Aalborg Universitet

Timed Broadcast via Off-The-Shelf WLAN Distributed Coordination Function for Safety-Critical Systems

Malinowsky, B. ; Grønbæk, Lars Jesper; Schwefel, Hans-Peter; Ceccarelli, A.; Bondavalli, A.; Nett, E.

Publication date:
2012

Document Version
Publisher's PDF, also known as Version of record

[Link to publication from Aalborg University](#)

Timed Broadcast via Off-The-Shelf WLAN Distributed Coordination Function for Safety-Critical Systems

B. Malinowsky[†], J. Grønbæk[†], H.P. Schwefel[†‡]
‡*Dept. of Electronic Systems*
Aalborg University, Denmark
†*FTW, Wien, Austria*
{malinowsky,groenbaek,schwefel}@ftw.at

A. Ceccarelli, A. Bondavalli
University of Firenze
Viale Morgagni 65
I-50134, Firenze, Italy
{andrea.ceccarelli, bondavalli}@unifi.it

E. Nett
University of Magdeburg
Universitätsplatz 2
D-39106, Magdeburg, DE
nett@ivs.cs.uni-magdeburg.de

Abstract—**Low cost wireless solutions for safety-critical applications are attractive to leverage safety-critical operation in new application areas. This work assesses the feasibility of providing synchronous and time bounded communication to standard IEEE 802.11 devices with low effort modifications. An existing protocol for time bounded communication in wireless systems is adapted to a generic safety-critical application with low bandwidth requirements, but strict bounds on time behavior. Experimental and simulation studies are conducted in which the protocol is implemented on top of IEEE 802.11e Distributed Coordination Function (DCF). The experimental results for packet loss ratio, communication delays, and broadcast completion are used to calibrate a stochastic simulation model that allows to extrapolate the expected long-term performance of the protocol. Both the experimental results and the simulation extrapolation show that necessary availability requirements can be met with 802.11e prioritization in the investigated cross-traffic and interference scenarios.**

Keywords: Safety-critical, Synchronous protocols, TDMA, 802.11e DCF, Experimental & Simulation

I. Introduction

Developing cost-efficient wireless communication solutions for high dependability and safety applications is attractive to 1) enable safety-critical automation for scenarios which currently rely on manual processes such as assisted vehicle driving and 2) to replace existing (wired) automation systems which lack flexibility in installation and are potentially too costly to deploy widely. An example case of the latter is presented in this paper considering an automated alerting system for railway workers utilizing wirelessly connected personal devices to issue alerts about approaching trains.

Wireless communication link reliability depends on the actual distances between the nodes, which varies in mobile settings, on the presence of obstacles, and on possible interference from other transmitting devices or noise sources. When transmitting on a shared communication medium, with a possibly unbounded number of participating nodes and various procedures for channel access, timing faults have to be accepted. Executing critical distributed services and achieving the required reliability, safety and timeliness of the communication through portable and wireless devices in harsh environments requires the design of resilient services and networks [1], [2].

A driver for wireless communication solutions has in the past decade been the IEEE 802.11 standard. Today it represents the *de-facto* technology to enable ubiquitous wireless access solutions for private and enterprise users. However, the standard is now also defining the basis for new automation solutions such as wireless communication infrastructure for intelligent transportation systems 802.11p [3] and is intended for data collection, monitoring and control in industrial applications.

In general, low cost solutions based on IEEE 802.11 are manufactured primarily with focus on performance. Thus, integral features needed for high dependability and safety are lacking such as synchronous and time bounded communication. How to provide such properties in wireless communication is a well studied topic, see [4]. However, many existing solutions require low level modifications to hardware and software which increases cost.

In this work, a practical approach is taken to evaluate a timed reliable communication solution operating on top of the standard 802.11 protocol stack with off-the-shelf hardware. The approach entails:

- Use the Distributed Coordination Function (DCF), a carrier sense multiple access with collision avoidance technique for sharing the channel among multiple nodes, in addition to the widely available 802.11e for medium access priority control.
- Use an existing reliable group communication protocol, and adapt its assumptions and properties (validity, agreement, integrity, ordering) to our use-case scenario.
- Evaluate communication reliability and performance in an experimental study with and without contending nodes.
- Develop a stochastic simulation model of the protocol for a feasibility analysis to extrapolate the results, and assess expected long-term protocol behavior in a full scenario setup.

IEEE
computer
society

The protocol is intended to operate in the context of the ALARP (A railway automatic track warning system based on distributed personal mobile terminals [5]) project, where real-time and safety-critical communication is required to transmit information to wireless devices carried by railway trackside workers. The wireless communication layer needs to deliver both life-critical messages (in particular, notification of trains approaching the worksite), as well as non-life critical messages. The protocol presented in this paper is devised starting from the ALARP requirements and a previous protocol, the Real-time Group Communication Protocol (RGCP, [6] [7]).

The paper is organized as follows: Section II introduces the use-case of the alert system for railway workers, and points out the key requirements driving the communication protocol design. That section also provides the background on the IEEE 802.11 Distributed Coordination Function (DCF) protocol and channel access prioritization options using IEEE 802.11e. Section III introduces the synchronous communication reference protocol, and describes the adaptations of this protocol to the ALARP scenario. The experimental and model based evaluations of the modified protocol realization are introduced in Section IV to assess its measured and expected performance. This part also constitutes the focus of this paper. Finally, Section V concludes the results and provides an outlook to future safety-oriented analysis of the protocol.

Related Work: Several approaches exist for applying Quality of Service (QoS) for real-time communication in IEEE 802.11 networks. Reference [8] evaluates different schemes, also covering techniques for low latency real-time communication like Blackburst [9], which requires jamming the wireless medium. Also, new software-based time slots coordination mechanism were proposed [10], with the main target being Voice-over-IP communication [11].

II. SAFETY-CRITICAL COMMUNICATION SCENARIO

This section introduces the safety-critical communication scenario, which guides our further studies of experimental test setups and developing our simulation models. The scenario covers the safety of railway workers at work sites near railway tracks, which can be easily generalized to similar setups. Hence, the motivation of studying such use-cases is not limited to railway-specific scenarios.

Safety of workers in railway scenario is a serious concern, since vehicles are constrained to tracks and drivers have much less margins to react in case of emergencies; therefore, trackside workers are exposed to injuries and fatalities [12], [13]. The ALARP [5] project proposes to design, develop and validate an Automatic Track Warning System to improve the safety of trackside workers. This system will be able to inform the trackside workers about approaching trains on the track. Additionally, it will keep track of the status

and position of the workers, to identify those at risk, not responding, or to suggest escape routes [14], [15].

The ALARP architecture is based on the following components realized on top of Commercial Off-The-Shelf (COTS) hardware: i) the trackside Train Presence Alert Device (TPAD), able to sense an approaching train on the interested track without interfering with the signaling system, using long-range multi-spectral cameras and eavesdropping the train-network communication, ii) a set of distributed, low-cost, wearable, context-aware, robust, trustable and highly reliable, wireless Mobile Terminals (MTs) to inform the workers about possible approaching trains and/or other events that could put at risk their safety, and iii) infrastructure for wireless communication. The MT will be able to generate alarms, and to communicate and interact through wireless connections with other MTs and the trackside train presence alert devices [14], [15].

High reliability, timeliness and safety, despite the possible harsh conditions are mandatory requirements of the ALARP communication, as alarms raised by the TPAD are safety-critical messages that need to be timely delivered to all the workers. For safety reasons, violations of timing bounds will have to be detected and will (after remediation actions) finally lead to sending subset of workers to safe zones; this would impact the working procedures leading to loss of productivity.

A. Communication Architecture

The overall communication architecture in ALARP follows a centralized communication setup. This enables better predictability of the communication timing and simplified realization of synchronous communication channels at the worksite, see Figure 1. This setup is primarily based on a fixed coordinator located at the worksite, with all MTs communicating via the coordinator [14]. The deployed timed reliable wireless communication protocol is using the coordinator to implement its centralized communication algorithm and maintain allocation of necessary communication resources. Communication links between TPAD and the worksite might be enhanced by helper infrastructure in form of additional relay nodes (or repeaters) at the transmission path. At the worksite, TPAD information is disseminated to MTs via the coordinator.

The overall communication solution maintains the communication layer (of the MTs and TPADs) in a known communication state. It also adheres to safety requirements, and timely decides on missing nodes as well as nodes deviating from the expected operation behavior, enforcing communication timeouts. The process of sending and delivering a message is bounded by a maximum time delay requirement; typical values are in the order of several seconds. In ALARP, this time requirement is set to 10 seconds.

The communication protocol adopted in ALARP shall allow the communication layer to reliably distribute messages.
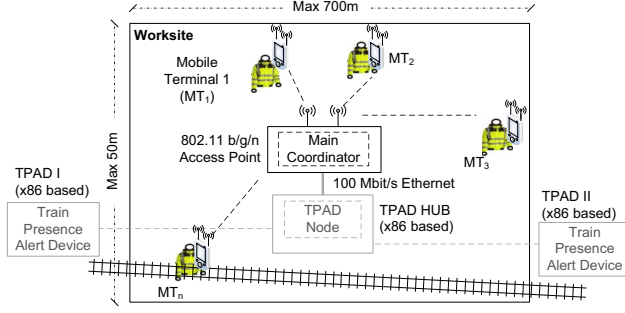
145

Figure 1. Communication architecture of the alert system for railway workers: workers are equipped with a Mobile Terminal (MT).

It offers broadcast (distribution to all nodes), multicast (one-to-many communication), and unicast (message exchange between two dedicated nodes) communication primitives. The reliability requirement is defined to match three different message criticalities, coming with a specific resilience degree for messages and probability of message delivery. For this, the communication layer offers the three criticality levels, or *message classes*: *high* for safety-critical messages (life-critical messages as notification to the workers of a train approaching the worksite, health problems of a worker, etc.), *medium* for messages which are not safety-critical but if not timely received may affect system availability (e.g., the message which notifies that a dangerous situation ended and work can restart), and *low* for messages with no special requirements (best-effort messages, e.g., the message an MT transmits to notify all MTs that it is being switched off by its user).

In the following sections, we focus on the timed reliable broadcast approach that is operated between the coordinator and all MTs over COTS Wireless Local Area Network (WLAN) technology.

B. Background on 802.11 Coordination Functions

The IEEE 802.11 standard specifies two channel access coordination functions, the Distributed Coordination Function (DCF), and the Point Coordination Function (PCF), summarized in the following.

Distributed Coordination Function: WLAN 802.11 uses a carrier-sense multiple access scheme with collision avoidance on the Media Access Control (MAC) layer, a sublayer of the Open Systems Interconnection (OSI) model data-link layer (Layer 2). We here summarize the basic mechanism, see [16] for details. 802.11 nodes that have a Layer 2 frame ready to transmit, first sense the channel: if it is idle for a certain time interval (referred to as Inter-Frame Space (IFS)), the fragment is transmitted immediately. Otherwise, the nodes enter into a random backoff procedure, where the backoff counter is chosen uniformly between two Contention Window (CW) bounds (CW_{min}

and CW_{max}), referring to the lower and upper bound, respectively. Values for the CW are given in slot time units, with the duration of a single slot specified by 802.11. The upper bound is increased upon unsuccessful transmissions in order to provide adaptivity to congestion levels, i.e., number of contending nodes. Time periods during which the medium is used by other nodes lead a node to temporarily pause its count-down to access the medium. When the backoff counter reaches 0, the frame is transmitted and the sending node listens for an Acknowledgment (Ack) from the receiver. The Acks are prioritized on the channel, as the medium access procedures are performed with lower IFS for those. If the sending node does not correctly receive an Ack in a certain time interval, it assumes that the transmission has failed; it then increases the upper bound (CW_{max}) for the backoff counter, and retransmits the fragment with the same MAC procedure. This MAC DCF procedure is symmetric, all nodes implement it in the same way. Due to the symmetric contention scheme, access delays vary and depend on the number of contending nodes as well as on transmission errors (due to the impact on backoff counter bounds).

Point Coordination Function: In order to make medium access delays more predictable, a scheme called Point Coordination Function (PCF) has been defined as a part of the IEEE 802.11 specification. The access scheme uses the same MAC procedures, but a coordinator (the access point) provides centralised medium arbitration, prioritizing channel access by utilizing shorter inter-frame spaces for part of the communication. With such higher prioritized frames, it polls the set of nodes for up-link data, i.e., data to be transmitted from a node to the access point. During that Contention Free Period (CFP), the access point grants exclusive access to the medium by transmitting a polling message to some node in the set of nodes. Although this technique avoids contention, the problem of message losses still remains. Further details can be found in [17].

However, PCF does not provide strict time guarantees on its CFP repetition interval and on delays for channel access with multiple PCF coordinators operating on the same channel. Besides, PCF mode is nowadays seeing extremely limited support by available COTS WLAN devices, making it not the preferable choice.

An alternative is to implement the polling for medium access coordination from the higher layer protocols; the second feature, prioritization of polling messages compared to other DCF traffic, can then be achieved by utilizing features of the 802.11e standard. Two features are of particular interest: 1) the use of different inter-frame spaces, and 2) the modification of bounds of random backoff intervals. Both of these can be implemented via 802.11e mechanisms, see [18] for details. The way these are used in the experimental implementation is described in Section IV-A.

III. THE TIMED RELIABLE COMMUNICATION PROTOCOL

This section presents the Timed Reliable Communication protocol (TRC), the real-time communication protocol we designed for 802.11 communication between devices within the ALARP worksite. This protocol is built by modifying the Real-time Group Communication Protocol (RGCP [6], [7]). After presenting the RGCP, the TRC protocol is described, in terms of changes applied to the RGCP protocol. While the RGCP is based on the assumptions of 802.11 PCF, our adapted protocol is targeted for use on top of 802.11 DCF, with the differences explained in Sections II-B and III-C.

A. The base protocol RGCP

The RGCP provides reliable and efficient group communication services, relying on the IEEE 802.11 PCF and on a centralised coordinator (Access Point (AP)). The coordinator schedules channel access for a group of nodes. The AP grants exclusive access to the medium by transmitting a polling message to some node in the group. Hence, the RGCP is based on the assumption of the PCF providing its contention free node slotting mechanism.

Before entering into details of the protocol we present its assumptions and properties. The RGCP is based on the following fault assumptions:

- Messages sent during the CFP are either lost (omission fault), or are delivered correctly within a fixed time-bound tm. The losses may be asymmetric i.e., some nodes may receive a broadcast message and some may not. The number of consecutive message losses is assumed to be bounded by a so-called *Omission Degree* OD [19] that denotes an upper bound on the number of omission faults that may affect a single message. In the case of a broadcast message, this means that after $OD + 1$ transmissions *every* receiver has received *at least one* of these transmissions.
- Nodes may suffer crash failures or, due to their mobility, may leave the reach of the AP at arbitrary points in time (without crashing). This resembles a permanent crash of the link between the AP and the leaving node.
- The AP is stable, i.e., not subject to any kind of error.

We here briefly discuss, and later analyze, the assumption on the omission degree OD. Every protocol that intends to guarantee reliable transmission and real-time properties on an unreliable medium must base on this [6] or define a fault-detection procedure for fault coverage. OD in our setting can be seen as a parameter that can be optimized in order to address the following trade-off: i) selecting an OD large enough to ensure that the coverage of the assumption on the semantics of the omission degree is large enough, and ii) minimizing the additional overhead created by an overly large choice of OD and assuring that the protocol can be executed within the required time-bounds.

The RGCP satisfies the properties of:

- *Validity*. If a correct station broadcasts a message m then it eventually delivers m [20].
- *Agreement*. If a correct station delivers a message m then all correct stations eventually deliver m [20]. Note that validity together with agreement ensures that a message broadcast by a correct station is delivered by all correct stations.
- *Integrity*. For any message m every correct station delivers m at most once and only if m was previously broadcast by its sender [20].
- *Total order*. If the messages m_1 and m_2 are delivered by stations s_1 and s_2, then station s_1 delivers message m_1 before message m_2 if and only if station s_2 delivers message m_1 before m_2.

The RGCP communication is structured into rounds. During each round, the AP polls each node of the group exactly once. After being polled, a node returns a broadcast request message to the access point, which assigns a sequence number to that message and broadcasts it to the node group. The broadcast request message is also used to acknowledge each of the preceding broadcasts by piggy-backing a bit field on the header of the request message. Each bit is used to acknowledge one of the preceding broadcasts.

By this, one round after sending a broadcast message, the access point is able to decide whether each group member has received the message or not. In the latter case, the access point will retransmit the affected message. By the assumptions made above, a message is successfully transmitted after at most OD+1 rounds.

If the AP does not receive the request message within a certain period of time after polling the node, it considers the request message (or polling message) to be lost, and transmits the last broadcast message of the not responding node if it has not yet been acknowledged by all nodes. If the AP does not receive the request message from a node for more than OD consecutive times, it considers that node to have left the group and broadcasts a message indicating the change in the group membership.

To enable the user to improve the timing guarantees, the parameter *resiliency degree* is introduced, to allow the user specifying the maximum number of retransmissions of the messages. This resiliency degree $res(c)$ represents a bound on message retransmissions, which may vary for different *message classes* c. A value $res(c)$ smaller than OD allows trading reliability of message transmission for shorter transmission delays.

The modification introduced in the protocol to consider the resiliency degree are as follow. If a message m is acknowledged by all nodes after at most $res(c) + 1$ rounds, the AP issues the decision to deliver m to the applications, through the broadcast of a *decision message*, which is re-transmitted $OD+1$ consecutive times (to guarantee reception by all correct nodes). If, however, the AP does not receive the acknowledgement of any node after $res(c) + 1$ rounds,

a decision not to deliver m is issued, again through the broadcast of a decision message.

It can be noted that the introduction of the resiliency degree brings advantages of shorter delivery time for a message, obtained by reducing the maximum number of retransmissions for a broadcast message to $res(c)$ times, but comes at the cost of violation of the validity property. Now, a message requested to be broadcast by a correct node may not be received by all the other nodes and therefore not delivered. However, the *agreement*, *integrity* and *total ordering* properties are retained (total ordering and agreement are preserved thanks to the introduction of the decision message).

B. The Timed Reliable Communication (TRC) protocol

The TRC protocol for communication within the worksite is presented by showing modifications and differences w.r.t the RGCP (with resiliency degree) that were applied to fit the ALARP requirements. Three key modifications are identified and described in what follows.

Modification 1. In ALARP there is no need to guarantee total ordering and agreement. Regarding total ordering, it is not required that the broadcasted messages are delivered according to a total order because each message is independent from the others.

Regarding agreement, in ALARP a message may not be delivered to all correct nodes. In fact, each node of the worksite contains its own means to detect the potential loss of critical messages, and react accordingly (e.g., by using timeouts to detect message loss and by safely notifying to the worker the potential hazardous condition due to such loss). Moreover, as soon as a node receives a message, it can start to process it. This implies, that two different nodes can have an inconsistent view of the ALARP system at a given point in time.

Consequently, the TRC protocol relaxes agreement and ordering properties by removing the decision message, that is no more transmitted by the AP. This further shortens the delivery time of a message.

Modification 2. While RGCP is intended only for broadcast, in ALARP reliable broadcast, multicast and unicast are required. The TRC protocol introduces unicast and multicast using a scheme very similar to broadcast. In case of multicast, the procedure is the same but a polled node needs to: i) specify that it wants to transmit a multicast message, and ii) transmit the multicast mask, which enlists the recipients of the multicast. In case of unicast, once polled by the AP, a node simply needs to reply that it has a unicast message for a specific recipient Y.

It is important to note that introducing multicast and unicast requires to further extend the header of the messages, to include the additional information needed (the recipient of the unicast message, or the multicast mask). This modification also requires to further extend the logic

of the AP. In the RGCP, each messages is retransmitted until it is acknowledged by all nodes (or the bounds on allowed retransmissions are reached). This requires that the AP simply needs to collect acks from all nodes. Instead in TRC, to define the delivery status of each message, the AP needs to remember if the acknowledged message is a unicast, multicast or broadcast (and the intended recipients), and define accordingly the delivery status of the message.

Modification 3. The formulation of the RGCP introduces the resilience degree $res(c)$ for different reliability classes, and actually consider the possibility of using different messages classes (reliability levels); however this is not applied explicitly in the usage of the protocol in [6], [7]. The resiliency degree $res(c)$ in ALARP is set to three different values that match the three corresponding message classes introduced in Section II-A:

1) Level high = $res(high)$;
2) Level medium = $res(medium)$;
3) Level low = $res(low)$.

It is expected that $res(high) \geq res(medium) \geq res(low)$.

Information on the message level has to be communicated by the broadcasting node (after receiving the poll from the AP), by properly extending the header of the message.

Table I summarizes the main differences between TRC and RGCP.

Table I
COMPARISON OF RGCP AND TRC.

	RGCP (with $res(c)$)	TRC
Provides Validity	no	no
Provides Agreement	yes	no
Provides Integrity	yes	yes
Provides Total Ordering	yes	no
Accepts levels of reliability for message delivery	yes	yes
Supports Broadcast	yes	yes
Supports Multicast	no	yes
Supports Unicast	no	yes

C. TRC Protocol Modifications for DCF

The TRC protocol for the ALARP system is based on 802.11 DCF, and not PCF mode as the base RGCP. This requires additional TRC functionality, because the protocol cannot rely on PCF contention free periods, built-in maintenance of the node polling list, and the node polling mechanism.

Therefore, the TRC coordinator is enhanced with a scheduler which takes care of polling a defined set of nodes for the required protocol time-slotting, similar to PCF. Within a protocol slot, the coordinator's poll packet or the node's request packet can be lost, or a node might not reply at all. To cover those cases, an additional configurable intra-slot time parameter, the Poll-Request Failed Timeout (*P+R Failed Timeout*) is introduced. This parameter specifies the longest time delay within one slot before the coordinator

sends the broadcast packet. This ensures that even with a failed poll or request transmission, a minimum slot length is left to enable the broadcast packet to be sent.

DCF communication in a selected 802.11 channel will introduce time delay of sending nodes due to the channel sensing and backoff algorithm in the presence of other contending nodes. Although the set of protocol communication nodes (within one ALARP worksite) will be scheduled according to a predefined sequence known to the coordinator, and hence, will not contend, this does not hold for any other node. With 802.11 DCF operating in a shared medium and an a-priori unknown number of contending nodes in the same channel, no upper time bound with transmission guarantee is provided. To mitigate that serious drawback, TRC communication is prioritized as far as possible. To get an advantage in terms of more aggressive access to the channel, the TRC protocol is applying a Quality of Service (QoS) scheme of IEEE 802.11e. In particular, transmission is using a shorter IFS and a lower CW_min setting. This conforms with our approach of avoiding any hardware and low-layer driver software modifications. The detailed parameter settings are described in Section IV-A.

In summary, the DCF version of the timed reliable broadcast protocol includes the following features:

- Explicit polling and maintenance of node membership by additional AP software on top of WLAN Layer 2.
- Introduction of a timeout mechanism at the AP in order to send broadcast messages within a slot despite missing requests messages from the polled node.

The realization of the protocol on top of 802.11 uses the following WLAN configuration settings:

- Prioritization of messages via shorter inter-frame spaces and shorter backoff window sizes using 802.11e.
- Deactivation of Layer 2 retransmissions in order to reduce the variability of Layer-2 transmission times.

The following evaluation section investigates to what extent different parameter settings influence broadcast availability and reliability of this extended approach. One important aspect is the coverage of the OD assumption, whose violation is assumed to lead to a disconnect of the node. The analysis addresses the question, whether a time to disconnect can be achieved which is in a feasible range of at least several hours for realistic parameter choices and settings.

IV. Experimental and Model Study

To evaluate the TRC protocol in an IEEE 802.11 DCF setting both experimental and simulation-model based studies have been conducted. The main objective of the study is to characterize the protocol performance for the DCF realization. The protocol is implemented in a testbed with a limited number of clients. From experimental results in the testbed, the protocol behavior in a real-world scenario can be established. An important part of this analysis is to define how DCF channel contention aspects affect the in-slot packet delay characteristics. This characterization, in the form of transmission delay distributions, is then applied in a stochastic simulation model developed by the authors for this purpose. Subsequently, this stochastic simulation model enables extrapolated analysis to assess the expected long-term operation behavior of the protocol scaled to different scenarios. In the following section the TRC protocol testbed realization is described.

A. Protocol Realization

The TRC protocol implementation consists of two parts: one part is executed on a coordinator, and the other one at each protocol communication node. The coordinator does the node polling and takes care of distributing the messages as explained in the protocol description section. For the experimental evaluation, only broadcast messages are distributed, to delimit the analysis. The protocol performance is evaluated by sending high-priority messages, i.e., using the highest resilience class. Here, a node is then defined to become disconnected when it cannot be reached after OD + 1 rounds. Such disconnection cases are measured in the experimental system and further analyzed in the simulation model.

For realizing the coordinator, an 802.11 Access Point (AP) is used. The main selection criterion for the AP is that it provides an Atheros chipset for wireless communication; this chipset allows an easy configuration and adaption to our needs using an Open-Source driver implementation. The selected AP uses the Atheros AR9287-BL1A chipset, supporting b/g/n communication and having two external detachable antennas.

The standard firmware of the AP has been replaced with OpenWrt [21], a Linux distribution for embedded devices, using the latest trunk version. The adapted TRC protocol is targeted to run on the AP, and therefore implemented using the OpenWrt Linux application programming interfaces. Messages are sent as Internet Protocol (IP) packets using Linux socket interfaces. The sources have been compiled to an OpenWrt package and installed on the AP.

The communication protocol client nodes are executed on a netbook and laptop, respectively, both running Debian, Kernel version 2.6.32-5. The external wireless communication device uses a RALINK RT2870F chipset and is attached via USB cable to the computer.

To ensure sufficient processor time slices, all protocol process priorities are increased to the highest possible.

For realizing TRC communication with 802.11e QoS, both coordinator and client nodes have to support and enable the 802.11e Wi-Fi Multimedia extensions. An early experimental communication test setup conducted in the beginning of the tests have provided the appropriate set of 802.11e configuration values to compete in cross-traffic scenarios. Based on those results, the 802.11 Enhanced

Distributed Channel Access (EDCA) category parameter is set to voice (AC_VO), with the Arbitration Inter Frame Space (AIFS) to 1, and Contention Windows (CW) to 4 (CW_{min}) and 7 (CW_{max}). The wireless communication uses the 802.11b mode with a physical rate of 11MBit/s for all measurements (including both protocol member nodes and contending nodes for cross-traffic). Data-link layer retransmissions are disabled on both the coordinator and client side to mitigate time delay variations within one slot caused by the standard 802.11 binary exponential backoff algorithm; therefore, all retransmissions are performed via the TRC protocol. Because the AP chipset does not allow to reset the maximum retransmission retry count to 0, all AP packets are sent as broadcast frames, thereby avoiding retransmissions. (Hence, with no retransmission, the CW_{max} value does not affect channel access.) Re-enabling Layer 2 retransmission in a future deployment does not affect TRC protocol functionality, but can adversely influence the intra-slot packet delays. The TRC protocol parameter *P+R Failed Timeout* is set to $40\,ms$.

B. Evaluation Approach and Metrics

The study of the protocol realized on DCF is taking a starting point in the execution of two scenarios: 1) *No cross-traffic* - the protocol is executed in an environment with limited traffic (caused by a few distant access points) in the same channel. This setting will provide a reference for the protocol operation under good conditions. 2) *With cross-traffic* – other wireless nodes are introduced in the channel to contend heavily for the channel resources. This setting is relevant to clarify the impact on the packet loss and message delay in scenarios when the TRC protocol will not have exclusive channel rights.

Obtaining long-term behavior of the TRC protocol is important to establish the suitability of the protocol realization for the described application scenario. Obtaining such assessments from an experimental setup is, however, time consuming and impractical as the execution conditions can be difficult to control. Instead, a model of the protocol logic has been implemented in a MATLAB simulation model. The model uses packet loss rates ($\widehat{PLR}$) and empirical transmission delay distributions estimated from the experimental studies. Subsequently, the model enables extrapolations of the measurement results to different protocol parameterizations.

To characterize the communication channel properties and the overall TRC protocol behavior, the following metrics are measured in the experimental setup:

*Packet Loss Ratio (**PLR**):* The packet loss ratio describes the ratio of packets lost in relation to transmission attempts sent. The PLR is used as an estimate of the packet loss probability. The same packet loss probability is assumed for each protocol packet sent, independent on whether it is a poll, request or broadcast. An estimate of PLR is defined


Figure 2. Experimental setup used for performance measurements

as $\widehat{PLR}_{PR}$ and is obtained from a statistic of failed poll-requests at the access point.

Poll+Request Transmission Time [ms]: This metric describes the time within a slot measured at the AP from the moment of sending a poll until receiving the corresponding request. The metric includes poll transmission time, poll processing time at a node and the request transmission time. The metric is calculated at the access point for each successful poll-request pair.

Request to BC Completion Delay (δ_{RCDC}) [ms]: This delay characterizes the time duration from the moment the coordinator (access point) has received a broadcast message for dissemination, until the time the access point has received an acknowledgment from all protocol member nodes. Hence, this metric assesses when the coordinator assumes a broadcast to be completed and when it can continue to process a new message from the given source node.

Request to BC Reception Delay (δ_{RRDD}) [rounds]: The broadcast reception delay is a distributed measurement using round numbers as granularity. It describes the delay measured in terms of number of rounds from when the coordinator has received a broadcast message for dissemination, until the message has been received first time by all protocol member nodes. As each broadcast message contains a retransmission number, this metric can be derived as a discrete round time metric based on node and AP recordings.

C. Experimental Evaluation Scenario

The setup used for the experimental studies is introduced in Figure 2. It consists of two sets of communication devices: 1) The devices implementing the TRC protocol given by the *Main Coordinator* access point and two *Protocol Member Nodes* and 2) a node set of contending nodes communicating via the AP. All nodes are operating in the same wireless channel (channel 1). Thus, the protocol member nodes must contend on the channel with the *Contending Nodes*.

Setting	Value	Setting	Value
Contending Nodes	2	Slot length	$50\,ms$
Protocol Member Nodes	2	res(high)=OD	15
Experiment Duration	$30\,min$	Wireless Mode	802.11b
BC Payload	$58\,bytes$	PHY Rate	11Mbit/s
BC Frame Size	$92\,bytes$	802.11e setting: AC_VO'	$AIFSN = 1$, $CW_{min} = 4$ slot times, $CW_{max} = 7$ slot times

The contending nodes are configured as standard DCF nodes (without 802.11e). These stereotypes are the most common in existing deployments (in offices, homes etc.). In the cross-traffic scenario both nodes are running with a saturated transmission buffer constantly trying to transmit 1506 byte frames to represent a worst-case load. It must further be noted that all wireless devices are located in close proximity. This means that signal conditions are good and that most experienced losses can be attributed to collisions and interference. Interference is likely, as the testbed setup has been deployed in an office scenario with around 12 other access points operating in the 2.4 GHz frequency band. The channel selected for the tests has been chosen such that it contains the fewest other access points. Each experimental execution run takes 30 minutes.

Detailed information on the parameters used for the experimental results are provided in Table II.

D. Experimental Results

The experimental results have been obtained from an execution with and without cross-traffic. The access point notifies a protocol member node when its previous broadcast has been completed in a poll. In the same slot, the node prepares and sends a new broadcast in the request. With this approach, a node performs only one broadcast at the time, and we can thus disregard any queuing in the following analysis (although this aspect must be considered in future work).

The overall outcome of the experiments is summarized in Table III. The main findings are:

- A high packet loss rate has been observed with cross-traffic enabled of almost 18%. This high increase is mainly expected to stem from collisions with the contending nodes. As no re-transmissions are enabled on layer 2, the node does not get the advantage of implementing a smaller contention window compared to the contending nodes (which have retracted form the medium as well) for a fast retransmission. Future work should consider if 1-2 fast (Layer 2) retransmissions could provide an advantage without increasing

transmission delays significantly. Despite the high loss rate the protocol largely runs fine still. The simulation analysis clarifies this aspect in more detail.

- Even for the case without cross traffic a substantial loss rate is observed of 3.5%. These losses may stem from other interfering nodes in wireless neighbor channels or hidden nodes.
- In the average case, the request to broadcast completion delay increases by around $70\,ms$ for δ_{RCDC} for a node when cross-traffic is activated.
- The maximum value δ_{RRDD} (worst case) in Table III with cross-traffic OFF stems from the transmission irregularities, where a nodes requires that many rounds to finally receive the broadcast.
- When studying the general protocol execution an irregularity has been identified in the communication from the mobile nodes participating in the TRC. The irregularity consists of a systematic transmission interruption of 1 seconds, approximately every 60 seconds. The interruption is TRC protocol independent as it also occurs when running periodic message transmissions with test tools (using $mgen$ and $wireshark$ monitoring traces). This interruption is presumably a driver issue that must be solved in future implementations. The irregularity to a large extent does not influence the measurement data for the processing of intra-slot delay metrics. But it has an impact on BC reception delay, see further below.
- An impact of the node transmission irregularity can best be observed in relation to disconnects. Recall, that the coordinator will consider a node to be disconnected if it has not responded for OD+1 rounds. In the cross-traffic case two of such disconnects have been observed. These are expectedly a result of the transmission irregularity that can stop the node from responding for a large part of a round. Together with the added delays and losses from cross-traffic, this provokes disconnects. This is especially the case with a small set of nodes and small slot sizes, because a single irregularity will severely affect the nodes ability to communicate within the protocol's upper bound limit. For the similar run without cross-traffic, no disconnects are observed. Here the redundancy of the protocol is sufficient to avoid the interruption.

The effects of the cross-traffic scenario are best observed in Figure 3. It shows the distribution of the poll+request transmission times (and associated processing time). It is seen that in both cases, this message flow is completed within a few ms. Even for the cross-traffic case the transmission times are far from the intra-slot timeout of $40\,ms$. This aspect can partially be described by the aggressive channel access scheme applied through 802.11e and the fact that only a single transmission is attempted. In normal DCF, eleven

Table III
EXPERIMENTAL RESULTS.

| Experiment	Parameter	Value
Cross-traffic OFF	$\widehat{PLR_{PR}}$	0.035
	Disconnects	0
	δ_{RRDD} (avg./max)	0.2/12 [rounds]
	δ_{RCDC} (avg./max)	117/1405 [ms]
	corresponding to	2.3/28 [slots]
Cross-traffic ON	$\widehat{PLR_{PR}}$	0.177
	Disconnects	2 (Node 1)
	δ_{RRDD} (avg./max)	0.12/10 [rounds]
	δ_{RCDC} (avg./max)	188/1480 [ms]
	corresponding to	3.8/29 [slots]

Figure 3. Transmission time distributions for successful poll and request message pairs in a 50 millisecond slot (joint over both nodes).


Figure 4. Comparison of broadcast delays with (lower figures) and without (upper figures) cross-traffic. The gray color shows the time until broadcast reception at the MT, while the dark color shows the time until all Acks have been received by the AP. The right part shows a zoom into the left curves. The curves show a step function with width of the steps equal to the slot-time of 50ms, as the reception events are evaluated at the end of each slot.

retries are standard, which in conjunction with increasing channel backoff times can lead to excessive delays.

These results are encouraging to evaluate shorter slot times. This aspect will be studied in the simulation model evaluation which uses these delay distributions as an input.

Having established in Table III the fundamental settings of the experimental results, the cumulative distribution functions in Figure 4 depict several interesting aspects regarding the broadcast transmission delays. The upper figure shows the case without cross-traffic. Considering initially δ_{RCDC} (dark color), it is clear that most broadcasts (around 99%) are successfully completed after round 3. In practice, most nodes have received the broadcast already in second round (in around 99% of the cases). All messages are successfully completed within the 16 (OD+1) rounds. This implies that no disconnects are observed.

Studying the cross-traffic case in the lower part, it can be observed that all nodes receive most broadcasts (around 99%) already in round 3. However, only at round 9, the coordinator has received all acknowledgements. This impact clearly stems from the increased packet loss rate. In the cross-traffic case, we observe that OD+1 has been exceeded two times, i.e., two disconnects have been observed.

E. Simulation Based Results

In order to extend the observations of the experimental results, the TRC simulation model has been parameterized from the experimental results. The parametrization has been

based on the estimated packet loss ratios provided in Table III, and the delay distributions of Figure 3. In the simulation model packet losses are assumed to be independent. Also, a broadcast is considered to be received by all nodes or no node at all, resembling the experimental test setup with spatially close nodes where packet losses are only resulting from collisions which affect all receiving nodes equally. In terms of transmission delays, the poll and request durations are sampled from the empirical delay distribution within the range up to the 90 percent quantile. Both with and without cross-traffic, long tails have been observed in the measurements with delay samples up to $40\,ms$ in rare cases. Also, as the slot moves on with a broadcast after $40\,ms$, longer delays that may have occurred are not in the distributions. As these longer delays are important to simulate, the tail of this distribution has been fitted separately. To mimic the high variability, a 2-parameter Pareto distribution has been fitted to match the largest 10% of the data of the empirical delay distribution. This fitted Pareto distribution is then used in the simulations for the tail samples.

Regarding the parts of the delay distribution before the Pareto tail, the experiments provide a delay distribution of the poll and request pair. To define a delay distribution for the individual transmission of a broadcast, the following assumption has been made: a poll and a request individually experience delays sampled from an independent and identically distributed shifted exponential (shifted as there

CDF comparing the Broadcast Request to Completion Delay Times

- - - Measured Request to Completion Delay c–t
—— Simulated Request to Completion Delay c–t
- · - Measured Request to Completion Delay
—— Simulated Request to Completion Delay

Figure 5. Comparison of experimental and model results of the broadcast request to completion delay.

is a minimum transmission delay). Thus, it is assumed that the joint distributions is a shifted 2-stage Erlangian, and that the parameters of the exponentials can easily be determined via least square fitting. Combining this shifted exponential sampling with the Pareto tail sampling, the simulation model can now be executed with parameters matching the experimental results.

Validation of simulation model: To assess the results from the simulation model in comparison to the experimental results, a comparison of the 'request to completion' delay produced by both is shown in Figure 5.

The results show generally a good correspondence between the simulation model and experimental results. However, in the cross traffic case the simulation model is generally a little more optimistic. A part of the explanation may stem from the fact that packet loss ratios to the individual nodes are not completely alike and also not independent identically distributed in reality (leading to some shifts in times to when a broadcast can be considered completed). The simulation model does not capture this effect. Another aspect is the systematic transmission interruptions in the simulation runs that make the tails a bit longer for the experimental results compared to the model (which does not include these long transmission interruption delays). As was observed that the transmission interruption delays are not related to the TRC design and implementation on DCF, we concluded that the slightly more optimistic model produces valid evaluation results.

Extrapolation of results: Having parameterized the simulation model, several extrapolation analysis experiments can be conducted via the simulation. Most interesting is to clarify how the protocol will operate in scaled up scenarios (to match the scenarios introduced in Section II) with more mobile terminals participating in the timed reliable broadcast. Further, the application-specific worst-case requirement for delivering a safety-critical message is 10 seconds (Section II-A). According to the experimental

results above, an OD of 15 is sensible to the retransmission attempts experienced, and makes it possible to tolerate significant interruptions in the transmission – see Figure 4. This leads to analyzing shorter slot sizes below $50\,ms$. With the TRC protocol worst-case execution time corresponding to $2 \times OD + 1$ rounds, we analyze cases of:

- S1: 20 nodes, slot size = $15\,ms$ (P+R Failed Timeout of $5\,ms$)
- S2: 16 nodes, slot size = $20\,ms$ (P+R Failed Timeout of $8\,ms$)
- S3: 12 nodes, slot size = $25\,ms$ (P+R Failed Timeout of $8\,ms$)
- S4: 10 nodes, slot size = $30\,ms$ (P+R Failed Timeout of $10\,ms$)
- S5: 6 nodes, slot size = $50\,ms$ (P+R Failed Timeout of $10\,ms$)

The case S1 with 20 nodes is interesting because it matches one of our use-case scenarios with its typical maximum set of communicating nodes. Here, the slot size has to be reduced to $15\,ms$ to still be able to complete a broadcast within 10 seconds (at the given value of $OD = 15$). The second scenario increases the slot size to $20\,ms$, S3 to $25\,ms$, S4 to $30\,ms$, and S5 to $50\,ms$, leading to number of nodes 16, 12, 10, 6, respectively, such that the broadcast can finish in the time bound of 10 seconds. According to Figure 3, a large part of the Poll and Request exchanges will finish well before $30\,ms$.

From a safety-critical perspective, a hazard event can be considered when a message is not successfully transmitted within the time deadline or when a node is disconnected (and cannot receive safety-critical event notifications). In the studied case these two events are actually the same. To assess to which extent such events can be avoided in the operation period of the communication system, we analyze the probability estimate that a 12-hour workday passes without any disconnections. The test setup is such that after completion of an ongoing broadcast, the node will immediately schedule a new broadcast in its next slot.

The following extrapolation results are shown for nodes under the cross-traffic scenario. For nodes under the non cross-traffic scenario, the selected slot durations and *P+R Failed Timeouts* of settings S1 to S5 do not significantly affect the disconnect probability results, due to the comparably short Poll-Request delays, see Figure 3. Therefore, the non cross-traffic case is omitted here.

Figure 6 shows the probability estimators created from each 200 simulation runs of 12 hours simulated time under the cross-traffic scenario. The settings with the longest slot-times (and less than 20 nodes), S4 with $30\,ms$, S3 with $25\,ms$ as well as S5 with $50\,ms$ slot duration, complete the workday with 0 and 2 disconnects in all 200 runs, respectively. S2 has 6 disconnects. The scenario S1 with $15\,ms$ slot durations runs exhibit a high probability of disconnect. For each estimator, the 95 % confidence interval

Figure 6. Probability to success estimators of the 12 hour runs with 95% confidence intervals, for cross-traffic scenario.


Figure 7. Number of disconnects per hour.

is given. The number of broadcasts in each setting before disconnect still increases from S1 to S5, with (rounded) 12127, 45808, 53011, 58967, and 67868 in the mean.

A further illustration of the disconnection behavior of Scenario S1 ($15\,ms$ slot duration) is shown in Figure 7. The figure shows the histogram of the time to disconnect resulting from the 200 simulation runs for S1, and the 95 % confidence interval for each disconnect estimator. Almost 25 % of runs disconnect during the first hour, and slightly over 15 % disconnect in the second hour. During hour 11 to 12, around 2,5 % of TRC protocol runs result in a disconnect of the node from the coordinator. Overall, only 9 runs out of 200 did not disconnect for the simulated 12-hour workday.

The results of S1 to S5 indicate that stable settings can be found, but that optimizations must be considered to further scale the solution. One option may be to re-introduce Layer 2 retransmissions, using only one or two frame-retransmissions to limit the frame time delay variability.

V. CONCLUSION AND OUTLOOK

This paper describes the design and performance analysis of a time-bounded protocol implemented on top of IEEE 802.11e DCF. A protocol based on 802.11 PCF for reliably broadcasting messages in WLAN is adapted to the use-case of a safety-critical warning system. A subset of these adaptions is implemented for an experimental study. The motivating use-case is message dissemination in a safety-critical warning system, using off-the-shelf communication equipment and technologies. Although the initial use-case description and the performed evaluation targets a safety-critical system for railway workers, the overall protocol design is generic enough to apply to other scenarios with similar safety and timing requirements. The performance analysis is based on experimental measurement results that are used to calibrate a stochastic simulation model.

The challenging part in using 802.11 with distributed access to the communication channel is the inability to make any guarantees on the number of other communicating nodes on the same channel. Besides, wireless transmission exhibits a high packet loss-rate compared to wired communication, and are subject to interference of neighbour channels. DCF adds message send delays which are variable due to the default channel-sensing and backoff algorithm specified in 802.11. Therefore, the TRC protocol implementation on top of DCF did not make use of link-layer retransmissions, to better focus on the actual TRC protocol performance for our evaluation. Instead, packet retransmissions were performed by the broadcast protocol itself. Link-layer transmission was employing 802.11e QoS prioritization. For the experimental test setup, the QoS settings of 802.11e have been chosen to allow an aggressive channel access behavior of the protocol.

The experimental study is conducted using one coordinator and two terminal nodes. The metrics for evaluation included packet loss ratio, poll+request transmission times and the time from broadcast request to broadcast completion. The test runs are both done with high-load cross-traffic on the same channel, and without cross-traffic. With cross-traffic on, a relatively high packet loss rate of 18% is observed; despite the high packet loss, the protocol is still able to successfully distribute and complete broadcasts. On average, the request to broadcast completion time increases by $70\,ms$ with cross-traffic enabled.

The stochastic simulation model serves both for validating the chosen experimental setup parameters, and allows scaling the node and slot size parameters. The model uses a mixture of the empirical delay distribution and a parametric fit of the delay tails to draw samples for message delay. Furthermore, the observed packet loss rate from the experimental setup is used in the model. The simulation model shows to correspond well to the experimental measurements. It performs slightly more optimistic, which is caused by independence assumptions on delay and packet loss, as well

as by the irregularities of the node behavior caused by the driver software in the experimental setup. The simulation model is then used to analyze scenarios of extrapolated settings with slot sizes from 15 to $50\,ms$, together with a higher number of nodes. The node number is chosen in a way that the application-specific worst-case broadcast completion time of 10 seconds is met. The results show that disconnections due to failed broadcasts within a workday of 12 hours happen very rarely (0 to 2 in 200 runs) in scenarios of 30 to $50\,ms$ slot sizes; such a disconnect requires the affected worker node to move to a safe zone. The time to disconnect is getting significantly shorter than 12 hours for the other scenarios. Therefore, the current realization can support at most 12 nodes. Future work will include the development of analytic models for the disconnection probability, such that the analysis of scenarios with very low disconnection probabilities becomes computationally feasible; alternatively, rare event techniques could facilitate an efficient simulation analysis of such settings.

A future detailed study could analyze and compare the protocol using various IEEE 802.11e settings, also in the presence of other interfering nodes having QoS features enabled. Additionally, the probable positive impact on packet transmission success when allowing one or two link-layer retransmissions should be evaluated.

REFERENCES

[1] M. Satyanarayanan, "Pervasive computing: vision and challenges," *Personal Communications, IEEE*, vol. 8, no. 4, pp. 10 –17, Aug. 2001.

[2] B. Hughes, R. Meier, R. Cunningham, and V. Cahill, "Towards real-time middleware for vehicular ad hoc networks," in *Proceedings of the 1st ACM international workshop on Vehicular ad hoc networks*, ser. VANET '04. New York, NY, USA: ACM, 2004, pp. 95–96. [Online]. Available: http://doi.acm.org/10.1145/1023875.1023894

[3] S. Eichler, "Performance Evaluation of the IEEE 802.11p WAVE Communication Standard," *2007 IEEE 66th Vehicular Technology Conference*, pp. 2199–2203, Sep. 2007.

[4] C. Basile, M. Killijian, and D. Powell, "A Survey of Dependability Issues in Mobile Wireless Networks," LAAS CNRS Toulouse, Tech. Rep., 2003.

[5] "ALARP – A railway automatic track warning system based on distributed personal mobile terminals – FP7-IST-2010-234088 http://www.alarp.eu/."

[6] M. Mock, E. Nett, and S. Schemmer, "Efficient Reliable Real-Time Group Communication for Wireless Local Area Networks," in *EDCC*, 1999, pp. 380–400.

[7] E. Nett and S. Schemmer, "Reliable Real-Time Communication in Cooperative Mobile Applications," *IEEE Trans. Computers*, vol. 52, no. 2, pp. 166–180, 2003.

[8] A. Lindgren, A. Almquist, and O. Schelen, "Evaluation of quality of service schemes for IEEE 802.11 wireless LANs," in *Local Computer Networks, 2001. Proceedings. LCN 2001. 26th Annual IEEE Conference on*, 2001, pp. 348 –351.

[9] J. Sobrinho and A. Krishnakumar, "Quality-of-service in ad hoc carrier sense multiple access wireless networks," *Selected Areas in Communications, IEEE Journal on*, vol. 17, no. 8, pp. 1353 –1368, Aug. 1999.

[10] A. Leonovich and H.-W. Ferng, "A time slots coordination mechanism for IEEE 802.11 WLANs," *Communications Letters, IEEE*, vol. 14, no. 4, pp. 360 –362, Apr. 2010.

[11] F. Guo and T. Chiueh, "Software TDMA for VoIP Applications Over IEEE 802.11 Wireless LAN," in *INFOCOM 2007. 26th IEEE International Conference on Computer Communications. IEEE*, may 2007, pp. 2366 –2370.

[12] "Office of Rail Regulation, Annual Report on Railway Safety, 2005 - http://www.rail-reg.gov.uk/upload/pdf/296.pdf."

[13] D. Druidi, "Railroad-related work injury fatalities, Monthly Labor Review, 2007."

[14] A. Seminatore, L. Ghelardoni, A. Ceccarelli, L. Falai, M. Schultheis, and B. Malinowsky, "ALARP (A Railway Automatic Track Warning System Based on Distributed Personal Mobile Terminals)," in *submitted to Transport Research Arena (TRA) - Europe 2012*, 2012.

[15] ALARP Consortium, "ALARP, D1.2 – Requirements Specifications," 2010.

[16] G. Bianchi, "Performance analysis of the IEEE 802.11 distributed coordination function," *Selected Areas in Communications, IEEE Journal on*, vol. 18, no. 3, pp. 535–547, 2000.

[17] B. Sikdar, "An analytic model for the delay in IEEE 802.11 PCF MAC-based wireless networks," *Wireless Communications, IEEE Transactions on*, vol. 6, no. 4, pp. 1542–1550, 2007.

[18] G. Bianchi, I. Tinnirello, and L. Scalia, "Understanding 802.11e contention-based prioritization mechanisms and their coexistence with legacy 802.11 stations," *Network, IEEE*, vol. 19, no. 4, pp. 28–34, 2005.

[19] G. Grünsteidl and H. Kopetz, "A Reliable Multicast Protocol for Distributed Real-Time Systems," in *Proceedings of the 8th IEEE Workshop on Real-Time Operating Systems and Software*, 1991.

[20] V. Hadzilacos and S. Toueg, "A Modular Approach to Fault-Tolerant Broadcasts and Related Problems," Ithaca, NY, USA, Tech. Rep., 1994.

[21] "Openwrt–Wireless freedom. Project website: openwrt.org."