRakabawy, A. Klemm, and C. Lindemann. TCP with Adaptive Pacing for Multihop Wireless Networks. In *Proc. ACM International Symposium on Mobile Ad Hoc Networking and Computing (MobiHoc)*, Urbana-Champaign, IL, USA, May 2005.

[41] A. Fladenmuller and R. De-Silva. The effect of mobile IP handoffs on the performance of TCP. *Mobile Networks and Applications*, 4(2):131–135, 1999.

[42] S. Floyd, V. Jacobson, C. Liu, S. McCanne, and L. Zhang. A Reliable Multicast Framework for Light-Weight Sessions and Application Level Framing. *IEEE/ACM Transactions on Networking*, 5(6), December 1997.

[43] Z. Fu, X. Meng, and S. Lu. How bad TCP Can Perform In Mobile Ad Hoc Networks. In *Proc. Symposium on Computers and Communications (ISCC)*, Taormina, Italy, July 2002.

[44] H. Füßler, M. Mauve, H. Hartenstein, M. Käseman, and D. Vollmer. A Comparison of Routing Strategies for Vehicular Ad Hoc Networks. *Technical Report TR-02-003, Department of Computer Science, University of Mannheim, Germany*, July 2002.

[45] H. Füßler, M. Torrent-Moreno, R. Krüger, M. Transier, H. Hartenstein, and W. Effelsberg. Studying Vehicle Movements on Highways and their Impact on Ad-Hoc Connectivity. *Technical Report TR-2005-003, Department of Computer Science, University of Mannheim, Germany*, June 2005.

[46] H. Füßler, J. Widmer, M. Käseman, M. Mauve, and H. Hartenstein. Contention-Based Forwarding for Mobile Ad-Hoc Networks. *Elsevier Ad Hoc Networks*, 1(4):351–369, September 2003.

[47] G. Dommety and A. Yegin and C. Perkins and G. Tsirtsis and K. Malki and M. Khalil. Fast Handovers for Mobile IPv6, November 2001. draft-ietf-mobileip-fast-mipv6-03.txt, IETF Internet Draft, work in Progress.

[48] M. Gerla, K. Tang, and R. Bagrodia. TCP Performance in Wireless Multihop Networks. *Proc. IEEE Workshop on Mobile Computer Systems and Applications (WMCSA)*, February 1999.

[49] S. Giordano and M. Hamdi. Mobility Management: the Virtual Home Region. *Technical Report TR DSC 99-037, Ecole Politechnique Federal de Lausanne (EPFL), Swiss*, March 2000.

[50] H. Soliman and C. Catelluccia and K. El-Malki and L. Bellier. Hierarchical Mobile IPv6 mobility management (HMIPv6), December 2004. draft-ietf-mipshop-hmipv6-04.txt, IETF Internet Draft, work in Progress.

[51] Z. Haas and B. Liang. Ad hoc mobility management with uniform quorum systems. *IEEE/ACM Transactions on Networking*, 7(2):228–240, April 1999.

[52] Z. Haas and M. Pearlman. The performance of query control schemes for the zone routing protocol (ZRP). *ACM/IEEE Transactions on Networking*, 9(4):427–438, August 2001.

[53] A. Hanbali, E. Altman, and P. Nain. A Survey of TCP over Ad Hoc Networks. *IEEE Communications Surveys and Tutorials*, 7(3), 2005.

[54] Helsinki Univertiy of Technolog (HUT). MIPL - Mobile IPv6 for Linux. www.mipl.mediapoli.com.

[55] J. Hightower and G. Barriello. Location Systems for Ubiquitous Computing. *IEEE Computer*, August 2001.

[56] P.G. Hoel. Introduction To Mathematical Statistics, 5th Ed. *Published by John Wiley and Sons*, 1984. page 418, table II.

[57] B. Hofmann-Wellenhof, H. Lichtenegger, and J. Collins. *GPS Theory and Practice*. Springer-Verlag/Wien, New York, 1997.

[58] G. Holland and N. Vaidya. Analysis of TCP over Mobile Ad Hoc Networks. *ACM Wireless Networks*, 8(2), March 2002.

[59] M. Hope and N. Linge. Determining the Propagation Range of IEEE 802.11 Radio LAN's for Outdoor Applications. In *Proc. Local Computer Networks (LCN)*, Lowell, Massachusetts, USA, October 1999.

[60] T.C. Hou and V.O.K. Li. Transmission Range Control in Multihop Packet Radio Networks. *IEEE Transactions on Communications*, 34(1), January 1986.

[61] Internet Engineering Task Force, IETF. www.ietf.org.

[62] T. Imielinski and J. Navas. GPS-based Addressing and Routing, November 1996. Internet Engineering Task Force (IETF), RFC 2009.

[63] Institute of Electrical and Electronics Engineers (IEEE) Inc. IEEE802.11b Standard. http://www.ieee802.org/11/.

[64] Institute of Electrical and Electronics Engineers (IEEE) Inc. IEEE802.3 Ethernet Standard. http://www.ethermanage.com/ethernet/standard.html.

[65] J. Postel. User Datagram Protocol, August 1980. Internet Engineering Task Force (IETF), RFC 768.

[66] J. Postel. Transmission Control Protocol (TCP), September 1981. Internet Engineering Task Force (IETF), RFC 793.

[67] D. Johnson and D. Maltz. Mobile Computing, chapter 5, Dynamic Source Routing (DSR). *Kluwer Academic Publishers*, pages 153–181, 1996.

[68] K. Ramakrishnan and S. Floyd and D. Black. The Addition of Explicit Congestion Notification (ECN) to IP, September 2001. Internet Engineering Task Force (IETF), RFC 3168.

[69] E. Kaplan. *Understanding GPS*. Artech House, 1996.

[70] B.N. Karp. Geographic Routing for Wireless Networks. *PhD Thesis*, 2000.

[71] B.N. Karp and H.T. Kung. Greedy Perimeter Stateless Routing for Wireless Networks. In *ACM/IEEE International Conference on Mobile Computing and Networking (MobiCom)*, Boston, MA, USA, August 2000.

[72] M. Kazantzidis, M. Gerla, and S.-J. Lee. Permissible Throughput Network Feedback for Adaptive Multimedia in AODV MANETs. In *International Conference of Communications (ICC)*, Helsinki, Finland, June 2001.

[73] O.E. Kelly, J. Lai, N.B. Mandayam, A.T. Ogielski, J. Panchal, and R.D. Yates. Scalable parallel simulations of wireless networks with WIPPET: Modelling of radio propagation, mobility and protocols. *Kluwer Mobile Networks and Application*, 5(3):199–208, September 2000.

[74] D. Kim, C. Toh, and Y. Choi. TCP-BuS: Improving TCP Performance in Wireless Ad Hoc Networks. *Communications and Networking*, 3(2), June 2001.

[75] Y. Ko and N.H. Vaidya. GeoTORA: A Protocol for Geocasting in Mobile Ad Hoc Networks. In *Proc. 8th Annual International Conference on Network Protocols (ICNP)*, Osaka, Japan, November 2000.

[76] Y.-B. Ko and N.H. Vaidya. Geocasting in Mobile Ad Hoc Networks: Location-Based Multicast Algorithms. In *Proc. 2nd Workshop Mobile Computer Systems and Applications (WMCSA'99)*, pages 101–110, New Orleans, USA, February 1999.

[77] Y.-B. Ko and N.H. Vaidya. Flooding-Based Geocasting Protocols for Mobile Ad Hoc Networks. *Mobile Networks and Applications*, 7(6):471–480, 2002.

[78] C.E. Koksal and H. Balakrishnan. An Analysis of Short-term Fairness in Wireless Media Access Protocols (poster). In *Proc. ACM Measurement and Modeling of Computer Systems (SIGMETRICS)*, pages 118–119, Santa Clara, CA, USA, 2000.

[79] S. Kopparty, S.V. Krishnamurthy, M. Faloutsos, and S.K. Tripathi. Split TCP for Mobile Ad Hoc Networks. In *Proc. IEEE Global Telecommunications Conference (GLOBECOM)*, Taipei, Taiwan, November 2002.

[80] R. Krüger, H. Füßler, M. Torrent-Moreno, M. Transier, H. Hartenstein, and W. Effelsberg. Statistical Analysis of the FleetNet Highway Movement Patterns. *Technical Report TR-2005-004, Department of Computer Science, University of Mannheim, Germany*, July 2005.

[81] A. Leiggener. Evaluation of Path Characteristics in Vehicular Ad Hoc Networks. *Master's Thesis, Institute EURECOM, Sophia Antipolis, France*, September 2005.

[82] A. Leiggener, R. Schmitz, A. Festag, L. Eggert, and W. Effelsberg. Analysis of Path Characteristics and Transport Protocol Design in Vehicular Ad Hoc Networks. In *In (electronical) Proc. Vehicular Technology Conference (VTC), Spring 06*, Melbourne, Australia, 2006.

[83] J. Li, J. Jannotti, D.S.J. De Couto, D.R. Karger, and R. Morris. A Scalable Location Service for Geographic Ad Hoc Routing. In *Proc. 6th Annual International Conference on Mobile Computing and Networking (MobiCom)*, Boston, MA, USA, August 2000.

[84] W.-H. Liao, Y.-C. Tseng, K.-L. Lo, and J.-P. Sheu. GeoGRID: A Geocasting Protocol for Mobile Ad Hoc Networks Based on GRID. *Journal of Internet Technology*, 1(2):23–32, 2000.

[85] H. Lim and C. Kim. Multicast Tree Construction and Flooding in Wireless Ad Hoc Networks. In *Proc. ACM International Workshop on Modeling, Analysis and Simulation of Wireless and Mobile Systems (MSWIM)*, pages 61–68, Boston, Massachusetts, USA, August 2000.

[86] H. Lim, K. Xu, and M. Gerla. TCP Performance over Multipath Routing in Mobile, Ad Hoc Networks. In *Proc. Innovative Confinement Concepts (ICC)*, Seattle, Washington, USA, May 2003.

[87] J. Liu and S. Singh. ATCP: TCP for Mobile Ad Hoc Networks. *IEEE JSAC*, 19(7), July 2001.

[88] L. Lovasz. *On the Ratio of Optimal Integral and Fractional Covers - Discrete Mathematics*. Springer-Verlag, 2003. ISBN 0387955852.

[89] S. Lu and M. Gerla. Split Multipath Routing with Maximally Disjoint Paths in Ad Hoc Networks. In *Proc. Innovative Confinement Concepts (ICC)*, Helsinki, Finland, June 2001.

[90] M. Mathis and J. Mahdavi and S. Floyd and A. Romanow. TCP Selective Acknowledgement Options, April 1996. Internet Engineering Task Force (IETF), RFC 2018.

[91] C. Maihöfer. A Survey on Geocast Routing Protocols. *IEEE Communications Surveys and Tutorials*, 6(2), 2004.

[92] B.S. Manoj and C. Siva Ram Murthy. A High Performance Wireless Local Loop Architecture Utilizing Directional Multi-Hop Relaying. *Technical Report Department of Computer Science and Engineering, University of Madras, India*, June 2002.

[93] M. Mauve, J. Widmer, and H. Hartenstein. A Survey on Position-Based Routing in Mobile Ad-Hoc Networks. *IEEE Network*, 5(6), 2001.

[94] IST project Moby Dick - Mobility and Differentiated Services in a Future IP Network. http://www.ist-mobydick.org. Project Number: IST-2000-25394.

[95] J. Monks, P. Sinha, and V. Bharghavan. Limitations of TCP-ELFN for Ad Hoc Networks. In *Proc. Mobile and Multimedia Communications*, Tokyo, Japan, October 2000.

[96] M. Möske, H. Füßler, H. Hartenstein, and W. Franz. Performance Measurements of a Vehicular Ad Hoc Network. In *Proc. IEEE Vehicular Technology Conference (VTC'04 Spring)*, Milan, Italy, May 2004.

[97] P. Muehlethaler. An efficient simulation model for wireless LAN's applied to the IEEE 802.11 standard. *INRIA Rapport de Recherche*, No. 4182, April 2001.

[98] C.S.R. Murthy and B.S. Manoj. *Ad Hoc Wireless Networks - Architectures and Protocols*. Prentice Hall Communications Engineering and Emerging Technologies Series, 2005. ISBN 0-13-147023.

[99] S. Ni, Y. Tseng, Y. Chen, and J. Sheu. The Broadcast Storm Problem in a Mobile Ad Hoc Network. In *Proc. 5th Annual International Conference on Mobile Computing and Networking (MobiCom)*, Seattle, Washington, USA, August 1999.

[100] NISTNet, Network Emulation Package. http://snad.ncsl.nist.gov/itg/nistnet, September 2003.

[101] NoW - Network on Wheels. http://www.network-on-wheels.de/.

[102] V.D. Park and M.S. Corson. A highly adaptive distributed routing algorithm for mobile wireless networks. In *Proc. Sixteenth Annual Joint Conference of the IEEE Computer and Communications Societies (INFCOM)*, Kobe, Japan, April 1997. 9 pages.

[103] S. Paul, K. Sabani, J. Lin, and S. Bhattacharyya. Reliable Multicast Transport Protocol (RMTP). *IEEE Journal on Selected Areas in Communications*, 3(15), 1997.

[104] J.S. Pegon and M.W. Subbarao. Simulation Framework for a Mobile Ad-Hoc Network. In *Proc. OPNETWORK*, Washington DC., USA, Sept. 1999.

[105] W. Peng and X. Lu. On the Reduction of Broadcast Redundancy in Mobile Ad Hoc Networks. In *Proc. 1st ACM International Symposium on Mobile Ad Hoc Networking and Computing (MobiHoc)*, pages 129 – 130, Boston, Massachusetts, USA, August 2000.

[106] W. Peng and X. Lu. AHBP: An Efficient Broadcast Protocol for Mobile Ad Hoc Networks. *Journal of Scienc and Technology*, 2002. Bejing, China.

[107] X. Perez-Costa, C. Bettstetter, and H. Hartenstein. Mobicom Poster: Towards a mobility metric for reproducible and comparable results in ad hoc network research. *ACM MC2R*, 7(4), October 2003.

[108] X. Perez-Costa, M. Torrent-Moreno, and H. Hartenstein. A Simulation Study on the Performance of Hierarchical Mobile IPv6. *ACM SIGMOBILE Mobile Computing and Communications Review*, 7(4):5–19, 2003.

[109] C. Perkins. Ad-Hoc On-Demand Distance Vector Routing (AODV). In *Proc. 2nd IEEE Workshop on Mobile Computing Systems and Applications (WMCSA)*, New Orleans, Louisiana, USA, February 1999.

[110] C. Perkins and P. Bhagwat. Highly Dynamic Destination-Sequenced Distance-Vector Routing (DSDV). *Computer Communication Review*, October 1994.

[111] L.F. Perrone, Y. Yuan, and D.M. Nicol. Modeling and simulation best practices for wireless ad hoc networks. In *Proc. 35th Winter Simulation Conference (WSC)*, New Orleans, Louisiana, USA, 2003.

[112] A. Qayyum, L. Viennot, and A. Louiti. Multipoint Relaying: An Efficient Technique for Flooding in Mobile Wireless Networks. *Technical Report 3898, INRIA - Rapport de Recherche*, July 2000.

[113] R. Koodli. Fast Handovers for Mobile IPv6, July 2005. Internet Engineering Task Force (IETF), RFC 4068.

[114] Research Project Partly Funded by the German Ministry for Research and Education - BMBF. INVENT Research Initiative. http://www.invent-online.de.

[115] S. Deering and R. Hinden. Internet Protocol, Version 6 (IPv6) Specification, December 1998. Internet Engineering Task Force (IETF), RFC 2460.

[116] S. Floyd and T. Henderson. The NewReno Modification to TCP's Fast Recovery Algorithm, April 1999. Internet Engineering Task Force (IETF), RFC 2582.

[117] S. Ostermann. Tcptrace - Open Source Software Tool for Analysis of Tcpdump Files. http://www.tcptrace.org.

[118] S. Parker and C. Schmechel. Some Testing Tools for TCP Implementors, August 1998. Internet Engineering Task Force (IETF), RFC 2398.

[119] S. Recker. IST project WINE GLASS (IST-1999-10699), Deliverable D15: Evaluation of Wireless Internet Testbed and Research Results. http://wineglass.tilab.com, March 2002.

[120] S. Thomson and T. Narten. IPv6 Stateless Address Autoconfiguration, December 1998. Internet Engineering Task Force (IETF), RFC 2462.

[121] W.R. Stevens. *TCP/IP Illustrated, Volume 1*. Addison-Wesley Professional Computing Series, 1996. ISBN 0-201-63346-9.

[122] I. Stojmenovic, A.P. Ruhil, and D.K. Lobiyal. Voronoi diagram and convex hull based geocasting and routing in wireless networks. In *Proc. 8th IEEE International Symposium on Computers and Communications (ISCC)*, Washington, DC, USA, October 2003.

[123] L. Stojmenovic. Home Agent Based Location Update and Sestination Search Schemes in Ad Hoc Wireless Networks. *Technical Report TR 99-10, Computer Science SITE, University of Ottawa, Canada*, September 1999.

[124] J. Sucec and I Marsic. An Efficient Distributed Network-Wide Broadcast Algorithm for Mobile Ad Hoc Networks. *Technical Report 248, Rutgers University, New Jersey, USA*, September 2000.

[125] K. Sundaresan, A. Vaidyanathan, H.-Y. Hsieh, and R. Sivakumar. ATP: A Reliable Transport Protocol for Ad-hoc Networks. In *Proc. 4th ACM International Symposium on Mobile Ad Hoc Networking and Computing (MobiHoc)*, Annapolis, Maryland, USA, June 2003.

[126] T. Narten and E. Nordmark and W. Simpson. Neighbor Discovery for IP version 6 (IPv6), December 1998. Internet Engineering Task Force (IETF), RFC 2461.

[127] H. Takagi and L. Kleinrock. Optimal Transmission Ranges for Randomly Distributed Packet Radio Terminals. *IEEE Transactions on Communications*, 32(3), March 1984.

[128] M. Takai, J. Martin, and R. Bagrodia. Effects of wireless physical layer modeling in mobile ad hoc networks. In *Proc. 4th ACM International Symposium on Mobile Ad Hoc Networking and Computing (MobiHoc)*, pages 87–94, Long Beach, CA, USA, Oct. 2001.

[129] Tcpdump - Open Source Network Analyser (based on libcap library). http://www.tcpdump.org.

[130] The Global UMTS Alliance. UMTS TDCDMA. http://www.umtstdd.org/.

[131] M. Torrent-Moreno, X. Perez-Costa, and S. Sallent-Ribes. A Performance Study of Fast Handovers for Mobile IPv6. In *Proc. IEEE Local Computer Networks (LCN)*, page 89, Bonn, Germany, October 2003.

[132] UCB/LBNL/VINT. Discrete event network simulator ns, version 2. http://www.isi.edu/nsnam/ns/.

[133] VideoLanClient (VLC). A free cross-platform media player. http://www.videolan.org/.

[134] W. Stevens. TCP Slow Start, Congestion Avoidance, Fast Retransmit, and Fast Recovery Algorithms, January 1997. Internet Engineering Task Force (IETF), RFC 2001.

[135] F. Wang, C. Toh, and Y. Choi. Improving TCP Performance over Mobile Ad Hoc Networks with Out-of-order Detection and Response. In *Proc. 4th ACM International Symposium on Mobile Ad Hoc Networking and Computing (MobiHoc)*, Lausanne, Switzerland, June 2002.

[136] B. Williams and T. Camp. Comparison of Broadcasting Techniques for Mobile Ad Hoc Networks. In *Proc. 3rd ACM International Symposium on Mobile Ad Hoc Networking and Computing (MobiHoc)*, pages 194–205, Lausanne, Switzerland, June 2002.

[137] D. Wisely and E. Mitjana. Paving the Road to Systems Beyond 3G. *Journal of Communication and Networks - Special issue*, December 2002.

[138] S. Xu and T. Saadawi. Performance Evaluation of TCP Algorithms in Multi-Hop Wireless Packet Networks. *Journal of Wireless Communication and Mobile Computing*, 2(1), 2002.

[139] P. Yao, E. Krohne, and T. Camp. Performance Comparison of Geocast Routing Protocols for a MANET. In *Proc. 13th IEEE International Conference on Computer Communications and Networks (IC3N)*, Chicago, IL, USA, October 2004.

[140] J. Yoon, M. Liu, and B. Noble. Random waypoint considered harmful. In *Proc. 22nd IEEE Annual Joint Conference of the IEEE Computer and Communications Societies (Infocom)*, pages 1312–1321, San Franciso, CA, USA, March 2003.