Technical University of Denmark

DTU

Performance evaluation of a multi-radio, multi-hop ad-hoc radio communication network for Communications-Based Train Control (CBTC)

Farooq, Jahanzeb; Bro, Lars; Karstensen, Rasmus Thystrup; Soler, José

Link back to DTU Orbit

DTU Library
Technical Information Center of Denmark

Performance evaluation of a multi-radio, multi-hop ad-hoc radio communication network for Communications-Based Train Control (CBTC)

Jahanzeb Farooq, *Member, IEEE,* Lars Bro, Rasmus Thystrup Karstensen, José Soler, *Senior Member, IEEE*

Abstract—Communications-Based Train Control (CBTC) is a modern signalling system that uses radio communication to transfer train control information between the train and the wayside. A vast majority of CBTC systems worldwide use IEEE 802.11 Wi-Fi as the radio technology mostly due to its cost-effectiveness. The trackside networks in these systems are mostly based on conventional infrastructure Wi-Fi. It means a train has to continuously associate (i.e. perform handshake) with the trackside Wi-Fi Access Points (AP) as it moves. This is a time-consuming process associated with a certain delay. Additionally, these APs are connected to the wayside infrastructure via optical fiber cables that incurs huge costs. This paper presents a novel design in which trackside nodes function in ad-hoc Wi-Fi mode, which means no association has to be performed with them prior to transmitting. A train simply broadcasts packets to any nodes in its range. A node upon receiving these packets forwards them to the next node and so on, forming a chain of nodes. Following this chain, packets arrive at the destination. To make the design resilient against interference, transmissions are separated on multiple frequencies. Furthermore, redundancy is introduced in the design as a node forwards packets to not only one but two of its neighbors. This paper investigates the performance of the new design from the perspective of resiliency, redundancy and scalability, and presents the results both from a field experiment carried out using prototype hardware and an extensive simulations study.

Index Terms—Railway signalling, rail transport, communication-based train control, CBTC, radio communication, Wi-Fi, Wireless LAN, IEEE 802.11, ad-hoc, multi-radio, multi-hop, multi-frequency

I. INTRODUCTION

Communications-Based Train Control (CBTC) is a modern railway signalling system that uses radio communication to transfer train control information between the train and the wayside. This results in high resolution and real-time train control information which increases the line capacity by safely reducing the distance (headway) between trains running on the same track. A sharp increase in the popularity of the CBTC systems is currently being witnessed among mass-transit railway operators. Despite its short range and lack of support for mobility, the IEEE 802.11 WLAN, also known as Wi-Fi, has prevailed as the radio technology of choice for CBTC systems, mainly due to its cost-effectiveness [1].

To ensure a continuous wireless connectivity, hundreds of Wi-Fi Access Points (APs) are installed at the trackside. Each AP is next connected to the wayside (normally a Traffic Control Center (TCC)) via Ethernet. The train must associate (i.e. perform handshake) to an AP first to be able to transmit, just like in an ordinary infrastructure Wi-Fi network. However, there are a number of challenges. Firstly, installing cables to connect each AP to the wayside is expensive and time-consuming. The cost of installing optical fiber cable can be as high as around 30,000 EUR per kilometer. Secondly, the train on-board equipment must roam (handover) from one AP to the other as the train moves. The IEEE 802.11 technology was originally designed for users in stationary office environments and thus inherently lacks support for mobility. Therefore, complex roaming algorithms are employed by CBTC systems to solve this problem. Often different roaming algorithms are employed for train's front and rear antennas due to different alignments with respect to the antennas on the trackside APs [1]. Nonetheless, roaming results in delays in communication and limits the supported train speed as well.

Additionally, the roaming decision is often made based solely on the signal quality received from the AP. This might lead to a situation where train connects to an AP which, despite having an acceptable signal quality, is not fully functional, in that it is not capable of communicating to the wayside, for example due to a software bug, a broken Ethernet cable, an incorrectly configured firewall, etc.

This paper presents a novel design for an ad-hoc based radio communication network (patent pending [2]). In this design, there are no "APs". Nodes function as plain Wi-Fi nodes, in an ad-hoc manner. A node broadcasts packets to the nodes within its range. A nearby node, upon receiving a packet, re-transmits (forwards) the packet, which is then picked up by the next nearby node. A chain of nodes is thus formed, following which the packets reach at the last node in the chain, and are forwarded to the wayside backbone over a wired link. Thus, a train does not have to worry about first associating with an AP as well as roaming. Wired links between the nodes and wayside backbone are no longer needed, except the two

J. Farooq is with Siemens A/S, 2750 Ballerup, Denmark. He is also with DTU Fotonik, Department of Photonics Engineering, Technical University of Denmark, 2800 Kgs. Lyngby, Denmark (e-mail: jahanzeb.farooq@siemens.com).

L. Bro is with Nyantec UG, 10117 Berlin, Germany (e-mail: bro@nyantec.com).

R. T. Karstensen is with Siemens A/S, 2750 Ballerup, Denmark (e-mail: rasmus.karstensen@siemens.com).

J. Soler is with DTU Fotonik, Department of Photonics Engineering, Technical University of Denmark, 2800 Kgs. Lyngby, Denmark (e-mail: joss@fotonik.dtu.dk).

nodes at each end of the chain. Furthermore, to make the chain resilient against failures and interference, transmissions are separated on multiple frequencies and a node forwards packets to two of its neighbors rather than one. Additional advantage of this design is that a node can be placed anywhere at the trackside and not only at designated points where connections to the pre-installed Ethernet cable are accessible.

Note that while CBTC traffic itself does not require high data rates as discussed further in Section III, one objective of this work is to study how much data rate the proposed design can support as any excessive bandwidth can be utilized for providing modern, non-CBTC applications such as remote diagnostics and maintenance, remote software upgrade, CCTV, passenger infotainment, onboard Internet etc., which are likely to become more widespread in near future. Likewise, it is noted that despite originally intended for a CBTC trackside network, the application of the proposed design is not limited to it, e.g. it can serve as a superior alternative to the conventional "Wi-Fi over Long Distance" (WiLD) method used to provide low-cost, long-distance wireless access to rural areas.

It is worth noting that the proposed design is not limited to using Wi-Fi. For example, LTE, with the Device-to-Device (D2D) support in its upcoming releases, can be used as an alternative radio technology.

Two experimental studies were carried out to study the performance of the design primarily in terms of number of packets transferred across the chain, the resiliency and redundancy enabled by it, and its scalability. A field experiment [3] was carried out first to provide a proof-of-concept. A simulation study was performed next to verify the findings of the field experiment and to extend the study to a larger scale with various additional scenarios [4]. The simulation study was extended further to investigate potential improvements to the design. This paper provides an overview of the these studies and discusses results.

The rest of this paper is laid out as follows. Section II discusses related work. Section III presents an overview of the CBTC systems. Section IV provides an overview of the proposed design. Sections V and VI provide an overview of the field and simulation studies, respectively, and present their results. Finally, Section VII concludes the paper.

II. RELATED WORK

A multi-hop ad-hoc network formed as a chain of nodes presents a suitable candidate for a long-distance network. Most of the related work, however, focuses on networks where all nodes operate on a single frequency. Since nodes must forward packets for other nodes, the capacity degrades sharply with the growing size of the network as a node must contend with additional nodes than its two immediate neighbors. Thus, these networks offer only a fraction of the capacity achieved by a single-hop network, as the capacity drops to one-half with each hop and to 1/7 as the number of nodes increases beyond 10 [5], [6]. Additional reasons include the "hidden node problem" in which two nodes are in the transmission range of a common node but not in each other's range. The hidden node problem is well-known in the context of the

conventional infrastructure Wi-Fi networks. In a multi-hop scenario where two nodes communicating to another node are not necessarily in each other's range, the hidden node problem is inevitable. Furthermore, IEEE 802.11's Carrier Sense Multiple Access/Collision Avoidance (CSMA/CA) mechanism—which is based on carrier sensing—does not work ideally in wireless networks where the interference range is often larger than the transmission range, as the power sufficient to introduce noise in a transmission is much lower than that required for a successful transmission [7]–[9].

Unfortunately, there exist very limited research work on multi-hop ad-hoc networks in the context of vehicular communication, in particular CBTC, and with the objectives of our proposed design. The IEEE 802.11p [10] standard—also known as Wireless Access in Vehicular Environments (WAVE)—is targeted towards vehicular communication, in particular roadway safety, and is based on the Dedicated Short Range Communications (DSRC) standard. The works in [11]–[13] discuss the multi-radio/multi-channel feature of this technology, although the focus is single-hop communication. The works in [14]–[17] discuss multi-hop communication, although for vehicle-to-vehicle (V2V) communication rather than vehicle-to-infrastructure (V2I) communication, where roadway vehicles relay messages to each other. An advanced version of an ad-hoc based network is a Wireless Mesh Network (WMN) which employs a multi-radio design, and despite it merits, has not been considered for CBTC except for a limited application in [18]. The work in [19] discusses an IEEE 802.11-based WMN with the focus of studying handover delay, although not in the context of rail transport.

III. OVERVIEW OF CBTC SYSTEMS

A brief overview of CBTC is presented here. For a more detailed version, refer to [1]. In CBTC, radio communication is used to exchange train control information between the train and the wayside, enabling Automatic Train Control (ATC) functions. ATC is composed of the Automatic Train Protection (ATP), Automatic Train Operation (ATO) and Automatic Train Supervision (ATS) subsystems. The train regularly sends its state—which includes the current speed, direction, and location—to the wayside over the radio connection. The train determines its speed and location using devices such as speedometers, tachometers, transponders ("balises"), Doppler radar and odometers. Based on this information, the wayside ATC equipment calculates the "limit of movement authority" (LMA) information and sends it back to the train. LMA basically includes the maximum speed and distance the train is permitted to travel. Based on LMA, the onboard ATC equipment ensures that the train speed and the safety distance to the preceding trains conforms to the required limits. Thanks to this real-time communication between the train and the wayside, the so-called "moving block operation" is realized that allows trains to run closer to each other. Furthermore, the number of trackside equipment—such as color light signals and track circuits—required by the conventional signalling systems is minimized. Fig. 1 illustrates typical wayside—which includes trackside—components of a CBTC system. The wayside ATC

subsystems additionally perform functions including scheduling trains and determining their destination/dwell times. These subsystems are often collectively referred to as the Traffic Control Center (TCC) [1].
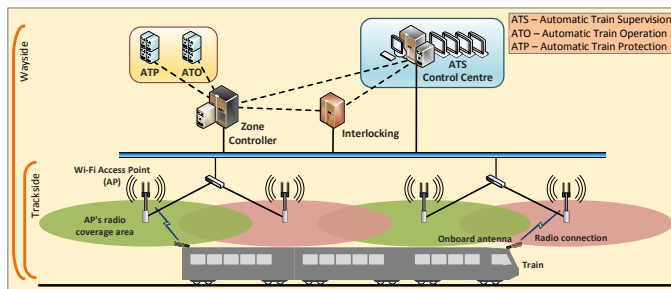

Fig. 1. CBTC wayside components

Multiple Wi-Fi APs are deployed at the trackside. Fig. 1 uses the green and red colors to differentiate between the APs' coverage areas. Each AP is connected (over a wired link) to the wayside components through the backbone network. A train communicates to an AP through a radio connection. It has to continuously search for a new suitable AP and re-associate as it moves along. To assist in roaming (handover), APs are placed in a way that their coverage areas overlap.

The radio communication system of nearly all of the CBTC systems worldwide use the license-free ISM (Industrial, Scientific and Medical) bands (2.4 and 5 GHz).

Normally, CBTC control messages are sent at regular, short intervals of 100-600 milliseconds. This guarantees that the two sides always receive the most updated information (i.e. train state and LMA) from each other. Data requirement for a CBTC system is typically in the range of 20-100 kbps. Additionally, these messages are often transmitted simultaneously on multiple radios. Therefore, IEEE 802.11 MAC retransmissions are often not necessary [1].

A. Handover in IEEE 802.11 Wi-Fi

The handover process in IEEE 802.11 is divided into 3 phases, namely scanning, authentication and re-association. Of these, the first two contribute greatly to the latency of the process. Scanning is the process of finding a suitable AP to connect to. Studies show that the latency of this phase accounts for approximately 90% of the total handover latency [20]. To minimize latency, in CBTC systems, active scanning is often adapted in which a node proactively announces that it is searching for APs rather than wait for announcements from APs. Nonetheless, to ensure a fail-safe function, onboard radios are occasionally configured to associate to any available APs regardless of what frequency they are operating on. Thus, scanning has to be done on all frequency channels, taking a significant amount of time. In the authentication phase, the node establishes its identity with the AP it has just found by exchanging special authentication messages. The latency of this phase is implementation dependent. It could take significant time if a centralized security mechanism is used that involves communicating with an authentication server [1].

The handover in IEEE 802.11 is the so-called "hard handover", in which the mobile node breaks the association to the current AP before establishing a new association to another AP. A critical aspect of roaming in CBTC thus is how the train smoothly switches from one AP to another, without causing interruptions in the communication. A large handover latency might result in a train failing to receive the location of the preceding train in-time [21], [22]. The distance between APs—referred here as "inter-node distance"—and the train speed determine the frequency of handovers. Handover time in CBTC is typically in the range of 70-120 milliseconds. Studies show that the number of packets lost due to handover is much larger than that due to radio propagation effects and put the number as high as 10% [23]. Normally a smooth transition is ensured by equipping a train with at least two radios, one at each end, such that one of these radios stays associated to the current AP while the other switches to a new AP [1].

IV. PROPOSED NETWORK DESIGN

Fig. 2 (a) illustrates the conventional network design for CBTC trackside. The idea behind the proposed design is to take advantage of the broadcast nature of radio communication to present a replacement for the conventionally wired trackside network. Thus, at its basic, a train broadcasts packets which are then picked up by a node in the chain and forwarded to its neighboring node, and so on. No AP scanning and association are thereby required. This is illustrated in Fig. 2 (b).
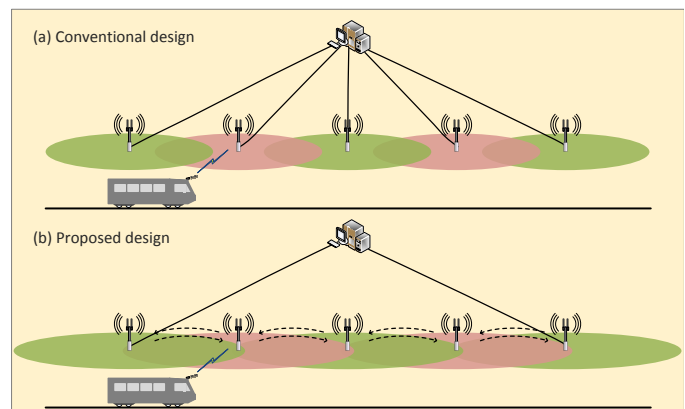

Fig. 2. CBTC trackside network: Conventional vs. proposed design

Note that this saves the train 70-120 milliseconds—the typical handover latency in CBTC as noted above—every time it transmits to a new node in the chain. Additionally, note that since the association phase is skipped, the proposed design uses the higher layer protocol IPSec for authentication and end-to-end encryption to secure the communication.

However, there lie two major challenges with this approach.
1) If all nodes transmit on the same frequency, the probability of interference rises sharply. Additionally relevant is the well-known "hidden node problem".
2) A single failed node results in a practically broken chain (i.e. a single point of failure).

In a "hidden node problem", two nodes are in the transmission range of a common node but not in each other's

range. Since they cannot hear each other, it effectively renders CSMA/CA used in IEEE 802.11 MAC to avoid collisions ineffective. In the conventional infrastructure Wi-Fi networks, this problem is solved by employing the Request to Send/Clear to Send (RTS/CTS) mechanism. There are two reasons why the RTS/CTS mechanism is not applicable in the context of our proposed solution. Firstly, for this mechanism to work optimally, all the nodes that might interfere must be in the transmission range of the node that sends the CTS message, such as an AP. However, as previously discussed in Section II, in an ad-hoc and multi-hop scenario, nodes are not necessarily in each other's range, such as the one shown in Fig. 2 (b). Secondly, the idea behind our proposed design is to broadcast, in which case the RTS/CTS mechanism is irrelevant.

A. Frequency separation and redundancy

To solve the interference problem, the proposed design uses three frequencies to ensure a certain separation between nodes transmitting on the same frequency just like the frequency-reuse in cellular networks. Each node is equipped with three radios, all on different frequencies. The two side radios use directional antennas one in each direction. These radios are used both for transmitting and receiving. Transmissions are made not only to the immediate neighbor node but also the following node. Thus, for a network with an inter-node distance of e.g. 600 meters, the transmissions must be received 1200 meters away. The third top radio is equipped with an omni-directional antenna and is used only for receiving. The three frequencies are then used in an alternating fashion on subsequent nodes. A predefined address included in each packet indicates the direction of the traffic flow. The three radios on a node are connected to each other via Ethernet and work cooperatively. As a radio receives a packet, depending on the direction of the traffic, it delivers the packet to the correct side radio (i.e. left or right) which transmits it further. This forms a "rope-like" interleaving.

Fig. 3 illustrates the mechanism, where the colors red, blue and green represent three frequencies. Note that the two transmission lines coming out of Node 1's right radio (blue) are shown only to emphasize that the transmissions are received on both Node 2 and Node 3. Nonetheless, in reality, it will be one broadcast transmission received at both nodes. The arrows on the lines indicate that a one direction flow (from left to right) is depicted. In an illustration of a two direction flow, there will be equivalent lines in the opposite direction, i.e. two lines from Node 3's left radio coming into Node 2's top radio and Node 1's right radio, and a line from Node 3's left radio coming into Node 2's top radio. As seen, the radios transmitting on the same frequency on two adjacent nodes face opposite—e.g. red radios on Nodes 1 and 2—thus ensuring frequency separation with the help of directional antennas. The 2-node transmission range solves the "single point of failure" problem and introduces redundancy to the design as a node receives the same packet from two nodes rather than one. It further solves the "hidden node problem" by ensuring that two nodes transmitting to a third common node are always in each other's range, e.g. Node 1 right radio and Node 3 left radio transmitting to Node 2 top radio.
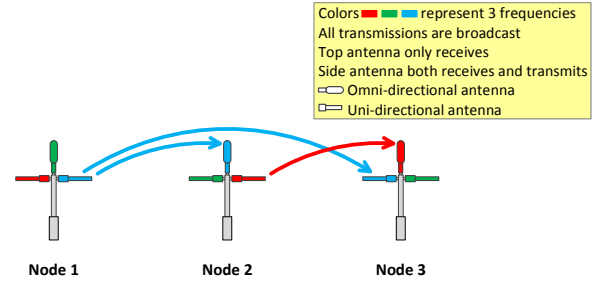

Fig. 3. The proposed design with a one-direction traffic flow (left to right)

B. Chain node vs. terminal node

The node type discussed above is referred to as a "chain node", as these nodes are what make the chain. A second type of node is the "terminal node", which is either a train or a TCC (Traffic Control Center). Basically, it is the node that uses the chain network to get its packets transferred to another terminal node at the other end of the chain. A train intends to send packets to the TCC, and a TCC intends to send packets to one or more trains. A train travels along the chain and broadcasts packets, which are then picked up by a chain node, and following the chain, arrive at TCC. Likewise, packets sent by TCC follow the chain in the opposite direction and are picked up by a passing train. Note that a TCC is a stationary server machine connected to a node at the end of the chain using a wired connection, typically through the wayside backbone network. Therefore, it does not use radio communication. Thus, the results discussed in this paper for a terminal node are in the context of the train terminal node.

While a chain node transmits only on two frequencies (in one direction each), a train transmits on all three frequencies for the following reasons. A train shall be able to communicate to the chain regardless of what direction or position the train is travelling relative to the chain. For example, let's suppose that the train transmitted on only one frequency. It might lead to a situation where the train meets a node whose respective antenna is a side antenna that faces opposite and thus cannot hear the train's transmission as it is a directional antenna. This is shown in Fig. 4 where Node 1 is not able to receive transmissions from the train. For packets flowing from train to the chain, it is not a problem as Node 2 is still able to receive from the train. However, note that the top radio is not used for transmitting but only receiving. Thus, for traffic flowing in the opposite direction, train will not be able to receive from either of the two nodes.
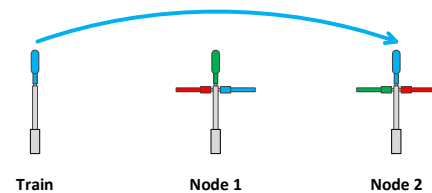

Fig. 4. A possible scenario if train transmitted on only one radio

Thus, the design requires that the train must transmit on minimum two frequencies, as in this way, a chain node with
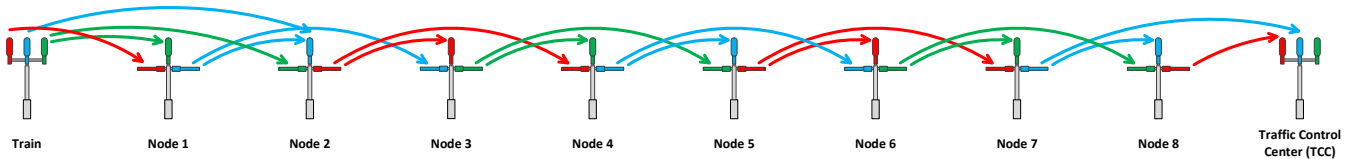
Fig. 5. A network of ten nodes with two terminal nodes, eight chain nodes, and a one-direction traffic flow (left to right)

any of the three possible frequency combinations will be able to receive from train on minimum one frequency. Nonetheless, to maximize the availability, the train transmits on all three frequencies. For the same reason, in contrast to a chain node, all three antennas on a train are omni-directional.

Fig. 5 shows a network where a terminal node (train) transfers packets to another terminal node (TCC) over a chain of eight nodes. Note that in real world, the train will travel along the chain, but for simplicity, here it is assumed that the train is located at the end of the chain. Likewise, for simplicity, a node identical to the train has been used to illustrate a TCC. Alternatively, it can be seen as a train transmitting to another train over a chain. This minor detail is not of any relevance when we discuss the results. Note how the three-frequency design ensures a frequency separation distance of 3 nodes (1800 meters), e.g. the red frequency is used by Node 2, Node 5 and Node 8.

C. Addressing and forwarding

There are two types of destination addresses involved in the design: the actual destination address—which will always be of a terminal node—and an address that indicates the direction of the packet flow. This so-called "direction address" is added to each packet and is one of "left", "right" and "both". When a train transmits, it uses "both" as the direction address. A chain node, upon receiving a packet from a train, replaces the direction in the packet with "left" or "right" and forwards the packet in both left and right directions. Thus, one of these packets is transmitted in the backward direction, a mechanism referred to as "backward forwarding" subsequently in this paper. Note that to ease the installation and maintenance efforts, a chain node must be deployed in a way that it is unaware of its location in the chain. Thus, forwarding the packet in both directions ensures that it takes the shortest path to TCC. The following chain node (in each direction) upon receiving this packet continues to forward it in one direction.

D. Duplicate packets

An inherent result of the redundancy in the design is the duplicate packets, which are both a requirement and a problem. Specifically, if each node forwards the duplicate packets, they quickly grow exponentially along the chain and congest the network. For example in Fig. 5, Node 1 will receive two copies of the same packet from the train. Next, Node 2 will receive four copies of the same packet, two forwarded by Node 1 and two received directly from the train. Likewise, Node 3 will receive six copies of the same packet, and so on. Therefore, duplicates are eliminated at each node with the help of a

unique sequence number included in each packet. Two types of duplicate packets exist:

1) Type 1: A node receives multiple copies of the same packet from two different nodes.
2) Type 2: A node receives multiple copies of the same packet from the same node.

In Fig. 5, example of duplicate type 1 is when Node 3 receives one copy of a packet from Node 1 (blue) and another from Node 2 (red). Example of duplicate type 2 is when Node 1 receives two copies from the train on its left (red) and top (green) radios. However, note that for type 2, the sending node is always a train, because only a train could be heard by the same chain node on more than one frequency. Node 2 in Fig. 5 receives duplicates of both types 1 and 2. It receives two copies from the train directly and one copy from Node 1.

Note that even though Fig. 5 shows TCC placed directly next to the chain, in reality, as described in Section IV-B, a TCC will be connected to the chain typically through the wayside backbone network. Thus, in the event that the two ends of the chain are connected to the same TCC—as shown in Fig. 2 (b)— the intermediate server machines in the wayside infrastructure will perform duplicate handling if two copies of the same packet are received following the two paths, before a packet is delivered to the TCC.

V. FIELD EXPERIMENT

The hardware node used in the field experiment was mainly based on ALIX 2D2 mainboards, with 500 MHz AMD Geode CPU and 2 mini PCI slots each. Each mainboard was equipped with an Atheros AR5414A WLAN radio card, which used IEEE 802.11a and operated at 54 Mbps data transmission rate. Three of these mainboards were mounted on a custom-made wooden mast and were connected together via Ethernet. Each board ran Linux (Debian Wheezy). Due to the limited availability of hardware, mostly old and leftover hardware was used for this purpose. HUBER-SUHNER Sencity SPOT-S antenna, operating in frequency band 5.150-5.875 GHz was used. The node mast is shown in Fig. 6.

The software component for the node model were written using Click Modular Router [24]—a framework for building configurable software-based routers—which facilitated in receiving, manipulating and forwarding packets. The sequence number contained in the Encapsulating Security Payload (ESP) header of IPSec was used for identifying duplicates and lost packets. Pre-defined MAC addresses were used to designate the directions "left", "right" and "both" discussed in Section IV-C and the address field in the IEEE 802.11 MAC header was used to hold this address. Note that since only broadcast transmissions are used in this design, the MAC addresses were
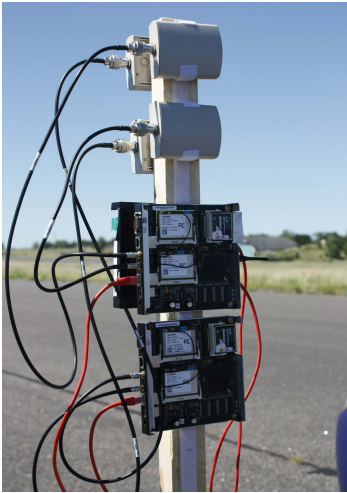
Fig. 6. Hardware node prototype

only used for determining the direction and not the destination. The address field in the IP header was used to specify the actual destination address, i.e. of a terminal node. By using these existing fields, the intention was to avoid a need for implementing a new protocol.

The experiment was carried out at an abandoned military airfield Flyvestation Værløse, formerly used by the Danish air force. The setup consisted of seven nodes. The nodes were placed 400 meters apart. The reasons for choosing this inter-node distance were that the radio hardware was not powerful enough to transmit signals at larger distances, and, it was the largest feasible distance to fit the whole chain on the airfield runway. Of these seven nodes, the nodes at the two ends of the chain—i.e. 1st and 7th nodes—were terminal nodes. The setup is shown in Fig. 7. The final test run was conducted over a period of 12 hours. Table I lists parameters and their values used in the experiment.

Fig. 7. Field experiment with a seven node chain

A. Results and discussions

In the test, one terminal node transmitted packets which were then transferred to the other terminal node over the chain. Note that for simplicity, in all results presented in this paper, the two terminal nodes are referred to as "train" and "TCC",

and in the one flow scenarios—like this one—it is assumed that packets flow from train to TCC. Nonetheless, regardless of the names, the results are equally applicable to both types of terminal nodes. Data rates ranging up to 8 Mbps were used.

TABLE I
FIELD EXPERIMENT PARAMETERS

Parameters	Value
WLAN technology	IEEE 802.11a OFDM at 54 Mbps
Frequency channels (MHz)	5735, 5800, 5865
Transmission power (dBm)	11.5
Receiver sensitivity (dBm)	-74
Antenna gain (dBi)	14
Packet size (bytes)	1000
Inter-node distance (m)	400
Nodes	7
Packet rate (per second)	1000
Run time (s)	40

As shown in Fig. 8, the results showed that for packet rates of 100 to 400 packets per second (800 kbps to 3.2 Mbps), 97.4 to 99.2 percent of packets were successfully delivered to TCC. However, the packet loss increased sharply at higher packet rates. As seen, at the rate of 1000 packets per second (8 Mbps), it rose to 36.69%. The unusually high number for Node 3 was partly due to a faulty radio on Node 2. This is visible in Fig. 9 that shows number of duplicate packets received at each node (a 100% means a duplicate of each packet was received). The large drop for Node 3 is because it only received the first copy of a packet (i.e. from Node 1) and not the redundant copy it was supposed to receive from Node 2. This faulty radio contributed significantly to the high packet loss seen at TCC.
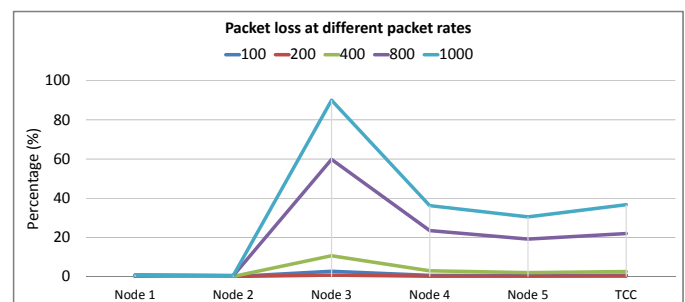

Fig. 8. Packets lost at each node

The field experiment successfully demonstrated the prototype of the design. The results showed that the chain network successfully transferred packets from one end to the other, as long as a too high packet rate was not used. It shall be noted that the field experiment suffered from a number of limitations. In addition to the limitation on the available hardware, time constraints imposed further challenges as one test run took several hours. Likewise, it was a tedious job to work with the node masts, due to the large distance between them, and, to collect data from all 7x3 mainboards which involved accessing them physically with a serial interface.

At a later point, an indoor demonstration was performed with six nodes and live video traffic was successfully transferred from one end to the other at a data rate of 5 Mbps. The objective mainly was to demonstrate that the development of the prototype was completed and that the prototype, particularly its software component, performed as intended. To reduce the signal loss and insulate the transmissions from any radio propagation effects, in this test, the antennas on the radios operating on the same frequency on the respective nodes were connected together with coaxial cables.
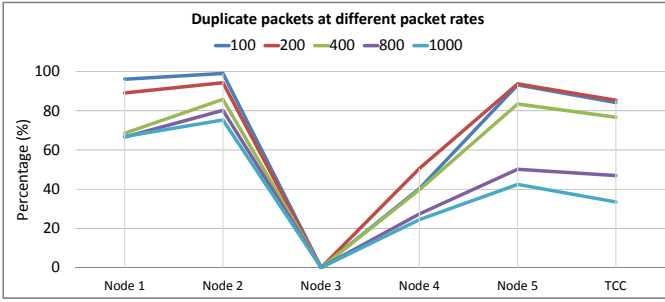

Fig. 9. Duplicate packets received at each node

VI. SIMULATION EXPERIMENT

OPNET Modeler 17.1 [25] was used for the simulation study. The simulation node model consists of three WLAN MAC models, each based on a modified version of the built-in OPNET WLAN model. In the real world, it will correspond to a node with three radio cards installed on it. Fig. 10 illustrates the node model.
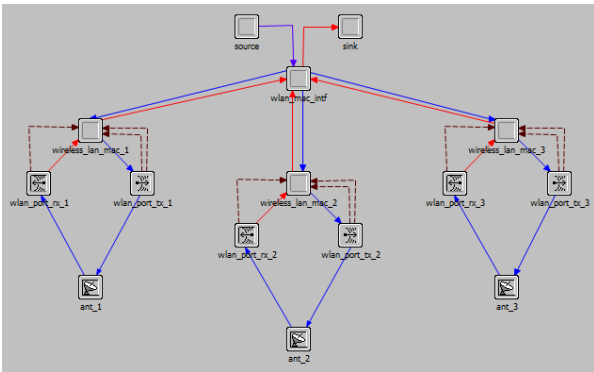

Fig. 10. Node model in OPNET

Table II lists generic simulation parameters and their values used in the simulations. Note that the inter-node distance depends on the node's signal range which in turn depends on various factors including transmission power, gain and receiver sensitivity of the antenna [1]. Nevertheless, an inter-node distance of 600 meters has been used in all simulations, as it could be directly related to the distance currently used in Copenhagen S-train based on the conventional CBTC technology. Thus, transmission power and receiver sensitivity were adjusted to transmit to a distance of 1200 meters. Likewise, a typical CBTC packet size of 512 bytes has been used.

TABLE II
GENERIC SIMULATION PARAMETERS

Parameters	Value
WLAN technology	IEEE 802.11a OFDM at 54 Mbps
Frequency channels (MHz)	5170, 5230, 5290
Transmission power (dBm)	7
Receiver sensitivity (dBm)	-76
Antenna gain (dBi)	14
Packet size (bytes)	512
Inter-node distance (m)	600
Nodes	100
Packet rate (per second)	1000
Simulation time (s)	60

The proposed design relies on the assumption that the separation provided by the three-frequency design is sufficient and signals from nodes beyond that distance will not interfere. However, in reality as relatively insignificant changes in the propagation conditions have shown to dramatically increase the signal range in railway environments [1]. Our simulation model uses OPNET's default Free-Space Path Loss (FSPL) propagation model, which, with the exceptionally large signal range it enables, provides the worst case scenario necessary to validate the proposed design. Using a more complex propagation model will improve the results as signal power will diminish faster, resulting in lower interference.

Nonetheless, to imitate more realistic radio propagation conditions, we introduce a random error in the system in which 2% of packets are marked erroneous. This number is based on the results from the field experiment in which an average error rate of 1 to 2 percent was observed [3]. The purpose is to study how the redundancy in the design guarantees a high number of packet transfer rate despite these errors.

A network size of 100 nodes has been used in our simulations as this will more likely be the largest network size used in the actual CBTC deployments both in terms of the number of nodes and the actual length (about 60 kilometers). In actual deployments, the chain will be divided into much smaller segments as the network infrastructure available at train stations will be used to wire the two nearest nodes.

Note that no IEEE 802.11 MAC layer retransmissions are made in this design due to the reasons that the redundancy in the design already ensures that a packet is sent to two nodes, and, retransmissions will have a negative impact on the bandwidth. Additionally, as a node will receive two ACKs instead of one in reply to one packet, the IEEE 802.11 MAC will behave unexpectedly as there is no way of knowing who sent which ACK and exactly which packet was acknowledged.

A. Key performance indicators (KPIs)

While a vast number of parameters could be studied, we are particularly interested in six performance indicators:

1) Total packets received: The number of total valid packets received at a node including duplicate packets.
2) Unique packets received: The number of unique packets received excludes duplicate packets. It serves as our

key parameter as it indicates how many packets are successfully transferred over the network. Note that this number for a node is essentially equivalent to the number of packets forwarded by the node.

3) Duplicate packets received: As discussed in Section IV-D, a duplicate packet is when multiple copies of the same packet are received at a node, from the same node or from different nodes.

4) Number of packets lost: The number of packets that, out of the original packets sent, were not received at the receiving end, for example due to errors.

5) Erroneous packet: The number of erroneous packets received at a node, mainly due to interference.

6) Number of collisions: The number of collisions refers to the situation in which a node receives multiple transmissions with the power level above the minimum required power level (i.e. receiver sensitivity) at the same time.

Additionally, we are interested in end-to-end delay, which is the delay that incurs from the time the packet is transmitted by the sending terminal node to the point it is received at the receiving terminal node. The long distance a packet must travel might result into a large end-to-end delay and thus might impact the timely delivery of the CBTC messages.

Note that while we discuss results for a select set of chain nodes as well, we are primarily interested in results for the terminal nodes. Additionally, note that as with the field experiment, in our simulations we prefer to use packet rate as a measure of data rate, as it facilitates in keeping track of the number of packets received at a node.

B. Scenario 1

In this scenario, one terminal node (train) transmits packets which are then transferred to the other terminal node (TCC) over the chain. A packet rate of 1000 packets per second—equivalent to 4.1 Mbps—is used.

Fig. 11 shows the results for the above six parameters against a select set of nodes displayed on x-axis. Note that the selection of nodes is not uniformly distributed. Specifically, first the five first nodes in the chain are listed—to highlight the interference the nodes near the train experience—and then every tenth node is listed. Y-axis shows the number of packets received in percentage. Note that at the rate of 1000 packets per second and the simulation time of 60 seconds, the number of packets sent by a single radio on the train during the whole simulation run is 60,000. Thus, a 100% unique packets received for a node on the figure implies that it received all 60,000 packets. Note that the total number of packets sent by the train is thrice this number, i.e. 180,000, and ideally the total number of packets received by a chain node is twice this number, i.e. 120,000.

The results show that 93.26% of the packets (red line) were successfully transferred to TCC, i.e. a packet loss of only 6.74% over a long chain of 100 nodes. As seen, the large and stable number of duplicate packets (dark blue line) received at each node highlights the effectiveness of the redundancy in the design. Furthermore, the frequency separation successfully minimizes interference as the number of erroneous packets is
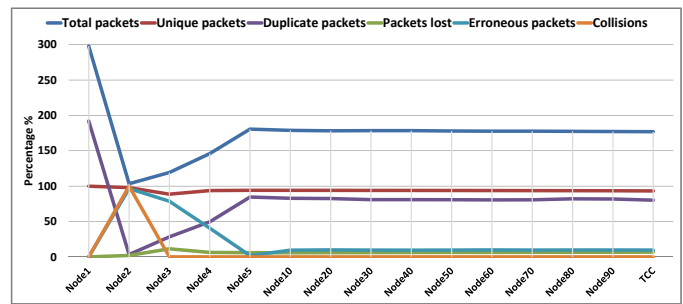

Fig. 11. Results for Scenario 1

minimum—except for at the first few nodes. As a result, only a negligible drop in the number of packets received—both total and unique—is seen at each subsequent node in the chain.

The results highlight a shortcoming of the design as well. As a train transmits on all frequencies in all directions in contrast to a chain node, the inherent frequency separation guaranteed otherwise in the chain is not fully achievable, resulting in interference on the nearby nodes. This is evident from the left part of Fig. 11 where a dramatic increase in the number of erroneous packets—and as a result a drop in the number of total and duplicate packets—is seen at Nodes 2 to 4.

At Node 2, train's transmissions result in collisions with those of Node 1. Note that Node 2 is the only node in this chain that is in the transmission range of two nodes transmitting on the same frequency, and thus the only node to experience collisions. As seen (orange line), nearly all transmissions from the train result in collisions on Node 2. While Nodes 3 and 4 are outside the transmission range of the train, they are still in the interference range. For example, at Node 3, train's transmissions interfere with those of Nodes 1 and 2. Nonetheless, as seen in Fig. 11, due to the redundancy in the design, only a minor drop in the number of unique packets received (red line) is seen at these nodes, except for Node 3. It is because while Nodes 2 and 4 suffer from interference primarily on only one of their radios, Node 3 does it on both of its radios.

Interference introduced by the train to Node 1's transmissions is particularly crucial. Due to the short distance between these two nodes, the insignificant difference in the received power of the two signals at Node 3 results in very low signal-to-noise ratio (SNR). Thus, 83% of the erroneous packets at Node 3 are received on its left radio. Notably, this phenomenon occurs only in the beginning of the chain where two nodes with a short distance between them (train and Node 1) transmit on the same frequency. Nonetheless, it is an important observation as it indicates that in case of an increased signal range due to improved propagation conditions, additional nodes in the beginning of the chain might face this problem. Likewise, the erroneous packets seen at Node 4, for example, are because the train is still in the interference range (it interferes with the transmissions of Node 2), although in this case the SNR will be comparatively higher.

Nonetheless, beyond this problematic initial part of the chain, i.e. as the interference from the train dies off, a stable number of packets received is seen at each node from Node 5

onward. Additionally, this implies that a network of a smaller size of e.g. 20 or 50 nodes would have fared the same.

As discussed in Section IV-C, the exceptionally high number of total packets received (i.e. above 200%) at Node 1 is because Nodes 2 and 3, upon receiving packets directly from the train, forward them in the backward direction as well, thus arriving back on Node 1.

Additional results—not presented here due to space constraints—showed a significantly low end-to-end delay of 2 milliseconds at TCC. Since transmissions are separated by frequencies, there can be at most one—for one flow—and two—for two flows one in each direction—nodes contending for the medium on one frequency at a given location in the chain. Thus, MAC contention delay as well as queueing delay are irrelevant. Notably, the IEEE CBTC standard [26] specifies a typical end-to-end delay of 500 milliseconds.

C. Scenario 2: A network with lower redundancy

The idea behind making the design redundant—i.e. by transmitting packets to two immediate neighbors in each direction—is to make it robust against random node failures. Failing a node and examining its impact on the network resiliency thus is an essential part of the evaluation. Thus in this scenario, first one node—out of the total 100 nodes—is purposely failed and its impact on the packet loss seen at TCC is studied. This is then repeated by increasing the number of failed nodes, one at a time, to a total of 10 nodes. Note that odd numbered nodes are failed, i.e. Node 1, 3, 5, and so on. Fig. 12 presents the results. Note that it shows the number of failed nodes on x-axis in contrast to the previous figures.
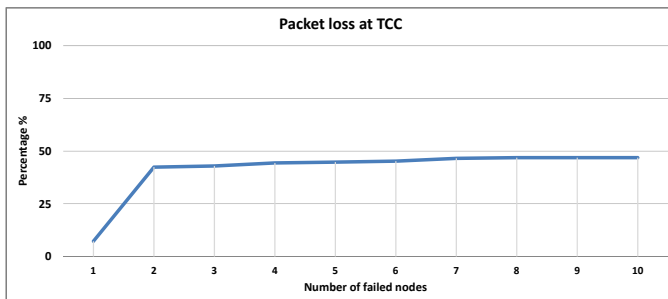

Fig. 12. Results for Scenario 2: Packet loss for a network with up to 10 failed nodes

The results show that when one node is failed (Node 1), the packet loss seen at TCC raises from the original 6.74% seen in Fig. 11 to 7.26%. The increase is insignificant due to the redundancy in the design, as train's transmissions are received not only by Node 1 but also Node 2. However, as one more node is failed next (Node 3), the packet loss increases sharply to 42.4% (i.e. an increase of approximately 35%). This is because as in Scenario 1 (Fig. 11), Node 4 receives around 40% erroneous packets from Node 2 due to the interference from the train. In Scenario 1, the redundant packets from Node 3 compensated for this. However, as Node 3 is not functional now, these erroneous packets result in lost packets, and this loss is not recovered throughout the chain. The shortcoming identified in Scenario 1 related to a train's transmissions thus

reappears here with a more pronounced impact. Nonetheless, after this point, as the interference from the train dies off, only a slight increase in the packet loss is seen at TCC as the number of failed node is increased incrementally to 10. A packet loss of only 0.93% is seen at each node on average.

Next, to present with the worst possible case, every second node in the chain is failed—i.e. 50 failed nodes. This essentially makes it a network with zero redundancy. The results are presented in Fig. 13. As expected, the number of duplicate packets has fallen to zero for all nodes—except for Node 2 that receives type 2 duplicate packets from the train. Thus, the number of total and unique packets has become equal for each node. As a consequence, a sharp drop in the number of packets received is seen at each subsequent node. Similarly, a sharp increase in the number of packets lost is seen, accumulating to 71.8% at TCC. However, as discussed above, the packet loss at Node 4 makes a large fraction of this number. Results presented later in Fig. 28 show how reduced interference from the train will lower this packet loss.
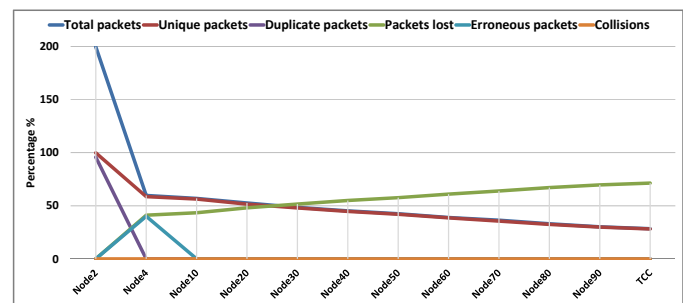

Fig. 13. Results for Scenario 2: A network without redundancy (i.e. every second node failed)

The results show that due to the redundancy in the design, the network sustains the failure of a remarkably large number of nodes (50 out of 100) as it still manages to transfer packets across the chain. Note that in a regular chain network, a single failed node can break the whole chain.

For the sake of comparison, Fig. 14 illustrates the number of unique packets received for the scenario with redundancy (Fig. 11) and the scenario without redundancy (Fig. 13). It emphasizes how the redundancy in the design ensures a stable number of packets received—on average 93.74% of packets—across the 100 nodes in the former, while in the latter, it sees a sharp drop. If the large drop seen at Node 4 due to the reasons discussed above is briefly ignored, the drop seen from Node 4 to TCC (a drop of approximately 30%) highlights how a network with no redundancy built into its design will perform.

D. Scenario 3: Traffic flows in two directions

In this scenario, Scenario 1 is extended with two flows, one in each direction. Note that this is equivalent to transmitting 8.2 Mbps. The results are presented in Fig. 15. Note that a significantly higher number of total and duplicate packets—on average 40.4% more packets, specifically—is received at each node in this scenario compared to the one flow scenario. These excessive packets are those flowing in the opposite direction.
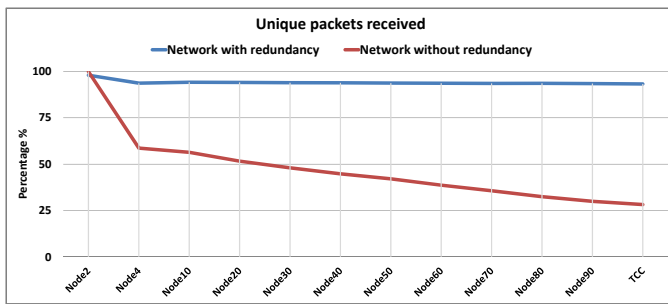
Fig. 14. Unique packets received: network with redundancy vs. network without redundancy

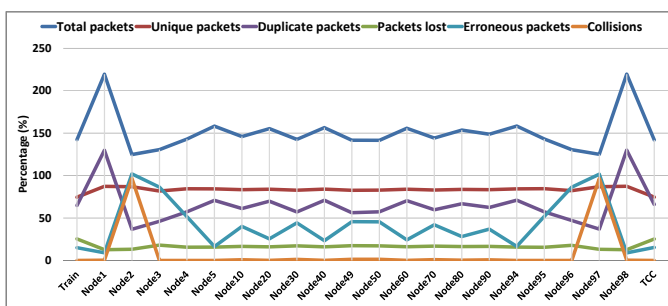Nonetheless, for brevity, the figure shows an average of these numbers for the two flows.

Fig. 15. Results for Scenario 3

As seen in Fig. 15, a stable number of unique packets received is maintained throughout the chain. Nonetheless, a significantly increased packet loss of on average 25.42% is seen at the two terminal nodes. Note that since traffic is flowing in both directions now, each top radio faces interference from nodes on its both sides. This includes collisions, which were limited to only Node 2 in the one flow scenario. This leads to increased interference as seen from the significantly higher number of erroneous packets at each node compared to Fig. 11. Specifically, on the middle nodes (Nodes 10 to 90), on average 17.21% erroneous packets are received per flow compared to 9.63% in the one flow scenario.

Note that the 25.42% packet loss over a large distance of 60 kilometers is still acceptable additionally due to the facts that (1) the CBTC traffic is redundant in nature as the train control information sent both ways is repeated at regular intervals, (2) the less favorable propagation conditions in the real word will lower the transmission range and thus the interference, and, (3) in real world, Quality of Service (QoS) differentiation will be applied to ensure that the CBTC traffic experiences lower packet loss. Nonetheless, the results imply that if exceptionally favorable propagation conditions are assumed, distant nodes might still be able to interfere despite the frequency separation, i.e. the interference range becomes larger than the 2-node transmission range and thus exceeds the 3-node frequency separation distance discussed in IV-B.

Additionally notable in Fig. 15 is the increased number of erroneous packets received at the two terminal nodes compared to Fig. 11. Note that the packets forwarded backward by e.g.

Nodes 1 and 2 upon receiving from the train, prove to be a major source of interference for the packets of the other flow, i.e. flowing from TCC to the train.

E. Scenario 4: Lower transmission power on the train

Two shortcomings of the design have been identified from the results discussed above.

1) A train's transmissions cause interference on the nearby nodes, as contrary to a chain node, a train transmits on all three frequencies in all directions, thus undermining the frequency separation. The impact of this shortcoming is seen on the nodes near the train.
2) In favorable propagation conditions, transmissions from nodes outside the transmission range still manage to interfere beyond the frequency separation distance. The impact of this shortcoming is seen throughout the chain.

A number of potential solutions to minimize the above two shortcomings are thus studied in this and the subsequent scenarios. Note that the solutions are studied in isolation to each other, to clearly identify the improvements that each solution can offer.

Note that the first of the two shortcoming is not particularly problematic if 100% redundancy is ensured in the chain, as seen in Scenario 1 (Fig. 11). However, the problem intensifies when there are failed nodes in the network, as discussed above in the results for Scenario 2 (Fig. 12 and Fig. 13). The problem is particularly critical due to the fact that as a train will be travelling along the chain, the effects currently seen on the first few nodes in the chain will be seen across each node in the chain. Multiple trains in close proximity will further worsen the situation. Thus, potential solutions for this shortcoming are studied first.

One potential solution is to use a different, lower transmission power on the train compared to that on a chain node. A lower power will reduce the interference as the transmissions will reach fewer nodes. Thus, Scenario 1 is repeated but this time with lower power values of 4, 5 and 6 dBm, in contrast to the original 7 dBm. Fig. 16 compares the number of erroneous packets received at each node for all four power values. It shows that as expected, compared to 7 dBm, the lower power values result in reduced interference on the nearby nodes. Likewise, as expected, the lowest power value (4 dBm) results in the lowest interference.
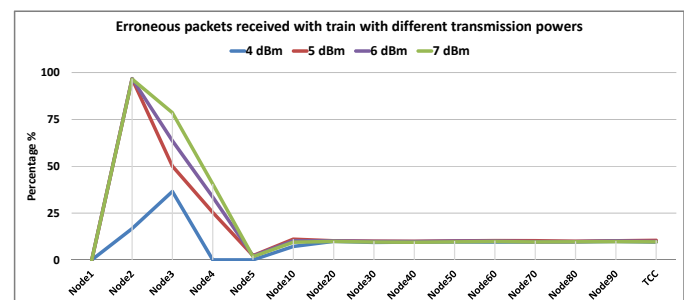

Fig. 16. Results for Scenario 4: Erroneous packets received with different transmission powers on the train

The results for the 5 dBm transmission power—which yielded the lowest packet loss at TCC—are presented in Fig. 17. Compared to Fig. 11, an increase in the number of total and duplicate packets received is seen in the figure for Nodes 3-4 as a result of lower interference from the train. As a result, the packet loss at TCC has decreased from the original 6.74% in Fig. 11 to 2.18%.
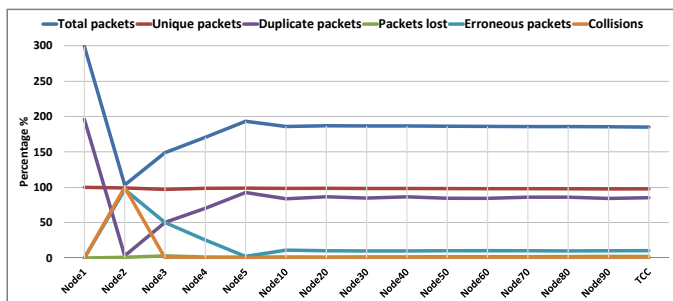

Fig. 17. Results for Scenario 4: Packet loss with 5 dBm transmission power on the train

The results for the 4 dBm transmission power—not presented here due to space constraints—showed that train's transmissions cannot reach Node 2 with this power. Nonetheless, they showed a packet loss of 7.99% at TCC, which is only slightly higher than the original 6.74% seen in Fig. 11. While 4 dBm transmission power further lowers the interference, the slight increase in the packet loss is due to the lowered redundancy as Node 2 is not able to receive redundant copies of the packets from the train. Nonetheless, as discussed subsequently, this implies that a design with a 1-node transmission range for the train might as well be feasible.

F. Scenario 5: Fewer radios on the train

As discussed in Section IV-B, as a train approaches a chain node, at most only two radios on the node will be able to receive train's transmissions because the third radio will be facing opposite. Thus, the design requires that a train must transmit on minimum two frequencies to ensure that a chain node is able to receive from it. Nonetheless, to further maximize the availability, a train transmits on all three frequencies. However, now that the problem caused by a train's transmissions is evident, in this scenario, the number of radios (or frequencies) on the train is reduced first to two and then to one. Note that chain nodes still use three radios.

1) Train with two radios: First, the feasibility of using two radios on the train is studied. As shown in Fig. 18, two frequencies on a train node will result in three different combinations relative to the frequencies on the nearest two chain nodes. The figure shows the combinations in the order—from most favorable to least favorable—as per the probability with which the chain nodes, preferably the first chain node, can receive transmissions from the train. As seen, for all three combinations, both chain nodes (Node 1 and Node 2) will be able to receive transmissions from the train. For frequency combinations 1 and 3, a total of three radios on the two nodes will be able to receive, in contrast to only two radios for the frequency combination 2. In the opposite direction, i.e.

traffic flowing from the chain nodes to the train, in frequency combination 1, the train will be able to receive from two nodes, in contrast to one for frequency combinations 1 and 2.
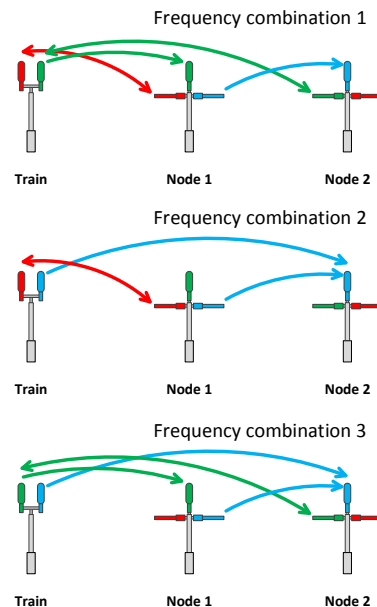

Fig. 18. Train with two radios and the three frequency combinations

Fig. 19 and Fig. 20 show the number of duplicate and erroneous packets received at each node, respectively, for the three frequency combinations.
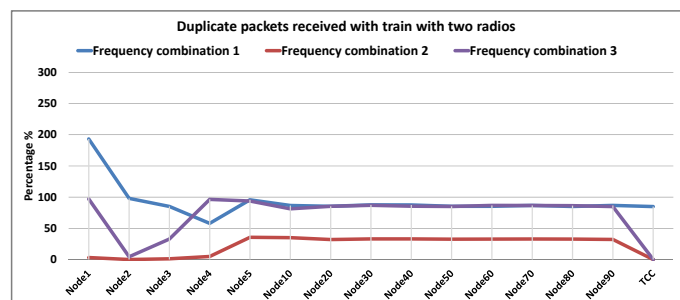

Fig. 19. Train with two radios: Duplicate packets received for the three frequency combinations
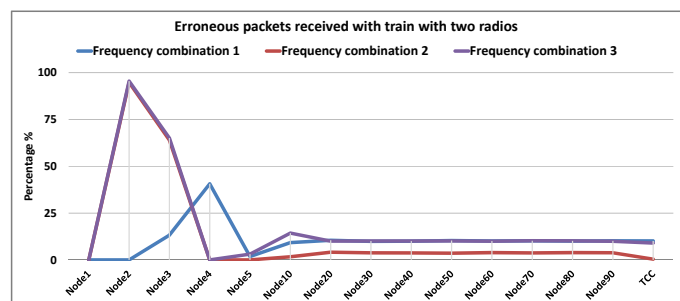

Fig. 20. Train with two radios: Erroneous packets received for the three frequency combinations

As seen in Fig. 19, frequency combination 1 results in the highest number of duplicate packets received at each node. As

seen in Fig. 20, a major contributor is the significantly lower interference at the nodes near the train compared to Fig. 11, particularly because the top radio on Node 2 does not receive transmissions from the train and thus does not experience interference (collisions) from the transmissions from Node 1. Fig. 19 shows that frequency combination 2 results in the lowest number of duplicate packets received at each node. As seen from the large number of erroneous packets at Node 2 in Fig. 20, it is due to the fact that nearly all (100%) transmissions from the train received at the top radio of Node 2 experience interference from the transmissions from Node 1. Since Node 2 does not receive any redundant copies of the packets from the train (i.e. on its right radio), these erroneous packets result in lost packets. Fig. 19 further shows that no duplicate packets are received at TCC for frequency combinations 2 and 3, as discussed above in the context of Fig. 18. Thus, it is noted that while a design with a reduced number of radios on the train decreases the interference at the nearby nodes, it decreases the ability of receiving duplicates for a train at the receiving end as well. Overall, from Fig. 19 and Fig. 20, it is concluded that frequency combination 1 is the most favorable and frequency combination 2, rather than 3, is the least favorable.

Fig. 21 presents the complete results for the most favorable frequency combination 1 for the one flow scenario, for a comparison with Fig. 11. It shows that the packet loss at TCC drops to 1.36% from the original 6.74%. Results for the equivalent two direction flow scenario—not presented here due to space constraints—showed a drop in packet loss from 25.42% seen in Fig. 15 to 21.57%.
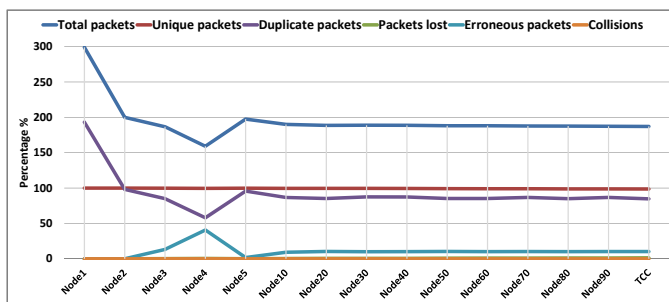

Fig. 21. Results for Scenario 5: Train with two radios and the most favorable frequency combination (combination 1)

Fig. 22 presents the complete results for the least favorable frequency combination 2 for the two direction flow scenario, for a comparison with Fig. 15. As discussed above, due to the interference experienced at the second node at each side (Node 2 and Node 97), and, a lack of duplicate packets received due to the reduced number of radios on the terminal nodes, a substantially high packet loss of 67.36% is seen at the terminal nodes compared to the original 25.42%. Thus, continuing from our discussion above, it is concluded that when the top radio is the only radio on a node that receives transmissions, in the absence of redundant packets, the impact of the interference from the transmissions of the preceding node increases dramatically.

The results show that the performance of the network depends on which two frequencies are used on the train.
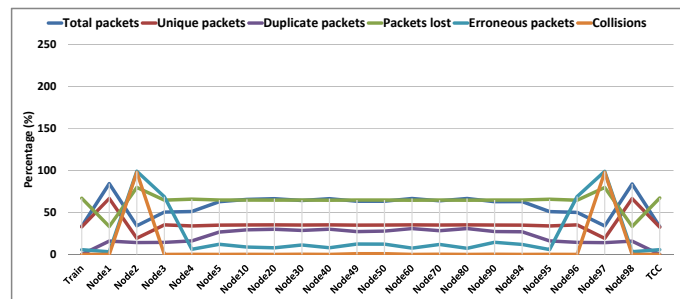

Fig. 22. Results for Scenario 5: Train with two radios and the least favorable frequency combination (combination 2), with two direction flow

Nonetheless, on the other hand, note that a train will be moving along the chain and thus will regularly encounter nodes with the favorable frequency combinations.

2) Train with one radio: As discussed in Section IV-B, the design requires that the train transmits on minimum two frequencies. Nonetheless, in this scenario, results for a train node with only one radio are presented to highlight the impact of using three frequencies on the interference experienced at the nodes near the train. Fig. 23 illustrates the three possible frequency combinations for a train with one radio. Note that in all combinations, at least one of the two chain nodes will be able to receive transmissions from the train. Combination 2 is particularly favorable as here both Node 1 and Node 2 will be able to receive from the train. For the same reason, it is the only combination in which Node 2 will receive duplicates. In the opposite direction, however, the train will be able to receive transmissions only in the first two combinations. This highlights the fact that while a train with one radio will be able to send packets to the chain, it might not always be able to receive from the chain.
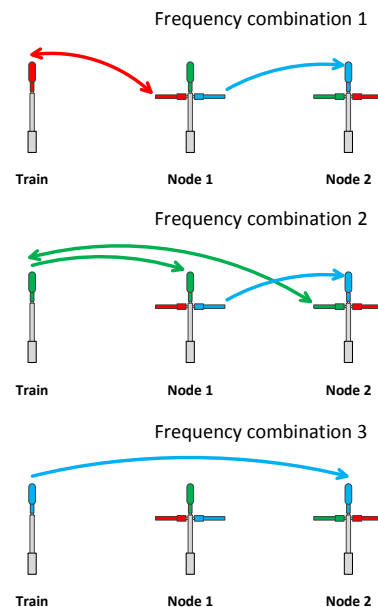

Fig. 23. Train with one radio and the three frequency combinations

Fig. 24 compares the number of duplicate packets received at each node for the three frequency combinations. It shows

that frequency combination 3 is the least favorable as no duplicate packets are received at any node before Node 4. This is because this is the only frequency combination where Node 1 does not receive train's transmissions, and thus does not send duplicate copies of packets. Note that the duplicate packets seen at Node 1 for frequency combination 2 are due to the backward forwarded packets from Node 2. Furthermore, no duplicate packets are received at TCC for any of the combinations as it has only one radio.
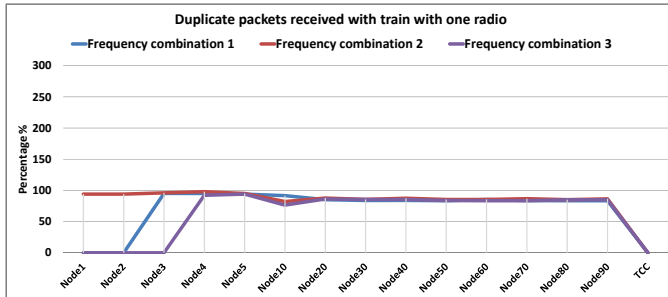

Fig. 24. Train with one radio: Duplicate packets received for the three frequency combinations

Fig. 25 shows the number of erroneous packets received for the three combinations. Note that due to the small numbers, the figure uses a smaller scale on y-axis. As expected, the erroneous packets seen on the first four nodes in Fig. 11 have disappeared as a result of the lower interference from the train.
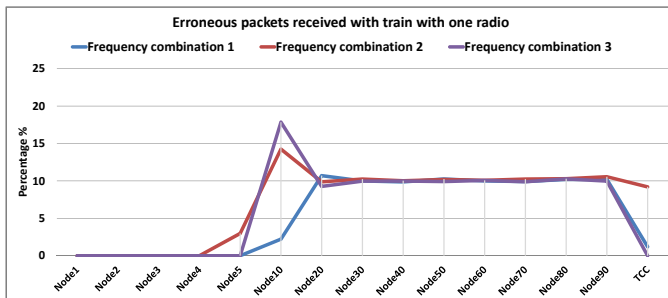

Fig. 25. Train with one radio: Erroneous packets received for the three frequency combinations

Fig. 26 presents the complete results for the most favorable frequency combination 1, for a comparison with Fig. 11, and shows that the packet loss at TCC drops to 5.27% from the original 6.74% due to the reduced interference at the nodes near the train. Results for the equivalent two direction flow scenario—not presented here due to space constraints— showed a drop in packet loss from 25.42% seen in Fig. 15 to 24.78%.

Note that the interference seen throughout the chain caused by the transmissions from the nodes outside the frequency separation distance (i.e. shortcoming 2) is comparatively a larger contributor of the high packet loss than the interference caused by a train's transmissions (i.e. shortcoming 1). Furthermore, as discussed above, reducing the number of radios to one also means that a terminal node is not able to receive any duplicates. Therefore, reducing the number of radios on the
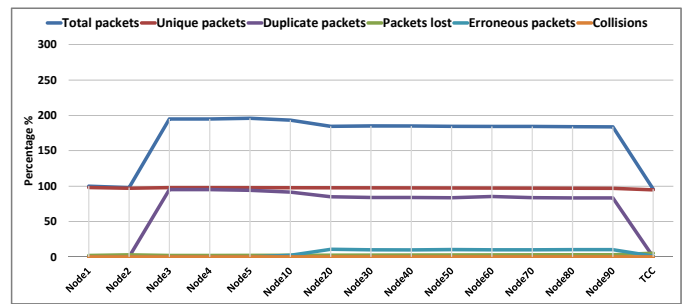

Fig. 26. Results for Scenario 5: Train with one radio and the most favorable frequency combination (combination 1)

terminal node alone does not offer significant improvements to the packet loss.

Fig. 27 presents results for the least favorable frequency combination 3, for a comparison with Fig. 15. As discussed above, no packets are received at either of the terminal nodes. Nevertheless, note that the chain is still able to transfer packets as seen from the average packet loss of 23.97% at the two second last nodes at the two ends (i.e. Node 1 and Node 98).
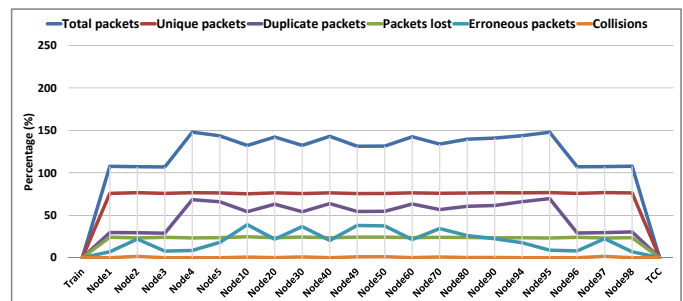

Fig. 27. Results for Scenario 5: Train with one radio and the least favorable frequency combination (combination 3)

In the results for Scenario 2, in which every second node in the chain was purposely failed (Fig. 13), a packet loss of 71.8% was seen at TCC, mainly as a result of the 40% occurred already at Node 4 due to the interference from the train. For the sake of comparison, Fig. 28 shows the packet loss seen in the original scenario together with that seen when a 5 dBm transmission power is used on the train (as discussed in Scenario 4), and, when one radio is used on the train (as discussed in Scenario 5). It shows that when a lower transmission power is used on the train, the packet loss at TCC drops to 63.9%—as the number of erroneous packets at Node 4 drops to 25%. When one radio is used on the train, the results show that the packet loss drops to 54.88%—as the number of erroneous packets at Node 4 drops to only 3.89%—which implies that about 17% packet loss seen at TCC in Fig. 13 is a direct result of the interference received at a single node (Node 4) due to train's transmissions. Thus, it is concluded again that the impact of interference caused by a train is more pronounced in a network with lower redundancy.

The results assert that a train node does not necessarily have to have three radios. A design in which the train has two radios, or even one, is equally feasible. This is again with a note that a train will be moving along the chain and
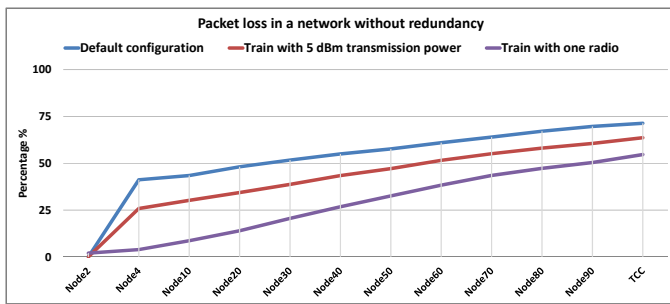
Fig. 28. Packet loss in a network without redundancy, with different train configurations

will nonetheless encounter a node with a favorable frequency combination two out of three times. However, for a standing train, the case of one radio and frequency combination 3 (Fig. 27) presents a problem as the train will not able to receive packets from the chain, unless it is in the range of multiple chain nodes. An alternative approach can be to still have three radios on the train such that all three are used for receiving but only one for transmitting. Another workaround could be to deploy nodes of all three frequency combinations on train stations, or maybe a specialized chain node with three omni-directional antennas like on a train in the original design.

G. Scenario 6: A more robust modulation and coding scheme

In this scenario, a more robust modulation and coding scheme is used for transmissions as it helps in coping with the impact of interference. Note that the solution is studied in isolation from the solutions in the previous scenarios, i.e. default transmission power and number of radios on the train is used. First the coding rate of 2/3, yielding a data transmission rate of 48 Mbps—instead of the coding rate of 3/4 for 54 Mbps—is used. Note that as seen in Table III, both transmission rates use the same modulation scheme, namely 64-QAM (Quadrature Amplitude Modulation). The results—not presented here due to space constraints—showed a dramatic decrease in the number of erroneous packets, both near the terminal nodes and throughout the chain, and as a result a drop in the packet loss from 25.42% seen in Fig. 15 to 14.39%.

TABLE III
A COMPARISON OF DATA RATES SUPPORTED BY DIFFERENT IEEE 802.11 TECHNOLOGIES PER MODULATION AND CODING SCHEME

Modulation scheme	Coding rate	Data rate (Mbps)		
		802.11a	802.11n	802.11ac
16-QAM	1/2	24	240	780
16-QAM	3/4	36	360	1170
64-QAM	2/3	48	480	1560
64-QAM	3/4	54	540	1755

Next, the most robust, QAM-based modulation and coding scheme, namely 16-QAM 1/2, is used. The results are presented in Fig. 29. The results show a further drop in the number of erroneous packets throughout the chain, as is evi-

dent from the "smoothed out" lines for the total and duplicate packets received, resulting in a packet loss of 13.07%.
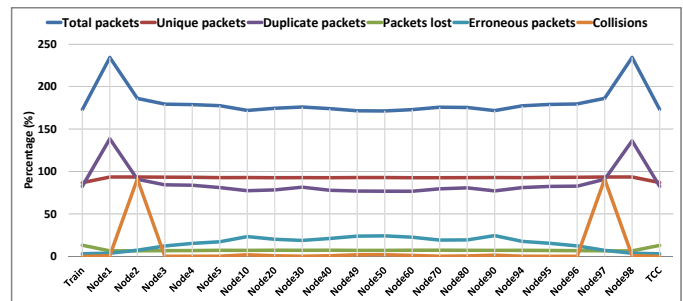

Fig. 29. Results for Scenario 6: Transmissions with the 16-QAM 1/2 modulation and coding scheme

Note that our experiments were carried out with IEEE 802.11a as (1) it is the technology officially approved by the company for the hardware for its CBTC project, and, (2) the available hardware that could be used for the experiment was limited as discussed in Section V. As seen in Table III, more advanced IEEE 802.11 technologies such as 802.11n and 802.11ac support substantially high data rates with the same two combinations of the modulation and coding schemes, specifically, data rates of up to 480 and 1560, and, 240 and 780 Mbps, respectively.

H. Scenario 7: Combined adjustment of multiple parameters

In this scenario, the impact of using various combinations of modulation and coding scheme, receiver sensitivity, and train's transmission power, are studied. Fig. 30 presents results for a combination of train's transmission power 5 dBm and receiver sensitivity of -75 dBm. As seen, the packet loss drops further to 12.73%, the lowest seen in our results yet.

As seen from the drops in the number of total and duplicate packets received at Node 2 and Node 97 compared to Fig. 29, these parameter values reduce the transmission range of the terminal node to one node. For the same reason, no jumps in the number of total and duplicate packets received are seen at the preceding two nodes (Node 1 and Node 98) as Node 2 and Node 97 do not forward packets backward.
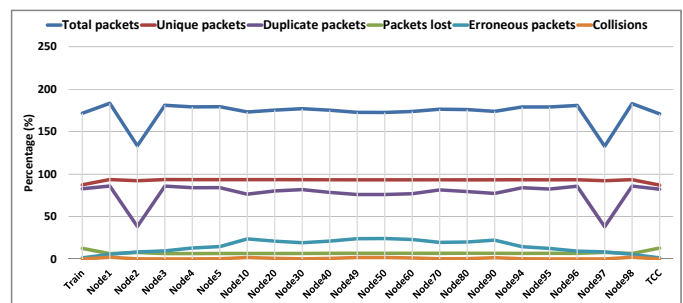

Fig. 30. Results for Scenario 6: Train with 1-node transmission range

Thus, together with the results discussed in Scenario 4 where a 4 dBm transmission power on the train resulted in a 1-node range, the results show that the presumption that a minimum 2-node range is required at the train is not valid. In

fact, a 1-node transmission range at the train together with the assurance that minimum two radios are able to receive from the train—i.e. two radios on the nearest node—will both result in lower interference on the nearby nodes and at the same time guarantee the redundancy aimed by the design.

VII. CONCLUSIONS

This paper presented a novel, ad-hoc based design for the trackside radio communication network for CBTC. A node in this design functions in ad-hoc mode, receiving broadcast packets and forwarding to its neighbors, thus forming a chain of nodes. The design offers redundancy and resiliency against node failures and interference by employing frequency separation, and transmitting to two neighbor nodes instead of one. As a result, the train does not have to associate with the nodes as it moves and the costly optical fiber cables connecting the nodes are no more needed. The paper presented two experimental studies to prove the feasibility of this design, a field study involving a hardware prototype, and an extensive simulation study in which a large network of 100 nodes (60 kilometers) was simulated. The results showed that compared to a regular network without redundancy, due to the redundancy inherent in this design, significantly large numbers of packets can be transferred across large networks with an acceptable packet loss. Reduced interference as a result of frequency separation further minimizes the packet loss. Likewise, end-to-end delay is minimum as frequency separation guarantees minimum contention for the wireless medium. The results exposed two limitations of the design as well. They showed that the terminal node (i.e. a train) undermines the frequency separation guaranteed by the chain nodes as it is required to transmit on all frequencies, and, the design under-estimates the interference produced by distant nodes in ideal propagation conditions despite the frequency separation. Nonetheless, the paper presented a number of potential solutions that minimize these limitations substantially by employing a lower transmission power on the train, a lower number of radios on the train, and more robust modulation and coding schemes.

ACKNOWLEDGMENT

REFERENCES

[1] J. Farooq and J. Soler, "Radio communication for communications-based train control (CBTC): A tutorial and survey," *IEEE Commun. Surveys Tuts.*, vol. 19, no. 3, pp. 1377–1402, 2017.

[2] Siemens AG, "Ad-hoc kommunikationsnetzwerk," Patent Application 10 2017 203 040.2, feb 24, 2017.

[3] R. T. Karstensen, "Reliable wireless infrastructure for bridging stationary and moving nodes," Master's thesis, Department of Photonics Engineering, Technical University of Denmark, 2015.

[4] J. Farooq, L. Bro, R. T. Karstensen, and J. Soler, "A multi-radio, multi-hop ad-hoc radio communication network for Communications-Based Train Control (CBTC)," in *Proc. IEEE 86th Vehicular Technology Conference (VTC 2017-Fall)*, accepted for publication.

[5] S. Xu and T. Saadawi, "Does the IEEE 802.11 MAC protocol work well in multihop wireless ad hoc networks?" *IEEE Commun. Mag.*, vol. 39, no. 6, pp. 130–137, jun 2001.

[6] J. Li, C. Blake, D. S. De Couto, H. I. Lee, and R. Morris, "Capacity of ad hoc wireless networks," in *Proc. 7th Annual International Conference on Mobile Computing and Networking (MobiCom '01)*. New York, NY, USA: ACM, 2001, pp. 61–69.

[7] K. Xu, M. Gerla, and S. Bae, "How effective is the IEEE 802.11 RTS/CTS handshake in ad hoc networks," in *Proc. IEEE Global Telecommunications Conference (GLOBECOM '02)*, nov 2002.

[8] F. Ye, S. Yi, and B. Sikdar, "Improving spatial reuse of IEEE 802.11 based ad hoc networks," in *Proc. IEEE Global Telecommunications Conference (GLOBECOM '03)*, dec 2003.

[9] G. Anastasi, E. Borgia, M. Conti, and E. Gregori, "IEEE 802.11 ad hoc networks: performance measurements," in *Proc. 23rd International Conference on Distributed Computing Systems Workshops*, may 2003.

[10] *Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications Amendment 6: Wireless Access in Vehicular Environments*, IEEE Std. 802.11p, 2010.

[11] L. Zhao, X. Hong, J. Zhang, Y. Zhang, and Q. Hao, "Feasibility analysis of multi-radio in DSRC vehicular networks," in *Proc. 16th International Symposium on Wireless Personal Multimedia Communications (WPMC)*, jun 2013.

[12] G. V. Rossi, K. K. Leung, and A. Gkelias, "Density-based optimal transmission for throughput enhancement in vehicular ad-hoc networks," in *Proc. IEEE International Conference on Communications (ICC)*, jun 2015.

[13] J.-H. Chu, K.-T. Feng, C.-N. Chuah, and C.-F. Liu, "Cognitive radio enabled multi-channel access for vehicular communications," in *Proc. IEEE 72nd Vehicular Technology Conference Fall (VTC 2010-Fall)*, sep 2010.

[14] M. Hadded, P. Muhlethaler, A. Laouiti, and L. A. Saidane, "TDMA-aware routing protocol for multi-hop communications in vehicular ad hoc networks," in *Proc. IEEE Wireless Communications and Networking Conference (WCNC)*, mar 2017.

[15] M. F. Feteiha and M. H. Ahmed, "Best-relay selection for multi-hop vehicular communication in highways," in *Proc. IEEE Global Communications Conference (GLOBECOM)*, dec 2015.

[16] M. J. Farooq, H. ElSawy, and M.-S. Alouini, "A stochastic geometry model for multi-hop highway vehicular communication," *IEEE Trans. Wireless Commun.*, vol. 15, no. 3, pp. 2276–2291, 2016.

[17] V. Shivaldova, T. Paulin, A. Paier, and C. F. Mecklenbräuker, "Performance measurements of multi-hop communications in vehicular ad hoc networks," in *Proc. IEEE International Conference on Communications (ICC)*, jun 2012.

[18] H. Rhee, "ATO data of train control system based on Wi-Fi mesh telecommunication," in *Proc. The 21st International Conference on Magnetically Levitated Systems and Linear Drives*, 2013.

[19] A. Srivatsa and J. Xie, "A performance study of mobile handoff delay in IEEE 802.11-based wireless mesh networks," in *Proc. IEEE International Conference on Communications (ICC '08)*, may 2008, pp. 2485–2489.

[20] A. Mishra, M. Shin, and W. Arbaugh, "An empirical analysis of the IEEE 802.11 MAC layer handoff process," *SIGCOMM Computer Communication Review*, vol. 33, no. 2, pp. 93–102, Apr. 2003.

[21] B. Bu, F. Yu, and T. Tang, "Performance improved methods for communication-based train control systems with random packet drops," *IEEE Trans. Intell. Transp. Syst.*, vol. 15, no. 3, pp. 1179–1192, jun 2014.

[22] L. Zhu, F. R. Yu, B. Ning, and T. Tang, "Handoff performance improvements in MIMO-enabled communication-based train control systems," *IEEE Trans. Intell. Transp. Syst.*, vol. 13, no. 2, pp. 582–593, jun 2012.

[23] H. Jiang, H. Zhao, and B. Zhao, "A novel handover scheme in wireless LAN in CBTC system," in *Proc. IEEE International Conference on Service Operations, Logistics, and Informatics (SOLI)*, jul 2011, pp. 473–477.

[24] E. Kohler, R. Morris, B. Chen, J. Jannotti, and M. F. Kaashoek, "The Click Modular Router," *ACM Transactions on Computer Systems*, vol. 18, no. 3, pp. 263–297, 2000.

[25] "OPNET Modeler," Riverbed Technology, https://www.riverbed.com/dk/products/steelcentral/opnet.html.

[26] *Communications-Based Train Control (CBTC) Performance and Functional Requirements*, IEEE Std. 1474.1, 2004.