

On the Feasibility of Collision Detection in Full-Duplex 802.11 Radio

Michele Segata, Renato Lo Cigno

Dept. of Information Engineering and Computer Science, University of Trento, Italy

{msegata, locigno}@disi.unitn.it

Abstract—Full-duplex radios are becoming a feasible reality thanks to recent advances in self-interference cancellation. Switching from half- to full-duplex requires a major re-design of many network features and characteristics, including the MAC layer. The literature provides several new proposals or improvements that are applicable in different topologies: centralized, distributed, and multi-hop Wireless LANs (WLANs). These proposals, however, mostly focus on directional, unicast communication. While the main goal of unicast-focused approaches is to get as close as possible to doubling the throughput, it is still unclear how to exploit full-duplex radios in broadcast-like environments such as the vehicular one or in general in WiFi-like scenarios where interference is the dominating impairment. In this work we analyze the possible benefits and drawbacks of exploiting self-interference cancellation in full-duplex radios to implement collision detection. We show that, if proved feasible, the required changes to the MAC layer of an 802.11-based transceiver would be minimal, and could largely improve the performance with respect to a standard collision avoidance mechanism. In addition, the paper discusses the tricky aspects and the parameters required to identify a collision in a wireless network and discusses the many differences between managing collisions in a wired and in a wireless environment.

I. INTRODUCTION

Collision detection in wireless random access protocols has traditionally been considered unfeasible due to the extreme difference between the transmitted and the received power. Recent achievements in mixed analog plus digital cancellation [1], [2] have made full duplex radios feasible, and spawned a burst of research on their use, including many new proposals for MAC protocols that we revise in Section II. Many proposals exploit bidirectional transmissions but, to the best of our knowledge, none of them explore feasibility and properties of collision detection with full duplex radios. The wireless scenario remains profoundly different from a coaxial cable or hub, where the attenuation is very limited and collisions are detected deterministically, and all stations detect them. In the wireless domain the collision can indeed be identified mainly by the transmitters, and other stations will simply see a malformed and interrupted frame: in the best case, the receiver can detect a collision by exploiting the capture effect [3], [4], but taking countermeasures is a duty of the transmitters. Moreover, the detection may be non-deterministic, and in some scenarios like vehicular networks or large open-air WiFi installations,

This research is partially supported by the European Commission, H2020-ICT-2015 Programme, Grant Number 688768 ‘netCommons’ (Network Infrastructure as Commons).

the notion of channel itself becomes blurred, as there may be intrinsic spatial reuse and stations very far one another are hidden with respect to stations in between. Still, the possibility of detecting collisions and avoid the waste of channel time is appealing and has only partially been investigated [5], [6].

This paper explores the possibility of exploiting full-duplex capabilities to perform collision detection and abort the transmission. In here, we focus on broadcast-based networks such as the vehicular one, as the gain in performance that full-duplex radios can bring in this domain has not yet been explored. Moreover, network overload in vehicular environments is still an open issue [7]–[10] and the ability to perform collision detection might become a game changer.

We start by analyzing the changes required to a standard 802.11 transceiver and the differences with detecting collisions in wired networks, identifying the parameters that can impact on the performance (Section III). We implement our solution in a network simulator and perform an extensive simulation campaign to understand the impact of the aforementioned parameters (Section IV). We first study the impact of different channel loads in a small setup with a constant number of nodes (Section IV-B). Then, we fix the channel load and explore the behavior in a spatially distributed network (Section IV-C). As an additional contribution, we analyze the impact of different contention window sizes and policies (Section IV-D). Finally, we conclude the paper by summarizing our findings and by highlighting some future investigation directions (Section V).

II. BACKGROUND AND RELATED WORK

A. PHY Layer Principles

The first step towards full-duplex wireless networking is the physical layer design. Indeed, classical radios are half-duplex because the self-interference generated by the own transmitted signal is several order of magnitudes stronger than the received one, so decoding the latter is unfeasible. Even though our focus is on the MAC layer, for the sake of completeness we briefly discuss how full-duplex is achieved at the PHY layer.

As briefly anticipated, full-duplex is not possible in common radios because of the self-interference caused by the own transmitted signal. The solution is thus to reduce self-interference by cancelling the transmitted signal from the receive chain. This operation, however, is non trivial and requires a combination of analog and digital cancellation. Analog cancellation works on the raw signal using analog hardware, while digital cancellation works on the baseband, digitized signal samples. The work

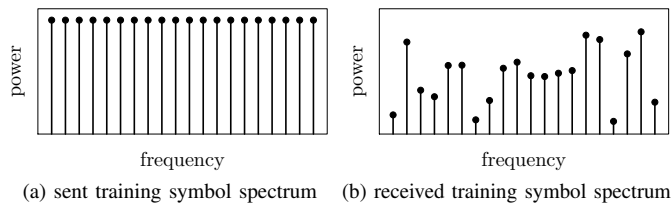


Figure 1: Channel frequency response estimation through training symbols.

in [11] described a simple example of an analog cancellation principle. By placing two transmit antennas at a distance d and $d + \frac{\lambda}{2}$ from the receiving antenna, with λ being the carrier wavelength, the transmitted signals will interfere destructively at the receiving antenna. As pointed out in [1], however, this design works well only at the center frequency: as we move from the center frequency the performance degrades, making it unfeasible for wideband signals. This example, however, clearly explains the analog cancellation principle. The works in [1], [2] propose more sophisticated designs.

After the analog cancellation phase, the signal is down-converted to baseband and sampled through an Analog to Digital Converter (ADC). At this point, however, the received signal is not yet completely clean, as the transmitted signal comes at the receiving antenna through multi-path propagation. The digital cancellation mechanism proposed in [1] estimates the channel frequency response of the transmitted signal. For example, in OFDM-based technologies like IEEE 802.11 we can exploit the training symbols in the preamble to estimate the channel. As the training symbols are known, it is possible to analyze the spectrum of the received symbols to determine channel frequency response. Figure 1 shows an example. If we send an OFDM symbol with a flat spectrum (Figure 1a) and we receive a symbol with a different spectrum (Figure 1b), we can estimate the channel frequency response. To further clean the received signal, we apply the estimated frequency response to the known transmitted signal and subtract it to the received one. For a better understanding of the working principle, Figure 2 shows a simplified representation of a full-duplex radio.

The combination of analog plus digital cancellation permits to cancel, depending on the implementation, between 70 dB to 110 dB of self-interference [1], [2], [12]. This permits to obtain a received signal with a high enough Signal to Noise Ratio (SNR) for successful decoding.

B. MAC Layer Solutions for Full-duplex Radios

PHY layer full-duplex alone permits the bidirectional communication between a single pair of nodes. In the wireless domain, however, multiple nodes share the same channel, so they need to coordinate their channel access to avoid interference and thus loss of packets. The MAC layer implements the coordination: for example, IEEE 802.11 implements a CSMA/CA protocol, meaning that the coordination is implemented through a “listen before talk” policy (CSMA), trying to avoid simultaneous transmissions by using a randomized access time (collision avoidance). Using a standard CSMA/CA approach on top of

a full-duplex radio would clearly give no additional benefit compared to standard half-duplex, as the protocol would forbid the transmission of a frame while receiving another.

There is thus the need to redesign the MAC layer taking into account the possibility of simultaneous transmission and reception. This task is, however, non-trivial. Full-duplex might induce to think that we can easily double the throughput of a wireless network, but this is only the case for a single pair of nodes, and sometimes not even in this very simple case due to residual self-interference. Even in the case of perfect cancellation, full-duplex does not always guarantee a double throughput. The analysis in [13] shows that the gain in performance is only slightly larger than 1. The main reason is the reduction in spatial reuse and asynchronous channel access.

To understand the first reason, consider the example in [13] (Figure 3). In Figure 3a the nodes $T1$ and $T2$ can simultaneously transmit because they are outside of each other interference range. As receivers $R1$ and $R2$ are close to transmitter $T1$ and $T2$, respectively, they can successfully decode their frame even if they are experiencing some interference. By enabling full-duplex, however, the node $R1$ would simultaneously transmit while receiving from $T1$, blocking the transmission of node $T2$ due to carrier sensing. In this example, the overall throughput is the same in both cases.

The second reason regards the asynchronous access scheme of standard CSMA/CA. Consider again the example scenario in Figure 3a. Imagine that $T1$ starts to transmit and that $R1$ has some data to send back. Before $R1$ can start transmitting it must wait for $T2$ to complete the transmission, perform the backoff procedure, and only then, if $R1$ wins the contention, it can start to transmit. This will cause the full-duplex transmission to be only partially overlapped, potentially wasting resources. In addition, the performance of the interference cancellation mechanism can have a huge impact on the throughput [14].

Obviously there are also very interesting benefits. For example, the scenario in Figure 3b can also be seen as a protection against hidden terminals. In addition, in infrastructure-based Wireless LANs (WLANs), full-duplex can reduce the bottleneck problem at the Access Point (AP), which in standard 802.11 gets the same share of the channel as associated nodes. This initial literature review, however, shows that enabling full-duplex does not immediately imply doubling the throughput. On the contrary, the MAC layer design will need to be smart enough to maximize the gain.

Some works build the MAC on top of their PHY layer. For example, the work in [1] proposes both a PHY and a MAC layer design. The idea in their work is to have a primary transmitter that initiates a frame transmission. When the receiver decodes the MAC addresses, if a frame for the primary transmitter is available it immediately starts a secondary transmission, protecting itself from hidden terminals and exploiting the full-duplex capability. If one of the two frames is shorter than the other, the remaining time is filled with a busy tone to maintain hidden terminal protection and to enable a simultaneous ACK transmission at the end of the frames. If the receiver has no frame for the primary transmitter, the secondary transmission

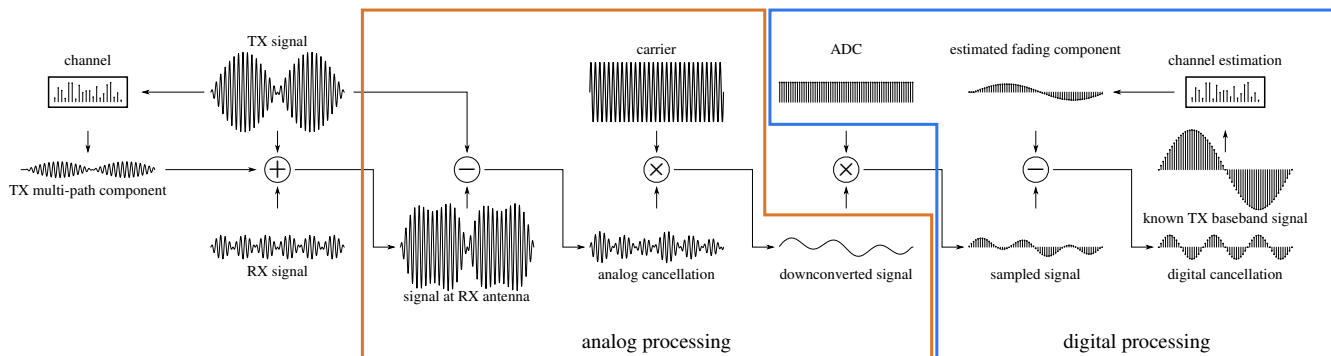


Figure 2: Simplified working principle of a full-duplex PHY.

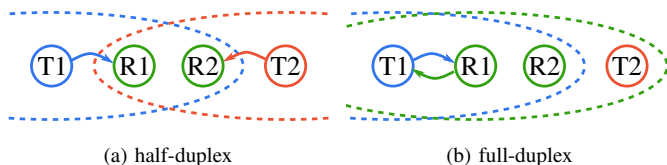


Figure 3: Reduction of spatial reuse in case of full-duplex transmissions. Dashed lines represent the interference range.

is only a busy tone for hidden terminal protection.

Busy tones are clearly a waste of resources, so the authors suggest a future modification that consider flexible MAC queues. To maximize the chances of having a frame for the primary transmitter, they propose a per-destination queue.

The work in [15] proposes Janus, a centralized MAC protocol for an infrastructure-based WLAN. The proposed idea is similar to the polling mechanism of IEEE 802.11 [16]. The AP starts each “poll” cycle by sending a Probe Request, asking which stations have frames ready to send. The interested nodes reply with a frame that includes the length of all the frames to be sent and the amount of interference they are experiencing from all other nodes. By using this information, the AP schedules transmission times and bit rates to maximize the throughput. In contrast to standard CSMA where simultaneous transmissions should never occur, Janus allows simultaneous transmission by choosing a bit rate that enables decoding given the amount of interference estimated during the scheduling phase.

For the sake of brevity we do not list additional MAC designs for full-duplex radios, but we briefly describe some collision detection approaches as they are more related to our work. The interested reader can refer to [17] for an exhaustive survey on PHY and MAC layer mechanisms for full-duplex radios.

In the literature we find some collision detection proposals for wireless networks, although not in the classical CSMA/CD philosophy. The authors of [3] exploit the capture effect to detect when a collision occurs. Frame capture continuously performs frame detection (for example, autocorrelation for OFDM) even while currently receiving. When detecting an incoming frame with a much stronger signal than the current one, the transceiver can switch to the new one, as there is a high chance of being able to decode it. A switch clearly indicates that a channel collision occurred, and the partially decoded

information from the first frame can be used to identify the transmitter. The paper, however, only focuses on the detection and does not propose how to exploit partially recovered data.

The work in [4] also considers the capture effect to distinguish between losses due to collisions and channel noise. Their proposal actively notifies other nodes about the occurrence of a collision, proposing a collision-aware backoff algorithm to improve the efficiency of the protocol.

An approach that performs collision detection and transmission interruption similarly to a cabled network is explored in [5]. The authors propose to use narrow, out-of-band signals to detect collisions and, in case, abort the transmission and retry. An out-of-band signal clearly requires to modify the hardware of the transceiver, which is undesirable. The work in [6] also performs detection and transmission abortion, but in this case they exploit a second receiver chipset that continuously performs autocorrelation (i.e., frame detection) to detect the presence of simultaneous frames. Compared to the solution presented in [5] the implementation is easier, as it does not require the design of specific hardware, but simply installing an antenna and the autocorrelation logic.

In this work we explore the use of a standard full-duplex radio for collision detection. Full-duplex clearly requires dedicated hardware, but once developed the transceiver can be exploited “as is” for the purpose.

III. ENABLING COLLISION DETECTION

As highlighted in the related work section, doubling the throughput in a full-duplex wireless network is not a trivial task and seems to be better suited for unicast traffic patterns. In pure (or mainly) broadcast wireless networks, like a vehicular network, it is not clear how to efficiently exploit full-duplex. The main problem of a broadcast network is the lack of feedback in case of failure, while 802.11 unicast traffic exploits acknowledgements to understand whether a communication is successful or not. As a broadcast communication is directed to any receiver within the reception range, the concept of acknowledgement is fuzzy, and indeed 802.11 does not use ACKs. As a consequence, when a transmission collides with another one, higher layers have no means to detect it, if not by implementing their own detection and recovery mechanism.

In addition, a problem affecting both unicast and broadcast communication is that the detection of an unsuccessful

transmission can only be performed at the end of the frame, because in standard radios the amount of self-interference does not permit to listen to the channel while transmitting. As a result the entire frame time is inefficiently wasted.

With a full-duplex radio we can cancel the self-interference and listen to the state of the channel. In the general setup, the idea would be to receive a frame while transmitting another, but another use case is understanding the state of the channel, and thus to detect concurrent transmissions. We propose to exploit full-duplex capabilities to detect collisions and interrupt an ongoing transmission, similarly to what is done in CSMA/CD protocols like Ethernet, to name the most famous.

The change to the standard IEEE 802.11 protocol is minimal and is listed in Listing 1. First of all, enabling collision detection provides the MAC layer with a negative feedback, and we can exploit it to retry a failed transmission and increase the contention window, which is clearly not done for standard 802.11 broadcast. We thus need to choose an amount of maximum retries after which the MAC layer will give up and continue with the next frame in the queue.

The first main change is on the procedure that handles the beginning of a frame reception. With a full-duplex network card we can start receiving a frame while transmitting. We can assume that an incoming frame is discovered through signal detection (i.e., autocorrelation) as when the network card is not transmitting. If that happens, we need to decide whether to interrupt the transmission or not. In contrast with a cabled network where a collision is deterministic, in a wireless network an incoming signal can have different power levels in a range that can span 50 dB. To cope with this, we need to set a power threshold δ (Line 5) above which a frame is considered a collision and the transmission should be aborted. By setting an extremely low value for δ , the protocol behaves in a conservative mode and aborts the transmission every time an incoming frame is detected. On the contrary, setting it very high simply disables collision detection and the protocol behaves as in standard CSMA/CA. The only difference is that the network card tries to decode the incoming frame thanks to its full-duplex capability.

After detecting the collision, we should abort the transmission (Line 6). In this case we wait for a certain amount of time before aborting the transmission: the aim is being sure that the other transmitter detects the collision as well. In here we assume that a slot time is enough, but this might depend on the hardware or on other compatibility issues (see Section V). In addition we assume no need for additional mechanisms (e.g., busy tones) to enable other receivers to detect the interruption.

After a slot time, the modified IEEE 802.11 MAC layer should stop the ongoing transmission and decrement the number of attempts (Line 8). When reaching the maximum number of attempts, we should drop the frame and notify upper layers about the failure. In addition, we need to reset the number of attempts and the contention window, as in legacy 802.11. In case of additional attempts left, the protocol should update the contention window as usual.

Finally, the MAC layer should assess the state of the channel

Listing 1: Proposed collision detection and handling mechanism for IEEE 802.11.

```

1 Event OnInit()
2   attempts = Amax
3   cw = CWmin
4 Event OnStartRX(frame)
5   if state is TX and frame.power >  $\delta$  then
6     schedule(OnStopTx, Tslot)
7 Event OnStopTx()
8   stopCurrentTransmission()
9   attempts = attempts - 1
10  if attempts == 0 then
11    dropFrame()
12    attempts = Amax
13    cw = CWmin
14  else
15    cw = min((cw + 1) × 2 - 1, CWmax)
16  performCCA()
17  performBackoff()

```

using the Clear Channel Assessment (CCA) function and perform the post transmit backoff procedure.

While the changes to the original MAC layer are relatively straightforward, the choice of some parameters is not trivial. First of all, the collision detection threshold δ can completely change the performance. A very low value is conservative and tries to avoid collisions at all. As a counter-effect, this might reduce spatial reuse. In contrast, a higher value tolerates a certain amount of collisions in favor of spatial reuse. In addition, low level parameters such as the contention window, number of attempts, and queue size can also highly influence the behavior. In the next section we perform a simulative analysis of the protocol variant under different parameter settings.

IV. SIMULATIVE ANALYSIS

We study the performance of the enhanced, collision-detection enabled 802.11 MAC by means of simulations. We modify the Veins vehicular networking simulator [18] to enable the possibility of interrupting an ongoing transmission and to introduce the changes in Listing 1. The reason why we choose a vehicular networking simulator such as Veins is because it implements a detailed IEEE 802.11p PHY and MAC layer [19], which enables us to realistically simulate low level phenomena such radio sensitivity, interference, SNR-based packet reception probability, etc. As the envisioned traffic pattern of vehicular networks is mainly broadcast, we believe that studying the behavior of collision detection in an 802.11p network is a good starting point.

We begin the analysis by considering a relatively small setup with 64 vehicles placed on 4 lanes (16 vehicles per lane) randomly spaced following an exponential distribution with mean $\mu \simeq 42$ m, corresponding to a spacing measured in

Table I: Simulation parameters.

	Parameter	Value
communication	Path loss model	Free space ($\alpha = 2.0$)
	Frequency	5.89 GHz
	Bandwidth	10 MHz
	Bit Rate	6 Mbit/s
	Transmit power	20 dBm
	CCA threshold	-65 dBm
	Noise floor	-95 dBm
	Minimum sensitivity	-94 dBm
	PHY model	IEEE 802.11p
	MAC model	1609.4 single channel (CCH)
	MSDU size	300 B
	Packet arrival	exponential $\lambda = 1$ Hz to 200 Hz
	Detection threshold δ	$-\infty, -85, -65, -45, \text{ and } +\infty$ dBm
	CW_{\min} and CW_{\max}	3 and 15
	Max attempts A_{\max}	3
MAC queue size	2	
scenario	Number of nodes	64
	Number of lanes	4
	Node distance (on lane)	exponential $\frac{1}{\lambda} = \mu \simeq 42$ m [20]

dense traffic scenarios [20]. Regarding the PHY and MAC configuration, we consider an 802.11p radio tuned in the 5.9 GHz band using a 6 Mbit/s bit rate (10 MHz bandwidth). The network setup considers the beacon generation process to be exponentially distributed with λ ranging between 1 and 200 Hz. As collision detection threshold δ we consider values of $-\infty, +\infty, -85, -65, \text{ and } -45$ dBm. Finally we set CW_{\min} and CW_{\max} to 3 and 15, respectively, A_{\max} to 3, and the queue size to 2 frames. Table I lists additional simulation parameters.

The choice of a small number of nodes lies behind the easiness of interpreting the results. By having 16 vehicles per lane with an average distance of 40 m, we obtain a scenario size in the order of 700 m, thus having all vehicles within the same interference range. Each simulation lasts 10 s and is repeated 10 times for increasing the statistical confidence.

A. Metrics

During the simulation each node logs statistics such as the number of correctly received frames, the number of collisions, the number of frames sent, etc. In the following we describe the logged information and the metrics we compute, to ease the interpretation of the graphs.

- **Received frames ($R_{i,s}$):** Number of frames correctly decoded by node i during simulation s .
- **Receive attempts ($RA_{i,s}$):** Number of frames that node i tried to receive. Reception might have been unsuccessful due to transmission abortion or interference.
- **Channel busy rate ($B_{i,s}$):** Fraction of time node i perceives the channel as busy, excluding overhead times such as Inter-Frame Spacings (IFSs) and backoff procedures.
- **Collision rate:** Amount of collided frames over all receive attempts $RA_{i,s}$. This metric is measured at the receiver and include frames that were trying to be received but either their transmission was stopped (collision detected) or the receiver was unable to decode them due to interference.
- **Transmission abort rate:** Average number of transmission abortion per frame due to collision detection.

- **Offered load:** Indicates the average number of frames per second generated by each vehicle, which is equivalent to the λ parameter of the exponential distribution driving the generation.
- **Total delivered traffic ($\tau_{i,s}$):** Number of frames per second received by each node, per node. For a particular simulation s this is computed as $\tau_{i,s} = \frac{R_{i,s}}{(N-1) \cdot t_s}$, where N is the number of nodes in the simulation and t_s is the duration of the simulation. The reason for dividing $R_{i,s}$ by $N-1$ and not using the standard throughput definition is that, as opposed to a cabled network, we cannot clearly define what the channel is, as the channel itself is spatially distributed. Nodes in different positions can perceive a completely different throughput. In addition, this definition of $\tau_{i,s}$ permits to compare it to the offered load.
- **Delivered traffic gain:** Gain in $\tau_{i,s}$ with respect to standard 802.11 computed as $\frac{\tau_{i,s}^{CD} - \tau_{i,s}^{802.11}}{\tau_{i,s}^{802.11}} \cdot 100$.

B. Offered Load Analysis

We begin the evaluation by showing the results for τ . Figure 4 shows the $\tau_{i,s}$ averaged over all vehicles and simulation repetitions, i.e.,

$$\tau = \frac{1}{N \cdot |S|} \sum_{s \in S} \sum_{i=1}^N \frac{R_{i,s}}{(N-1)t_s}, \quad (1)$$

where S is the set of repetitions for a particular simulation setup and N is the number of vehicles.

The first notable comparison in Figure 4 is between standard 802.11 and the collision-detection enabled version with $\delta = +\infty$ dBm. The two curves are very similar, with the CD version providing a slightly larger τ . The difference is simply given by the fact that the full-duplex radio is capable of receiving a frame while transmitting another. When a collision between two transmitters only occur, they will be able to receive each other frames but will not abort the transmission because the δ threshold is set to ∞ . Their capability of receiving each other frames results in a slight performance gain. This result is as expected and witnesses the correct implementation of the CD mechanism in the simulator.

The second comparison is between standard 802.11 and the aggressive CD strategies ($\delta = -\infty$ and -85 dBm). The two strategies behave very similarly due to the small scenario setup, and they both increase τ by a non-marginal amount. In the best case (i.e., for an offered load between 100 and 150 frames per second) the gain is around 40% (Figure 4b). Finally, for the less aggressive CD strategy ($\delta = -65$ dBm), the gain is mostly negative. By becoming less aggressive ($\delta = -45$ dBm), the performance increases again, and it is comparable to standard CSMA/CA. The trend thus suggests that the aggressive strategies are the most beneficial, strategies in the middle tend to worsen the performance, while permissive strategies converge to CSMA/CA.

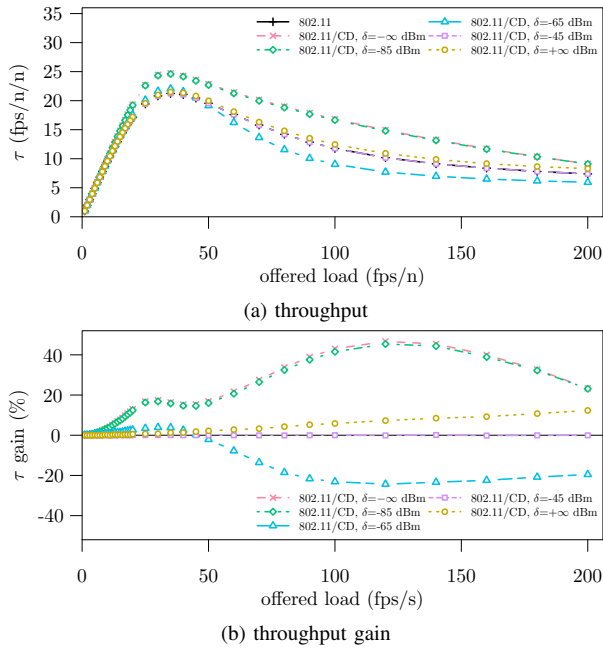


Figure 4: Throughput and throughput gain vs. offered load for the 64 nodes scenario.

To understand the behavior of the protocol and the results in Figure 4 we consider the channel busy ratio and the transmission abort rate (Figure 5). As for τ , each metric is averaged over all nodes and all simulation repetitions. We start by analyzing the transmission abort rate (Figure 5b). The first obvious result is that standard 802.11 and 802.11 with CD and $\delta = +\infty$ dBm result in no transmission abortion, as they both do not perform collision detection. When lowering the detection threshold for the collision detection mechanism the rate of aborted transmission increases. For $\delta = -65$ dBm the rate is moderate, as the threshold is high and thus a transmission is stopped only when a close neighbor is interfering. For the lowest values of δ , instead, the rate raises very fast with the offered load, as a minimal amount of interference causes the transmission to be stopped.

The large amount of attempts is also witnessed by the channel busy ratio in Figure 5a. The results for the aggressive CD strategies are rather counter-intuitive, as the busy ratio decreases with the offered load, for a load larger than 35 frames per second per vehicle. The reason lies behind the large number of attempts: for each transmission aborted the MAC layer needs to perform a new backoff procedure, increasing the overhead and thus reducing the effective channel utilization. Rather than being a waste of resources, however, coordinating channel access through an aggressive CD policy is definitely worth in terms of successful transmissions. This means that the increase in τ of an aggressive collision detection policy is obtained through a trial-and-error approach. In contrast, a more permissive strategy ($\delta = -65$ dBm) does not increase the benefit, suggesting that using more time to coordinate all nodes and completely avoid frame overlapping is very effective. A large overhead, however, also suggests that performing collision

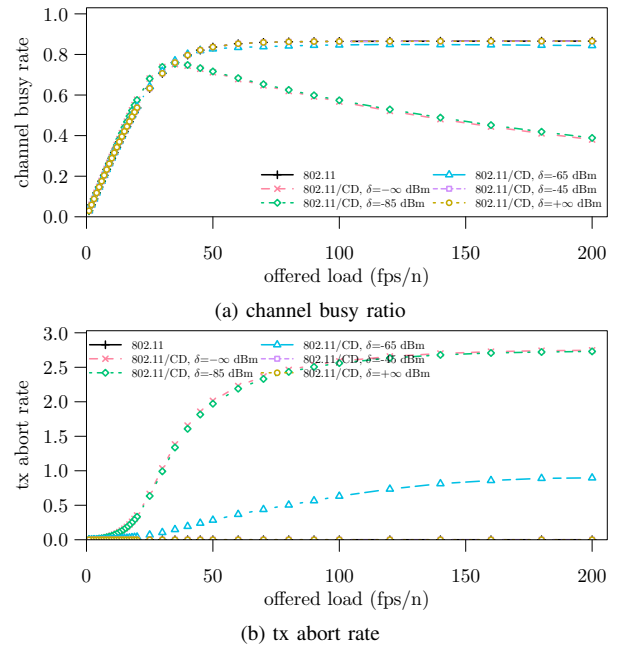


Figure 5: Channel busy ratio and transmission abort rate vs. offered load for the 64 nodes scenario.

detection for small frames might not be worth.

C. Impact of the Number of Nodes

The initial analysis considered a small scale scenario with a very large beacon generation rate. In this section, we fix the average beacon rate to 35 Hz, which maximizes τ in Figure 4, and change the number of vehicles from a minimum of 8 up to a maximum of 256, split over 4 lanes as before. By changing the number of nodes we increase the scenario size, including the case in which the scenario size is larger than the interference range. For 256 vehicles, the scenario extends on average for more than 2.5 km.

Figure 6 shows the results for τ and τ gain. The shape of the τ curve is different from the one observed in Figure 4. The reason is that each node generates the same load, and we change the number of vehicles. As the average frame rate is 35 Hz, in the best case we expect every node to receive 35 frames by each other node, which is indeed what we see for a small amount of nodes. For a small number of nodes the network is not overloaded and all nodes are located within the same interference range, so all protocol variants perform well.

As we increase the number of nodes, however, the performance quickly begins to drop. Standard 802.11 for a number of nodes around 50 performs 20% worse with respect to the aggressive CD strategies. Similarly to what we observed in Figure 4, the more permissive CD strategy does not provide much improvement, and performs worse than standard 802.11 for a number of nodes larger than 100.

Between 80 and 100 nodes the performance of the two aggressive strategies start to diverge. For $\delta = -85$ dBm the τ gain is mainly positive, but decreases with the number of nodes. Around 200 nodes, its gain becomes negative. For $\delta = -\infty$ dBm, instead, the performance gain reaches a maximum

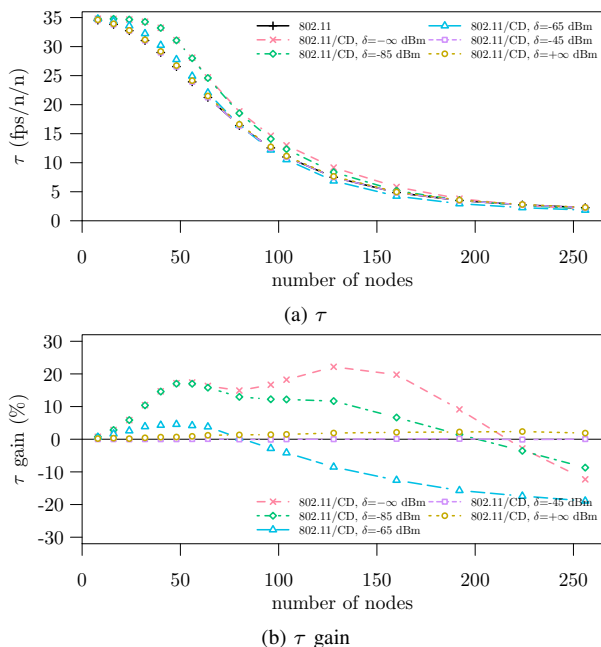


Figure 6: τ and τ gain vs. number of nodes for the 35 Hz scenario.

of 20% close to 150 nodes. The τ gain then starts to decrease and becomes negative between 200 and 250 nodes.

The negative performance gain for a large number of vehicles is due to the extension of the scenario. For networks which extend beyond the reception range, a CD strategy might interrupt a transmission even for a frame that the transceiver and most of its neighbors are not able to decode due to the large distance. Indeed in OFDM the short preamble can successfully be detected even for a negative SNR, but clearly the correct decoding of the payload is very unlikely. In such a setup the standard 802.11 approach might actually improve spatial reuse.

D. Impact of the Contention Window

Protocols with collision detection such as Ethernet are 1-persistent, meaning that when the channel becomes idle they immediately start to transmit and perform the backoff only in case of a collision. The rationale is that there is no need to waste time in performing a backoff if a collision can immediately be detected. In this section we test the behavior of the different Contention Window (CW) strategies using three different contention window growths. The first strategy is the one considered the previous analyses, i.e., 3, 7, and 15. We configure the other two strategies to implement Ethernet's 1-persistence. The second strategy uses the backoff window growth of Ethernet, i.e., 0, 1, and 3, while the last one uses a more conservative growth, i.e., 0, 3, and 15. The scenario under analysis is the one of Section IV-B, i.e., 64 nodes under increasing offered load.

Figure 7 shows the τ performance metric and the collision rate for the different CD strategies. The first noticeable fact is that the performance for different contention window strategies significantly changes only inside the offered load range around the maximum τ . Within this range, the strategy that provides

the highest performance is the one using an initial contention window of 4 slots, independently of δ (results for other values of δ omitted). This suggests that, for a moderate load, trying to immediately transmit when the channel turns idle can result in additional collisions, which in the end lowers the performance. This is confirmed by the collision rate analysis, which shows that a non-aggressive access strategy reduces the collision rate.

V. CONCLUDING DISCUSSION AND FUTURE WORK

The analysis in this work has shown that enabling collision detection through full-duplex radios can potentially improve the performance of a broadcast-based wireless network. The changes to the standard MAC protocol would be minimal, making it feasible for a real-world implementation. In contrast to pure CSMA/CD networks like Ethernet, however, we need to take into consideration some factors. First of all, in a cabled network determining whether a collision occurred is a yes/no decision. In a wireless system the answer is "continuous", meaning that two frames might collide at a particular receiver but might not at another. To cope with this we need to choose a collision detection threshold, which clearly influences the performance. In a relatively small setup an aggressive collision detection strategy is highly beneficial, but in a more spatially distributed network it is not yet clear: in this latter case the use of collision detection might even worsen the performance with respect to a standard CSMA/CA protocol.

The study has shown that collision detection obtains its performance gain through a trial and error approach, sacrificing a large portion of the available channel time for repeated backoff procedures. This means that collision detection might be worth only for large frame sizes. Yet we need to clearly define this threshold, and we plan to give an answer to the question through theoretical modeling. A model can help us understanding the impact of other MAC parameters such as contention window, number of retries, and queue size.

An additional problem is the compatibility with unicast traffic. When two stations send unicast traffic to each other, pure collision detection would make them abort the transmission and retry after the backoff. In such a case, however, it is much more convenient to continue the transmission, which would succeed thanks to the full-duplex capabilities of the radio. A solution to the problem would be to check the destination MAC address in the 802.11 header and interrupting the transmission depending on that address. This is feasible as 802.11 transceivers stream the bits to the upper layer as they are decoded, so the MAC does not need to wait for the end of the frame. This, however, would increase the collision detection time, as in the best case the MAC needs to wait at least 6 OFDM symbols (5 symbol times for the PHY preamble and header plus at least one symbol to obtain the address). In a 10 MHz channel this would require at least 48 μ s, so the efficiency of this mechanism should carefully be investigated.

Finally, even in a pure broadcast setting we need to understand the requirements for proper collision detection. In this work we assume that a single slot time is enough, but to confirm this assumption an experimental test is required.

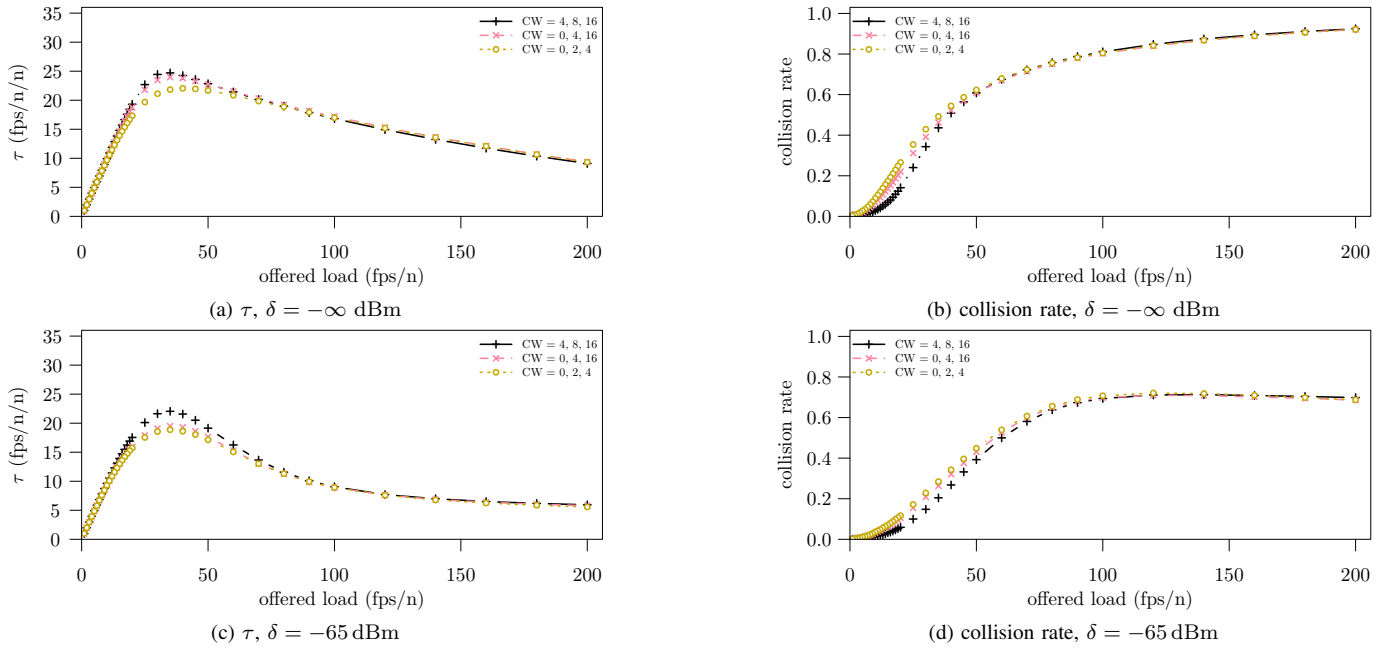


Figure 7: τ and C_r vs. offered load for the contention window analysis, 64 nodes scenario.

Moreover, we might need mechanisms to ensure that all nodes are aware of the collision (e.g., jamming sequence in Ethernet).

To conclude we believe that collision detection for broadcast-based 802.11 networks is a mechanism which is worth investigating, as it can bring several benefits compared to standard collision avoidance. Still, we need to clearly understand its real world feasibility together with its limits.

REFERENCES

- [1] M. Jain, J. I. Choi, T. Kim, D. Bharadia, S. Seth, K. Srinivasan, P. Levis, S. Katti, and P. Sinha, "Practical, Real-time, Full Duplex Wireless," in *17th ACM Int. Conf. on Mobile Computing and Networking (MobiCom 2011)*. Las Vegas, NV, Sep. 2011, pp. 301–312.
- [2] T. Chen, J. Zhou, N. Grimwood, R. Fogel, J. Maravsević, H. Krishnaswamy, and G. Zussman, "Full-duplex Wireless Based on a Small-form-factor Analog Self-interference Canceller," in *17th ACM Int. Sym. on Mobile Ad Hoc Networking and Computing (MobiHoc 2016)*, Demo. Paderborn, Germany, Jul. 2016, pp. 357–358.
- [3] K. Whitehouse, A. Woo, F. Jiang, J. Polastre, and D. Culler, "Exploiting the Capture Effect for Collision Detection and Recovery," in *2nd IEEE Workshop on Embedded Networked Sensors (EmNets 2005)*. Sydney, Australia, May 2005, pp. 45–52.
- [4] J.-H. Yun and S.-W. Seo, "Novel Collision Detection Scheme and its Applications for IEEE 802.11 Wireless LANs," *Elsevier Computer Communications*, vol. 30, no. 6, pp. 1350–1366, Mar. 2007.
- [5] J. Peng, L. Cheng, and B. Sikdar, "A Wireless MAC Protocol with Collision Detection," *IEEE Trans. on Mobile Computing*, vol. 6, no. 12, pp. 1357–1369, Oct. 2007.
- [6] S. Sen, N. Santhapuri, R. R. Choudhury, and S. Nelakuditi, "Moving Away from Collision Avoidance: Towards Collision Detection in Wireless Networks," in *8th ACM Workshop on Hot Topics in Networks (HOTNETS 2009)*. New York, NY, Oct. 2009, pp. 1–6.
- [7] J. Ahn, Y. Wang, B. Yu, F. Bai, and B. Krishnamachari, "RISA: Distributed Road Information Sharing Architecture," in *31st IEEE Conf on Computer Communications (INFOCOM 2012)*. Orlando, FL, Mar. 2012, pp. 1494–1502.
- [8] T. Tielert, D. Jiang, H. Hartenstein, and L. Delgrossi, "Joint Power/Rate Congestion Control Optimizing Packet Reception in Vehicle Safety Communications," in *10th ACM Int. Workshop on Vehicular Internet Networking (VANET 2013)*. Taipei, Taiwan, June 2013, pp. 51–60.
- [9] M. Segata, R. Lo Cigno, H.-M. Tsai, and F. Dressler, "On Platooning Control using IEEE 802.11p in Conjunction with Visible Light Communications," in *12th IEEE/IFIP Conf. on Wireless On demand Network Systems and Services (WONS 2016)*. Cortina d'Ampezzo, Italy, Jan. 2016, pp. 124–127.
- [10] M. Segata, F. Dressler, and R. Lo Cigno, "Jerk Beaconing: A Dynamic Approach to Platooning," in *7th IEEE Vehicular Networking Conf. (VNC 2015)*. Kyoto, Japan, Dec. 2015, pp. 135–142.
- [11] J. I. Choi, M. Jain, K. Srinivasan, P. Levis, and S. Katti, "Achieving Single Channel, Full Duplex Wireless Communication," in *16th ACM Int. Conf. on Mobile Computing and Networking (MobiCom 2010)*. Chicago, IL, Sep. 2010, pp. 1–12.
- [12] D. Bharadia, E. McMillin, and S. Katti, "Full Duplex Radios," *ACM SIGCOMM Computer Communication Review (CCR)*, vol. 43, no. 4, pp. 375–386, Oct. 2013.
- [13] X. Xie and X. Zhang, "Does Full-duplex Double the Capacity of Wireless Networks?" in *33rd IEEE Conf. on Computer Communications (INFOCOM 2014)*. Toronto, Canada, Apr. 2014, pp. 253–261.
- [14] Z. Tong and M. Haenggi, "Throughput Analysis for Full-duplex Wireless Networks with Imperfect Self-interference Cancellation," *IEEE Trans. on Communications*, vol. 63, no. 11, pp. 4490–4500, Nov. 2015.
- [15] J. Y. Kim, O. Mashayekhi, H. Qu, M. Kazandjieva, and P. Levis, "Janus: A Novel MAC Protocol for Full Duplex Radio," *Stanford Computer Science Tech. Rep.*, vol. 2, no. 7, p. 23, 2013.
- [16] P. Larcheri and R. Lo Cigno, "Scheduling in 802.11e: Open-Loop or Closed-Loop?" in *3rd IFIP Conf. on Wireless On demand Network Systems and Services (WONS 2006)*, Les Ménuires, France, Jan. 2006.
- [17] D. Kim, H. Lee, and D. Hong, "A survey of In-band Full-duplex Transmission: From the Perspective of PHY and MAC Layers," *IEEE Communications Surveys & Tutorials*, vol. 17, no. 4, pp. 2017–2046, Feb. 2015.
- [18] C. Sommer, R. German, and F. Dressler, "Bidirectionally Coupled Network and Road Traffic Simulation for Improved IVC Analysis," *IEEE Trans. on Mobile Computing*, vol. 10, no. 1, pp. 3–15, Jan. 2011.
- [19] D. Eckhoff and C. Sommer, "A Multi-Channel IEEE 1609.4 and 802.11p EDCA Model for the Veins Framework," in *5th ACM/ICST Int. Conf. on Simulation Tools and Techniques for Communications, Networks and Systems (SIMUTools 2012)*, Poster Session. Desenzano, Italy: ACM, March 2012.
- [20] N. Wisitpongphan, F. Bai, P. Mudalige, and O. K. Tonguz, "On the Routing Problem in Disconnected Vehicular Ad-hoc Networks," in *26th IEEE Conf. on Computer Communications (INFOCOM 2007)*, Anchorage, AK, May 2007, pp. 2291–2295.