Technical University of Denmark

DTU

A multi-radio, multi-hop ad-hoc radio communication network for Communications-Based Train Control (CBTC) with optimized frequency separation

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

DTU Library
Technical Information Center of Denmark

A multi-radio, multi-hop ad-hoc radio communication network for Communications-Based Train Control (CBTC) with optimized frequency separation

Jahanzeb Farooq[1,2], Lars Bro[3], Rasmus Thystrup Karstensen[1] and José Soler[2]
[1]Siemens A/S, Ballerup, Denmark
[2]DTU Fotonik, Technical University of Denmark, Lyngby, Denmark
[3]Nyantec UG, Berlin, Germany
jahanzeb.farooq@siemens.com, bro@nyantec.com, rasmus.karstensen@siemens.com, joss@fotonik.dtu.dk

Abstract—Communications-Based Train Control (CBTC) is a modern signalling system that uses radio communication to transfer train control information between train and wayside. The trackside networks in these systems are mostly based on conventional infrastructure Wi-Fi (IEEE 802.11). It means a train has to continuously associate (i.e. perform handshake) with the trackside Wi-Fi Access Points (AP) as it moves, which incurs communication delays. Additionally, these APs are connected to the wayside infrastructure via optical fiber cables that incur huge installation costs. Our earlier work presented a novel design in which trackside nodes function in ad-hoc Wi-Fi mode, which means no handshake has to be performed with them prior to transmitting. A node upon receiving packets from a train forwards these packets to the next node, forming a chain of nodes. Following this chain, packets reach the destination. To make the design resilient against interference between the nodes, transmissions are separated on multiple frequencies, ensuring a certain separation between the transmissions. Nonetheless, the results show that despite this separation, a significant amount of interference is experienced along the chain due to the interference range being greater than the frequency separation distance. This paper proposes an extension to the design in which additional frequencies are employed in an interleaving fashion to optimize the frequency separation distance and presents the results from an extensive simulation study.

Index Terms—Railway signalling, CBTC, radio communication, Wi-Fi, IEEE 802.11, ad-hoc, multi-radio, multi-hop

I. Introduction

Communications-Based Train Control (CBTC) is a widely popular 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. 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.

In these systems, hundreds of Wi-Fi Access Points (APs) are installed at the trackside to enable uninterrupted wireless connectivity. Each AP is connected to the wayside (normally a Traffic Control Center (TCC)) via optical fiber cables. Just like in an ordinary infrastructure Wi-Fi network, the train must first associate (i.e. perform handshake) to an AP to be able to transmit. However, there are a number of disadvantages of this design. Firstly, installation of cables to connect each AP to the wayside is time-consuming and incurs high costs. Secondly, the train must handover from one AP to other as the it moves. The IEEE 802.11 technology lacks the support for mobility as it was originally developed for users in stationary environments. This results in employment of complex handover algorithms in CBTC systems to enable seamless mobility. A completely seamless operation is still not feasible, leading to delays in communication as well as limitations on the maximum train speed.

In [1], we presented a novel design for an ad-hoc based radio communication network (patent pending [2]) in which there are no conventional "APs". Nodes function as ordinary Wi-Fi nodes, in an ad-hoc manner. A node broadcasts packets to any nodes within its range. As a nearby node receives the packet, it re-transmits (forwards) it, to be picked up by the next nearby node. This forms a chain of nodes. Following this chain, the packets reach the last node in the chain, which is typically connected to TCC. Thus, a train is not required to establish an association with an AP, and as a result, to handover between APs. Wired links between the nodes and wayside backbone are no longer needed except for the two nodes at each end of the chain. To make the chain resilient against failures, redundancy is introduced in a way that each node forwards packets to two of its neighbors in each direction instead of one.

In a conventional multi-hop ad-hoc network where all nodes operate on a single frequency, the capacity degrades sharply with the growing size of the network as a result of the increased interference as well as contention for the medium [3], [4], [5], [6], [7], [8], [9]. Thus, to make the chain resilient against interference, in the proposed design, transmissions are separated on multiple frequencies to introduce a certain separation between nodes transmitting on the same frequency.

The results from an extensive simulation study presented in [1] and in this paper verify the effectiveness of the design primarily in terms of resiliency, redundancy, and scalability. However, the results also highlight a limitation of the design. They show that despite the inherent frequency separation, the design under-estimates the interference produced by the nodes outside the frequency separation distance in ideal propagation conditions. As a result, a significantly high packet loss is seen as the packets are transferred across the chain. This paper investigates the problem further and proposes a mechanism to optimize the frequency separation distance by employing additional frequencies in an interleaving fashion. Additionally, it describes the proposed design in greater clarity.

The rest of this paper is laid out as follows. Section II presents a brief overview of CBTC systems. Section III provides an overview of the proposed design. Section V provides an overview of the extended design together with the simulation study and the results. Section VI discusses future work. Finally, Section VII concludes the paper.

II. OVERVIEW OF CBTC SYSTEMS

A brief overview of CBTC is presented here. For a more detailed version, refer to [10]. In CBTC, radio communication is used to exchange train control information between the train and the wayside, enabling Automatic Train Control (ATC) functions. The train regularly sends its state to the wayside over the radio connection. The state information includes the current speed, direction, and location of the train. Based on this information, the wayside ATC equipment calculates the "limit of movement authority" (LMA) information and sends it back to the train. LMA 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. Due to this real-time communication between train and wayside, the precise location of the trains can be determined. This enables the so-called "moving block operation" that allows trains to run closer to each other. Furthermore, the number of trackside equipment—such as color light signals and track circuits—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).

A large number of Wi-Fi APs are deployed at the trackside to guarantee that the train has a radio connection all the time. Each AP is connected (over a wired link) to the wayside components through the backbone network. A train has to continuously search for a new suitable AP (a process called *scanning*) and re-associate as it moves along. To assist in handover, APs are placed in a way that their coverage areas overlap. Fig. 1 uses the green and red colors to differentiate between the APs' coverage areas. A critical aspect of handover in CBTC is how the train smoothly switches from one AP to another, without causing interruptions in the communication.
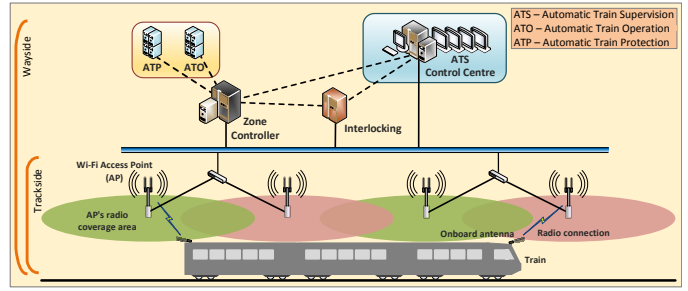

Fig. 1. CBTC wayside components

A large handover latency might result in a train failing to receive information about the minimum permitted distance to the preceding train in-time. 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 [10].

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 [10].

III. PROPOSED NETWORK DESIGN

A brief description of the proposed design—albeit revised to add more clarity—is presented here. Refer to [1] for details. Fig. 2 (a) illustrates the conventional network design for CBTC trackside. In the proposed design, 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, as illustrated in Fig. 2 (b). No AP scanning and association are thereby required.
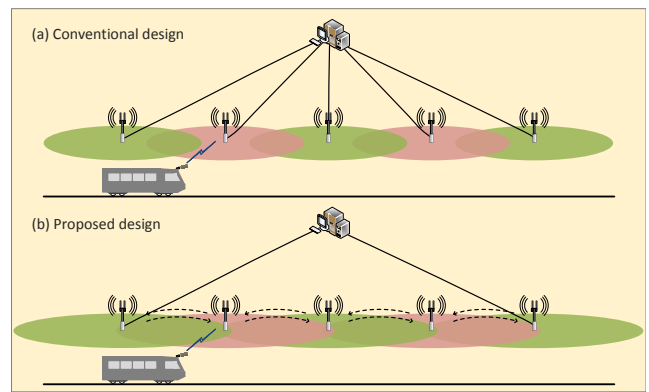
Fig. 2. CBTC trackside network: Conventional vs. proposed design

A. Frequency separation and redundancy

A conventional multi-hop ad-hoc network operating on a single frequency presents two major challenges. Firstly, as noted above, if all nodes transmit on the same frequency, the probability of interference rises sharply. Additionally relevant is the well-known "hidden node problem" in which two nodes are in the transmission range of a common node but not
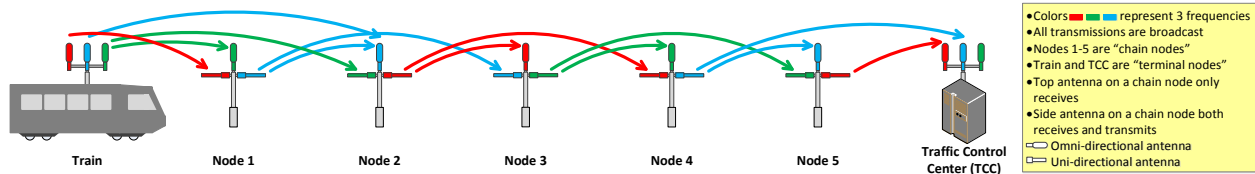
Fig. 3. A one-direction traffic flow where a terminal node (train) transfers packets to another terminal node (TCC) over a chain of five nodes

in each other's range. Since they cannot hear each other, it renders the Carrier Sense Multiple Access/Collision Avoidance (CSMA/CA) mechanism used in IEEE 802.11 MAC to avoid collisions ineffective. Secondly, a single failed node practically breaks the chain.

To solve the interference problem, the proposed design uses three frequencies to ensure a certain separation among nodes' transmissions. 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. 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 inside a node are connected to each other via Ethernet. 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.

Fig. 3 illustrates the mechanism, where the colors red, blue and green represent three frequencies. Note that the two transmission lines coming out of a radio—e.g. 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. 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.

B. Chain node vs. terminal node

The node type discussed above is referred to as a "chain node". A second type of node is the "terminal node", which is the actual sender/receiver of these packets and is either a train or a TCC. Basically, it is the node that uses the chain network to get its packets transferred to a terminal node at

the other end of the chain. A train sends packets to the TCC, and a TCC sends packets to one or more trains. Note that in reality, TCC is a stationary machine connected to the end of the chain using a wired connection. 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 node.

While a chain node transmits only on two frequencies (one in each direction), a train transmits on all three frequencies and in all directions, i.e. it uses three omni-directional antennas. This is necessary as a train shall be able to communicate to the chain regardless of what direction or position it is travelling relative to the chain. Note that one of the directional antennas on a chain node might be facing opposite and thus might not be able to receive from the train. Thus, by transmitting on all three frequencies, it is guaranteed that a chain node with any of the three possible frequency combinations is able to receive from train on a minimum of two frequencies. Fig. 3 shows a network where a terminal node (train) transfers packets to another terminal node (TCC) over a chain of five nodes. Note how the use of three frequencies ensures a frequency separation distance of three nodes, e.g. the red frequency is used by Node 2 and Node 5.

An inherent consequence of the redundancy in the design is the duplicate packets. Note that a node might receive multiple copies of the same packet either from the same node or from two different nodes. For example in Fig. 3, Node 1 will receive two copies of the same packet from the train. Node 2 will receive four copies of the same packet, two forwarded by Node 1 and two received directly from the train, and so on. If each node forwards the duplicate packets, they grow exponentially along the chain and congest the network. Therefore, duplicates are eliminated at each node based on a unique identifier.

IV. IMPROVED DESIGN WITH OPTIMIZED FREQUENCY SEPARATION

As concluded subsequently in Section V-A when discussing results for the original three-frequency design, despite the frequency separation distance enabled by the design, nodes beyond this distance still interfere due to their increased interference range in ideal propagation conditions. To minimize this interference, an improved design is proposed that extends the frequency separation distance by increasing the total number of frequencies used and employing them in an interleaving fashion at each node. A chain node still uses three frequencies but now each subsequent node uses a different set of three frequencies instead of repeating the same set. That is, at each

node, one of the existing three frequencies is replaced with a new one by iterating over the total number of frequencies.

Note that no modifications are required to be made to the chain node's equipment as it is still equipped with three radios. On the contrary, the train node is now required to be equipped with additional radios. Fig. 4 illustrates the mechanism.
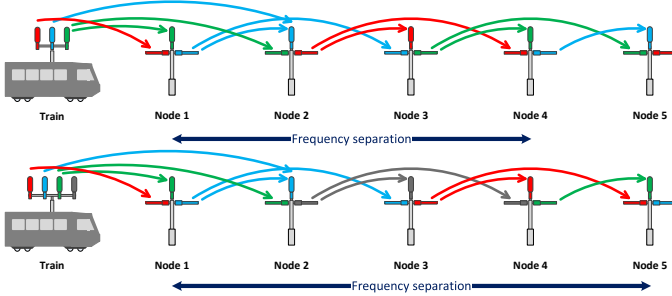

Fig. 4. Frequency separation guaranteed with 3 and 4 frequencies

The top part of Fig. 4 shows the original three-frequency design. Here, for example, the blue frequency is used by Node 1 and Node 4, i.e. a frequency separation distance of three nodes. The bottom part of the figure uses one additional frequency (black—used by Nodes 2, 3 and 4). As seen, this extends the frequency separation distance from three nodes to four nodes as the blue frequency is now repeated at Node 5 instead of Node 4. Increasing the number of frequencies to five and six further extends the frequency separation distance likewise.

A train node is required to be equipped with additional radios because, as discussed in III-B, it must use all frequencies to be able to transmit to the chain and receive from it. Note that it is normal to employ various—e.g. up to four—radios per train in conventional CBTC systems in order to ensure high availability [10]. Thus, this additional radio on the train does not necessarily increase the system's cost.

V. SIMULATION STUDY

Simulations were carried out using a discrete-event simulator [11]. Table I lists the key simulation parameters and their values used in the simulations. 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 the Copenhagen's S-train CBTC system based on the conventional CBTC technology. As the design requires that a node be heard by two of its neighbors, transmission power and receive sensitivity were adjusted to transmit to a distance of 1200 meters.

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, this is far from reality as minor changes in the propagation conditions have shown to dramatically increase the signal range in railway environments [10]. Our simulation model uses the simulation tool's default Free-Space Path Loss (FSPL) propagation model. The FSPL model assumes a free space between the sender and receiver and therefore does not

Parameters	Value
WLAN technology	IEEE 802.11a OFDM at 54 Mbps
Frequency channels (MHz)	5170, 5230, 5290, 5735, 5795, 5815
Transmission power (dBm)	7
Receive 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

consider signal loss that occurs due to obstacles. Thus, it enables exceptionally large signal range which provides the worst case scenario necessary to validate the proposed design.

In the simulation scenarios discussed, two terminal nodes—train and TCC—placed at the two ends of the chain transmit packets which are then transferred to the other terminal node over the chain. This results in two packet flows, one in each direction. A packet rate of 1000 packets per second is used (per flow), resulting in a combined data rate of 8.2 Mbps. A network size of 100 nodes has been used in our simulations as it 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 existing network infrastructure available at the train stations will be used to provide a wired connection to the two nearest nodes of the chain. Thus, chain length will be much smaller.

When discussing results, we are particularly interested in six performance indicators, namely unique packets received, duplicate packets received, total packets received, collisions, erroneous packets received, and, packets lost. Number of unique packets received serves as our key parameter as it indicates how many of the original unique packets (i.e. excluding duplicates) sent by the train 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. Total packets received includes duplicate packets. Erroneous packets are a result of interference between transmissions from different nodes (including collisions). These packets are discarded and do not count towards the aforementioned packets received numbers. Packets lost is the number of packets that, out of the original unique packets sent, were not received at the receiving end, for example owing to errors.

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.

A. Results for the original design

In the first part of the study, the performance of the original design with three frequencies is studied. Fig. 5 shows the results for the above mentioned six parameters against each

node listed on the x-axis. Note that only a select set of nodes are listed for readability. First the first five nodes are listed—to study the interference seen on the nodes near the train—and then every tenth node is listed. The y-axis shows the number of packets received as a percentage of the unique packets sent. Note that with the rate of 1000 packets per second and a simulation time of 60 seconds, the number of packets sent by a single radio on a terminal node during the whole simulation run is 60,000. Thus, a 100% unique packets received by a node implies that it received all 60,000 packets. Note that the total number of packets sent by a terminal node is thrice this number—as it sends on three radios. And ideally, the total number of packets received by a chain node in one direction is twice this number—as it receives packets from two of its preceding nodes.
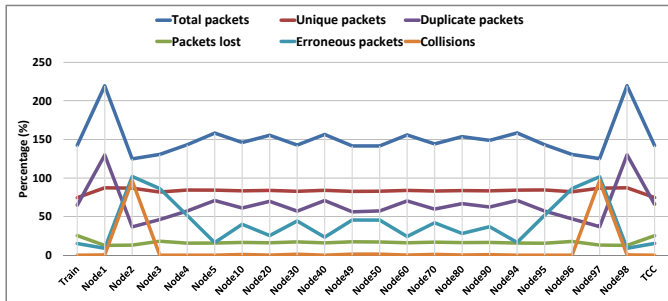

Fig. 5. Results for the original three-frequency design

Fig. 5 shows that on a middle node (i.e. Nodes 10 to 90), on average 149.36% total packets (blue line) are received. This is a low number as a node shall ideally receive two copies of the same packet, i.e. 200%, as described above. Likewise, the number of duplicate packets received (purple line) is on average 63.93%, a much lower number than the ideal 100%. Note that for brevity, the figure only shows an average of the number of total and duplicate packets received at each node. In reality, since two flows are involved, a significantly higher number of packets is received at each node as it includes the packets flowing in the opposite direction as well.

As seen, a stable number of unique packets received (red line) is maintained throughout the chain. Nonetheless, only 74.58% of packets are successfully delivered to the terminal nodes at the each end. These low numbers of total and duplicate packets received and the significantly higher packet loss (25.42%) are due to the interference produced by the nodes beyond the frequency separation distance. These nodes are still able to interfere because due to the favorable propagation conditions, their interference range exceeds the frequency separation distance. This is evident from the results for erroneous packets (light blue line) that show that on the middle nodes, on average 17.21% erroneous packets are received per flow. In comparison, in the one flow scenario presented in [1], only 9.63% erroneous packets were received on average. This increased interference is expected as traffic is now flowing in both directions. Thus, each top radio faces interference from nodes on its both sides.

The significant drop in the number of total and duplicate packets received at the nodes near the two ends of the chain (Nodes 2-4 and 95-97) is due to the sharp rise in the number of erroneous packets. This rise is due to another limitation of the design reported in [1]. Namely, as a train transmits on all frequencies in all directions, the inherent frequency separation guaranteed otherwise inside the chain is not fully achievable, resulting in interference on the nearby nodes. For example, at Node 2, the 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 to experience collisions. While Nodes 3 and 4 are outside the train's transmission range, they are still in its interference range. For example, at Node 3, the train's transmissions interfere with those of Nodes 1 and 2. Nonetheless, as seen in Fig. 5, due to the redundant design, only a minor drop in the number of unique packets received (red line) is seen at these nodes.

As discussed in [1], the result showed an end-to-end delay of 2.67 milliseconds at the terminal nodes, which is well below the typical end-to-end delay of 500 milliseconds specified in the IEEE CBTC standard [12]. Since transmissions are separated by frequencies, there can be at most only two nodes contending for the medium on a given frequency and location in the chain. Thus, MAC contention delay and queueing delay are irrelevant. Likewise, as reported in [1], no MAC layer retransmissions are made in this design. Given these reasons, end-to-end delay is not of particular interest for this work.

To summarize it, two limitations are identified here: (1) despite the frequency separation distance, nodes beyond this distance still interfere due to their large interference range, (2) a train's use of all frequencies in all directions for transmission causes interference on the nearby chain nodes. The focus of this paper is to minimize the first of these limitations.

B. Results for the improved design

To minimize the interference that a node causes beyond the frequency separation distance, in this second part of the study, the frequency separation distance is extended by increasing the total number of frequencies used in the system and employing them in an interleaving fashion at each node, as described in Section IV. The total number of frequencies is increased to four, five and six in steps to examine closely how it affects the interference. The results for the scenario with four frequencies are presented in Fig. 6.

The results show that when compared to Fig. 5, an extended frequency separation has resulted in a lowered number of erroneous packets (light blue line) on the middle nodes due to the lowered interference. The result is that the number of packets transferred to the terminal nodes has increased from 74.58% to 82.77%. In other words, the packet loss has dropped from 25.42% to 17.23%.

The results for five frequencies are presented in Fig. 7. They show a further drop in the interference, as is also evident from the "smoothed out" lines for erroneous, total and duplicate packets received. For example, the number of total packets
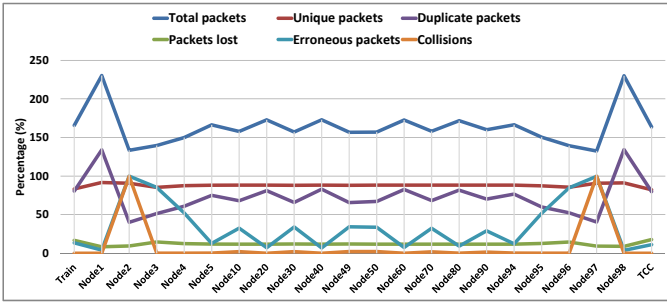
Fig. 6. Results for the design with 4 frequencies

received by a middle node has increased to 170.6% on average, compared to 149.36% seen in Fig. 5. As a result, the packet loss seen at the terminal nodes has further dropped to 16.22%.
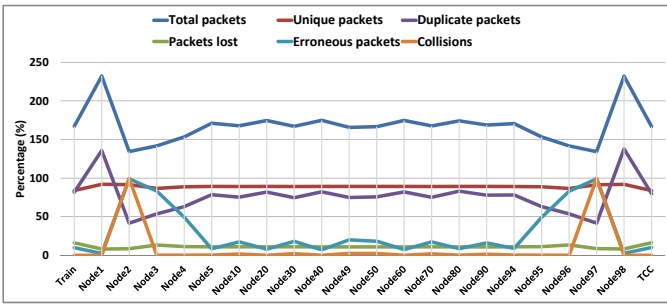

Fig. 7. Results for the design with 5 frequencies

Finally, Fig. 8 shows the results for the six-frequency design. As expected, a significant further drop in the erroneous packets is seen. Specifically, on the middle nodes, on average only 1.93% erroneous packets are received compared to 17.21% seen in the three-frequency scenario (Fig. 5). As a result, the packet loss at terminal nodes drops to 14.54%, a significant drop from 25.42% seen in Fig. 5.
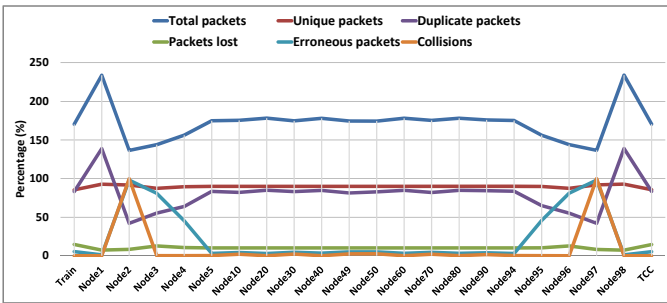

Fig. 8. Results for the design with 6 frequencies

Notably, despite that now the number of erroneous packets caused by the chain nodes due to the insufficient frequency separation has dropped to only 1.93% per node, the total packet loss seen at a terminal node is still 14.54%. This indicates that the interference caused by the terminal nodes at their nearby nodes (Nodes 2-4 and 95-97) is a greater contributor of the total packet loss. Employing additional frequencies might

thus not improve the situation significantly unless the interference from the terminal nodes is also minimized. Specifically, the results indicate that out of the 25.42% packet loss seen in the original three-frequency scenario, approximately 10.88% was introduced by the interference from the nodes inside the chain and 14.54% by the terminal node.

C. Comparison of the original and the improved design

For the sake of comparison, Fig. 9 presents the number of erroneous packets received for different number of frequencies. As seen, increasing the frequency separation distance minimizes the amount of interference dramatically.
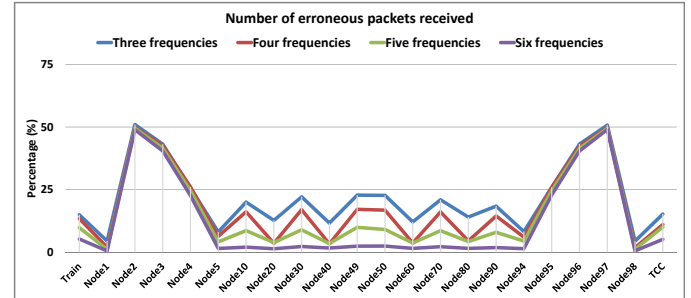

Fig. 9. Erroneous packets received for different number of frequencies

Additionally noteworthy in Fig. 9 is that despite a terminal node now transmits on a greater number of frequencies, it has not increased the interference on the nearby nodes. The reason is that a chain node still transmits on only two frequencies. Thus, the total number of frequencies interfering with a terminal node's transmissions has not changed.

Finally, Table II summarizes the results.

TABLE II
PERFORMANCE COMPARISON OF DIFFERENT NUMBER OF FREQUENCIES

Frequencies	Number of packets received at middle nodes (%)			Packet loss at terminal nodes (%)
	Total	Duplicate	Erroneous	
3	149.36	63.93	17.21	25.42
4	164.4	74.21	10.67	17.23
5	170.6	78.6	6.5	16.22
6	176.5	83.67	1.93	14.54

It is emphasized that the 14.54% packets loss is acceptable as the network size used is exceptionally large—60 kilometers—and the train control information in CBTC is redundant in nature and is sent both ways at regular intervals.

The results emphasize that the interference is a function of distance between nodes transmitting on the same frequency— or in other words, their signal ranges—and demonstrate the effectiveness of the design as it allows to extend this distance by employing additional frequencies. Note that the large number of frequencies used is not of a particular concern here as conventional CBTC systems are already known to employ multiple frequencies to improve availability [10]. Additionally, a vast majority of CBTC systems in operation today

work in the license-free Industrial, Scientific and Medical (ISM) frequency band [10]. Furthermore, the objective here primarily is to demonstrate that such a solution is feasible. As discussed above, our simulations use the simplistic FSPL propagation model. The results show that more realistic, less favorable propagation conditions will improve the performance by negatively affecting the signal range and thus lowering the interference. Thus, a design with a fewer number of frequencies, e.g. four, might as well yield the desired results.

VI. FUTURE WORK

Future work will focus on proposing solutions to minimize the interference caused by the train's transmissions as well as scenarios with a greater number of trains and mobility.

VII. CONCLUSIONS

This paper extends the previously presented design of an ad-hoc based trackside radio communication network for train to trackside communication in 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 train thus does not have to perform a handshake with the nodes (as in conventional infrastructure Wi-Fi) as it moves and the costly optical fiber cables connecting the nodes are no more needed. The design offers resiliency against interference by employing multiple frequencies. Nonetheless, the result show a significant amount of packet loss due to the interference inside the chain. This paper extends the design by employing additional frequencies in an interleaving fashion in order to optimize the frequency separation distance. The results show that increasing the frequency separation distance very effectively minimizes the interference. As a result, a significantly large numbers of packets can be transferred across large networks with only limited packet loss.

REFERENCES

[1] 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)*, sep 2017.

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

[3] 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.

[4] 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.

[5] 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.

[6] 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.

[7] Q. Dong, S. Banerjee, and B. Liu, "Throughput optimization and fair bandwidth allocation in multi-hop wireless LANs," in *Proc. IEEE International Conference on Computer Communications (INFOCOM '06)*, apr 2006.

[8] T. Tainaka, H. Masuyama, S. Kasahara, and Y. Takahashi, "Performance analysis of burst transmission mechanism for IEEE 802.11-based multi-hop wireless LANs," *IEEE Trans. Wireless Commun.*, vol. 10, no. 9, pp. 2908–2917, 2011.

[9] Z. Zeng, Y. Yang, and J. C. Hou, "How physical carrier sense affects system throughput in IEEE 802.11 wireless networks," in *Proc. IEEE International Conference on Computer Communications (INFOCOM '08)*, apr 2008.

[10] 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.

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

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