Enhancing the Real-Time Behavior of IEEE 802.11n

Federico Tramarin and Stefano Vitturi
Institute of Electronics, Computer and Telecommunication Engineering
National Research Council of Italy, CNR–IEIIT
Via Gradenigo 6/B, 35131 Padova, Italy
{tramarin, vitturi}@dei.unipd.it

Michele Luvisotto
Department of Information Engineering
University of Padova
Via Gradenigo 6/B, 35131 Padova, Italy
luvisott@dei.unipd.it

***Abstract*—IEEE 802.11 systems are drawing an ever increasing interest for wireless industrial communication, also thanks to the interesting features provided by the most recent and advanced amendments to this standard, such as IEEE 802.11n. Due to the intrinsic unreliability of the wireless medium, the current research efforts aim at improving both timeliness and reliability of such a protocol in view of its adoption for real–time applications. A significant issue in this context is represented by the reduction of the randomness that affects packet delivery times. An important benefit in this direction can be obtained by the deactivation of the standard legacy carrier sensing and backoff procedures. In this paper we show, through a simulative assessment, that a fine control of such features leads to improved real–time performance.**

I. Introduction

The adoption of IEEE 802.11 wireless LANs (WLANs) [1] in industrial applications is ever more extensive thanks to several advantages they provide [2], making them a suitable opportunity for the replacement or extension of real-time Ethernet segments. Furthermore, the recent IEEE 802.11n High Throughput (HT) amendment introduced several enhancements, at both the physical and MAC layers, that can be exploited to improve some significant performance figures for real–time networks, principally in terms of reliability and timeliness [3].

IEEE 802.11 operates mainly in the 2.4 GHz Industrial, Scientific and Medical (ISM) band, also adopted by other wireless technologies such as IEEE 802.15.4 and IEEE 802.15.1, likely resulting in coexistence problems and poor channel conditions. To avoid packet losses and ensure fair access to the medium, the Distributed Coordination Fucntion (DCF) of the MAC sub–layer adopts the Carrier Sense Multiple Access with Collision Avoidance (CSMA/CA) scheme. According to this strategy, devices use a combination of Carrier Sensing (CS) and random backoffs in order to minimize the collision probability.

Although CSMA/CA has proven to be effective in general purpose communication systems, its adoption in the industrial context can be actually detrimental rather than beneficial [4]. Indeed, industrial traffic is often characterized by real–time requirements, such as low jitter on cyclic operations and bounded latency on alarm packets that, clearly, can be seriously compromised by latencies and randomness introduced by both CS and backoff. To cope with these issues, some available industrial communication protocols (e.g. Wireless HART and ISA100.11a based on IEEE 802.15.4 [5]) adopt high layers services to resolve contentions and transmission errors, for example exploiting a master–slave relationship in a polling or TDMA-based scheme: hence, distributed and stochastic channel access schemes do result unnecessary, if not dangerous, since they might downgrade the overall performance.

Unfortunately, a similar approach is usually difficult to apply to commercial IEEE 802.11–based devices. Indeed, the use of such components, although justified by their affordability as derived from the high production volumes, often imposes very limited customization possibilities. In other words, IEEE 802.11 devices typically can not be adequately configured to achieve a satisfactory degree of timeliness. However, recently released ”Soft–MAC” chips (such as Qualcomm Atheros ones) allow a precise control of protocol parameters, since most parts of the MAC sub–layer are executed in software and easier to control.

Within this framework, this paper proposes the proper tuning of some significant protocol parameters, and presents a simulation–based assessment of their impact on the performance of an IEEE 802.11n network, employed in an industrial context. Specifically, we focus on the effects that the deactivation of the CS and/or the backoff procedures have on some significant indicators. While some previous work dealt with a rather similar issue for IEEE 802.15.4 networks [4], to the best of our knowledge the problem has not been already addressed for more complex IEEE 802.11n–based systems.

II. Carrier Sense and Backoff in IEEE 802.11n WLANs

The DCF channel access mode adopted by IEEE 802.11 is based on the execution of CS and backoff procedures, as described below.

A. Clear Channel Assessment

The CS mechanism is built upon a PHY primitive called Clear Channel Assessment (CCA), which returns an estimate of the channel state, that can be either *idle* or *busy*. For the specific case of IEEE 802.11n the CCA works as follows:

- A node not sending HT frames should return a busy channel if it detects a *valid* signal over the -76 dBm threshold.
- A node sending HT frames, over a 20 MHz channel, should return a busy channel if it detects a *valid* 20 MHz HT signal over the -82 dBm threshold, or *any* signal over the -62 dBm level;

978-1-4799-8244-8/15/$31.00 ©2015 IEEE

- A node sending HT packets, over a 40 MHz channel, should sense the channel (composed by two paired 20 MHz channels) as busy if it detects a *valid* 20 MHz HT signal with a power over -82 dB, a *valid* 40 MHz HT signal with a power over -79 dBm, or *any* signal over the -62 dBm level.

B. Access to the channel

In the DCF procedure, a node that has a packet to send firstly performs a CS operation by means of the CCA function. If the wireless medium is idle for a period greater than or equal to a DCF Inter Frame Space (DIFS), the node is allowed to immediately transmit.

Conversely, the node defers its transmission, continuously performing CCA, until the medium is determined to be idle for at least a DIFS. Subsequently, to avoid possible synchronization with other devices simultaneously trying to access the medium, it generates a random backoff period, for an additional deferral time, as

$$T_{BO} = rand \cdot T_{SLOT} \quad (1)$$

where T_{SLOT} is the *slot time* defined in the standard, while *rand* is a random number uniformly drawn from $[0, CW]$. The Contention Window (CW) is initialized to CW_{min} and exponentially increased at each subsequent attempt as

$$CW(k) = 2 \cdot CW(k-1) + 1 \quad (2)$$

until the maximum value CW_{max} is reached.

During this backoff period, the node continuously keeps sensing the channel. If any activity is detected, the backoff timer is suspended and can be resumed only after the medium is found idle for a DIFS. When the timer finally expires, the node is allowed to send the packet.

The same random backoff procedure is adopted also to recover from a failed transmission, inferred from a missed reception of the ACK frame within a specific timeout.

III. Problem Statement

For the purposes of this performance assessment, we refer to a typical polling–based industrial protocol implemented on an Infrastructure IEEE 802.11n WLAN.

The protocol encompasses a *master–slave* relationship, in which the master cyclically polls each of the n slaves gathering some data (e.g. sensor readings). Both the master and the slaves are IEEE 802.11n compliant devices. Specifically, in an infrastructure network, the master is implemented by the WLAN Access Point (AP) while the role of slaves is relevant to WLAN Stations (STAs).

The network operations are graphically described in Fig. 1: the master sends a *request* packet of length L_{req} to the slave, which, after having confirmed reception with the ACK frame, sends a *data* packet of length L_{data} with the required information. The master acknowledges the reception with an ACK, before starting the following polling operation with another slave. Finally, in the considered industrial protocol, we bounded the time allocated for the polling of a single node to $T_{S,max}$: if the polling procedure exceeds this value, then the master considers that polling as failed and moves to the next slave.

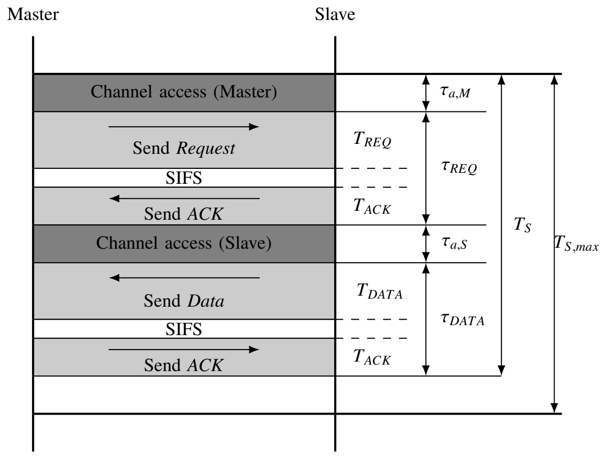

Fig. 1. Structure of a polling operation.

A. Performance indicators

For the performance assessment presented in this paper we refer to two principal indicators.

The first one is the *service time* for a specific slave i, $T_S(i)$, defined as the time required for the aforementioned successful data exchange, computed as

$$T_S(i) = \tau_{a,M} + \tau_{REQ} + \tau_{a,S} + \tau_{DATA} \quad (3)$$

where both the terms $\tau_{a,M}$ and $\tau_{a,S}$ account for the time required by the channel access procedure, as described in Sec. II. The time required for the successful delivery of the *request* packet, τ_{REQ}, is equal to the sum of T_{REQ} (the pure packet transmission time), a Short Inter Frame Space (SIFS) and T_{ACK} (a fixed ACK transmission time), in the best case. Finally, τ_{DATA} is relevant to the *data* packet, and behaves analogously to τ_{REQ}. All the terms in Eq. (3) are possibly random variables because of the effects of both CS and backoff.

A second meaningful indicator is the overall *cycle time* T_C, i.e. the time required for polling the whole set of slaves

$$T_C = \sum_{i=1}^{n} T_S(i) \quad (4)$$

this is a random variable as well, clearly upper bounded by the sum of the maximum service time $T_{S,max}$ of each slave.

B. Protocol parameters

IEEE 802.11n offers many different configuration options [6]. In this paper we assume that every node (master and slave) is equipped with two antennas. At the PHY layer, we selected the Modulation and Coding Scheme (MCS) 6, corresponding to a 64–QAM modulation with code rate 3/4. In addition, 40 MHz channels are employed and Space–Time Block Coding (STBC) is used to enhance robustness, providing a raw data rate of 121.5 Mbit/s. Other significant simulation parameters are reported in Tab. I.

TABLE I
SIMULATION PARAMETERS

Parameter	Description	Value
n	Number of slave nodes	5
$T_{S,max}$	Maximum polling duration	500 µs
L_{REQ}	Length of request packet	50 B
L_{DATA}	Length of data packet	10 B
$DIFS$	DIFS value in IEEE 802.11n	28 µs
$SIFS$	SIFS value in IEEE 802.11n	10 µs
T_{SLOT}	Slot time in IEEE 802.11n	9 µs
CW_{min}	Minimum value of contention window	15
CW_{max}	Maximum value of contention window	1023
P_{tx}	Default transmission power	100 mW

For the described configuration some meaningful time values introduced in previous subsection result as follows: T_{REQ} = 54 µs, T_{DATA} = 54 µs and T_{ACK} = 34 µs, yielding a total service time in ideal conditions of 252 µs. Finally, the maximum polling time $T_{S,max}$ has been set to 500 µs, to accommodate for delays that may be caused from channel access and retransmissions.

IV. Performance Assessment

A. Channel model

The typical impairments affecting a wireless medium can be mainly represented by channel fading and external interference, both being particularly considerable for wireless systems in the ISM band. For the aims of the present work, these impairments can be effectively modeled by means of fluctuations of the Signal-to-Noise Ratio (SNR) over the course of time.

Therefore, in our implementation, the channel state can assume one out of three values, namely *bad*, *average* or *good*, meaning that packets are corrupted with very high, medium and very low probability, respectively. The proposed model is actually a variation of the *Gilbert–Elliot* model, widely employed to emulate an industrial wireless channel, as validated by experimental measurements [7]. We introduced the third *average* state to model the likely occurrence of minor impairments often found in a real wireless channel, that actually degrades the transmission quality while still allowing some (small) success probability. The dynamic of the channel follows a stochastic process, in which the time spent in both the *bad* and the *average* states is modeled as a uniform random variable in the range [50, 100] µs, while the time spent in the *good* state is drawn from an exponential variable with mean of 200 µs. The aforementioned time intervals effectively provide a sort of worst–case for the channel behavior, being able to model both a (harsh) fading channel and the presence of an external interference. Indeed, the proposed channel model is characterized by a faster dynamic than the typical channel coherence time for the ISM band [8], and also it well represents the occurrence of packet transmissions due to a neighboring WLAN.

As a matter of fact, a very meaningful example in this direction is given by the scenario represented in Fig. 2. As can be seen, a simple external interfering network (nodes N and M) disturbs the node B of the industrial network under test.

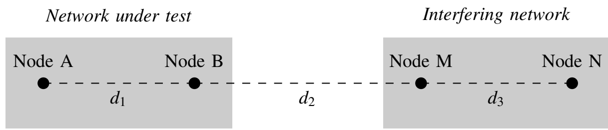

Fig. 2. A pictorial description of the network setup.

Specifically, the channel state perceived by node B is *bad* when node M transmits, *average* when node N transmits and *good* when the two nodes are silent.

Let us consider, for the setup in Fig. 2, that both d_1 and d_3 are equal to 5 m, d_2 is equal to 20 m, and the transmission power is the standard default given in Tab. I. Hence, we obtain that the SNR levels at node B are equivalent to 12 dB, 14 dB and 46 dB respectively, for the aforementioned three channel states, assuming the classic two–slope path loss model.

The Packet Error Rate (PER) corresponding to the SNR levels and packet sizes mentioned above has been obtained through experimental campaigns as described in [9], yielding the results reported in Tab. II. These values have been taken into account when simulating the system performance in the following subsection.

TABLE II
EXPERIMENTALLY MEASURED PER VALUES

Packet size	SNR		
	12 dB	14 dB	46 dB
10 Bytes	75.4%	16.1%	0%
50 Bytes	99.8%	75.4%	0%

B. Simulation results

The simulations have been carried out on a scenario similar to that presented in Fig. 2, in which the network under test comprised one master and five slave devices. The interference from the neighboring WLAN was supposed to affect either a single slave node, or the master node. The performance assessment has been carried out running more than 1000 network cycles.

Four different system configurations have been considered:

- Backoff + CS (legacy DCF).
- Only backoff (CCA procedure is ignored).
- Only CS (backoff procedure is ignored).
- No backoff + no CS (both CCA and backoff procedure are ignored).

The first set of proposed results is relevant to the number of failed polling operations, whose behavior is reported in Fig. 3.

As a general result, the deactivation of CS and backoff allows to sensibly reduce the number of missed polling operations. Furthermore, it can be observed that the worst case occurs when the interference is located on the master node since in this case it equally impacts on each polling operation.

The increased determinism that is achieved by disabling CS, backoff or both can be observed also in the Empirical CDF of service time for slave node 1 when it is disturbed by interference, reported in Fig. 4. In particular the ECDF in case of no backoff and no CS exhibits only three values

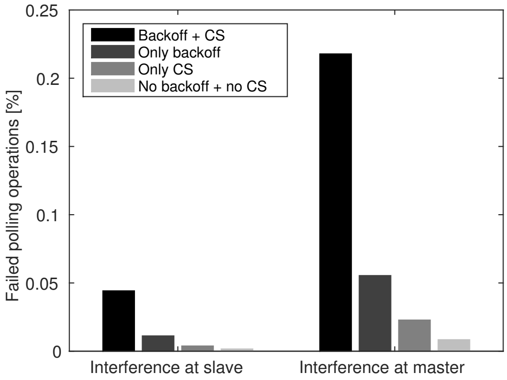

Fig. 3. Failed polling operations under different system configurations.

(252 µs, 349 µs and 446 µs), that are relevant to the cases of 1, 2 or 3 transmission attempts for the request packet: all other sources of randomness in the system have been removed.

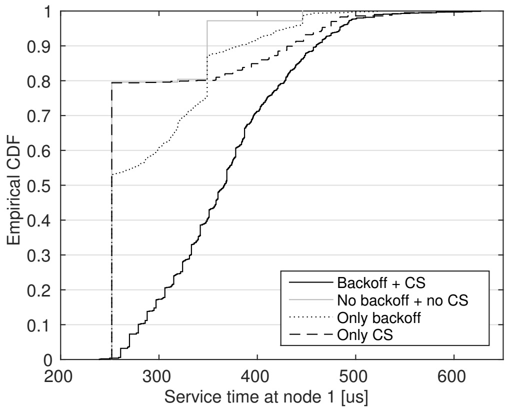

Fig. 4. Empirical CDF of service time for the node disturbed by interference.

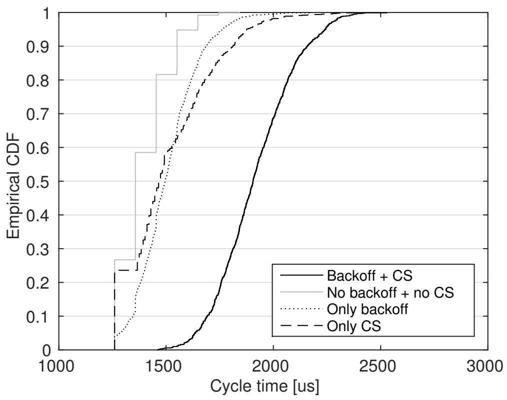

Fig. 5. Empirical CDF of cycle time with interference at master node.

A similar trend is visible in Fig. 5, where the ECDF of the cycle time in case of interference at master node is represented.

Finally, Tab. III shows the statistics (mean and standard deviation) for service time at node 1 when it is disturbed by external interference and for the cycle time when the master node is disturbed. The beneficial effects of disabling CS and/or backoff are evident.

TABLE III
SYSTEM PERFORMANCE UNDER DIFFERENT CONFIGURATIONS

Configuration	**Service time at node 1 (disturbed)**		**Cycle time (master node disturbed)**	
	Mean	**Std. Dev.**	**Mean**	**Std. Dev.**
Backoff + CS	368.6 µs	65.7 µs	1916.6 µs	178.9 µs
Only backoff	289.2 µs	77.2 µs	1507.5 µs	209.9 µs
Only CS	297.1 µs	59.1 µs	1506.1 µs	141.3 µs
No backoff + no CS	274.2 µs	46.7 µs	1395.1 µs	115.0 µs

V. CONCLUSIONS

The simulation outcomes discussed in this paper show that the deactivation of CS and backoff procedures allows to decrease the number of failed pollings, which reflects in an increased system reliability as well as in an improved timeliness.

A straightforward extension of this work is the validation of these results in an experimental setup, exploiting the possibilities offered by cited recent Soft–MAC devices. Moreover, the advantages provided by the proposed protocol tuning should be studied in combination with other mechanisms to improve the performance of an IEEE 802.11n network, such as for example the automatic rate adaptation, towards an ever growing utilization of this technology in the industrial environment.

REFERENCES

[1] *IEEE Std for Information technology–Telecommunications and information exchange between systems–Local and metropolitan area networks–Specific requirements. Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications*, IEEE Std., March 2012.

[2] A. Willig, "Recent and Emerging Topics in Wireless Industrial Communications: a Selection," *Industrial Informatics, IEEE Transactions on*, vol. 4, no. 2, pp. 102–124, May 2008.

[3] S. Vitturi, F. Tramarin, and L. Seno, "Industrial Wireless Networks: The Significance of Timeliness in Communication Systems," *Industrial Electronics Magazine, IEEE*, vol. 7, no. 2, pp. 40–51, June 2013.

[4] M. Bertocco, G. Gamba, and A. Sona, "Is CSMA/CA really efficient against interference in a wireless control system? An experimental answer," in *Proc. of IEEE ETFA*, Hamburg, Germany, September 2008.

[5] S. Petersen and S. Carlsen, "WirelessHART Versus ISA100.11a: The Format War Hits the Factory Floor," *Industrial Electronics Magazine, IEEE*, vol. 5, no. 4, pp. 23–34, December 2011.

[6] E. Perahia and R. Stacey, *Next Generation Wireless LANs: 802.11 n and 802.11 ac*. Cambridge university press, 2013.

[7] A. Willig, M. Kubisch, C. Hoene, and A. Wolisz, "Measurements of a Wireless Link in an Industrial Environment using an IEEE 802.11-Compliant Physical Layer," *Industrial Electronics, IEEE Transactions on*, vol. 49, no. 6, pp. 1265–1282, December 2002.

[8] T. Paul and T. Ogunfunmi, "Wireless LAN Comes of Age: Understanding the IEEE 802.11n Amendment," *Circuits and Systems Magazine, IEEE*, vol. 8, no. 1, pp. 28–54, First 2008.

[9] F. Tramarin, S. Vitturi, M. Luvisotto, and A. Zanella, "The IEEE 802.11n wireless LAN for real-time industrial communication," in *Proc. of IEEE World Conference on Factory Communication Systems, WFCS 2015*, to be published. May 2015.