

Analyzing and Boosting the Performance of Power-Line Communication Networks

Christina Vlachou*, Albert Banchs†, Julien Herzen*, Patrick Thiran*
* EPFL, Switzerland † IMDEA Networks Inst. and Univ. Carlos III of Madrid, Spain
* firstname.lastname@epfl.ch † banchs@it.uc3m.es

ABSTRACT

Power-line communications are employed in home networking to provide easy and high-throughput connectivity. IEEE 1901, the MAC protocol for power-line networks, employs a CSMA/CA protocol similar to that of 802.11, but is substantially more complex, which probably explains why little is known about its performance. One of the key differences between the two protocols is that whereas 802.11 only reacts upon collisions, 1901 also reacts upon several consecutive transmissions and thus can potentially achieve better performance by avoiding unnecessary collisions.

In this paper, we propose a model for the 1901 MAC. Our analysis reveals that the default configuration of 1901 does not fully exploit its potential and that its performance degrades with the number of stations. We derive analytically the optimal configuration parameters for 1901; this drastically improves throughput and achieves optimal performance *without* requiring the knowledge of the number of stations in the network. In contrast, to provide a similar performance, 802.11 requires knowing the number of contending stations, which is unfeasible for realistic traffic patterns. Our solution can be readily implemented by vendors, as it only consists in modifying existing MAC parameters. We corroborate our results with testbed measurements, unveiling a significant signaling overhead in 1901 implementations.

Keywords

Power-line communications; HomePlug AV; IEEE 1901; CSMA/CA; Performance Evaluation; Enhancement.

Categories and Subject Descriptors

C.2 [Computer-Communication Networks]: Network Protocols, Local and Wide-Area Networks.;
C.4 [Performance of Systems]: Modeling techniques

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
CoNEXT'14, December 2–5, 2014, Sydney, Australia.
Copyright 2014 ACM 978-1-4503-3279-8/14/12 ...\$15.00
<http://dx.doi.org/10.1145/2674005.2675016>

1. INTRODUCTION

Power-line communications (PLC) are developing rapidly. HomePlug, the leading alliance for PLC standardization, proposes different solutions for home automation and high data-rate local networks, with physical rates up to 1 Gbps. It is estimated that by 2013 over 120 million HomePlug devices had been shipped worldwide [3]. The annual growth rate for HomePlug AV devices (aimed for high data-rate networks) is expected to be about 30% until 2016 [4]. Although 95% of PLC devices follow the HomePlug specification [3], the MAC layer of this specification has received little attention so far in the research community, in contrast to the PHY layer. In particular, no work has investigated how far from optimality this MAC protocol is. This work explores and justifies the design choices adopted for the MAC layer of PLC networks.

As PLC technology is becoming an important component in home networks¹, residential buildings are expected to host networks with a high number of PLC stations. These PLC stations interfere with each other, because – in contrast to wireless technology that can employ different communication channels – PLC utilizes the entire available bandwidth (1.8 - 80 MHz) for communication. The recent HomePlug AV2 standard provides physical data rates of more than 1 Gbps, by employing efficient PHY layer techniques and parameters as well as small inter-frame timings [17]. However, in addition to data rate improvements, there is also the need for enhancements at the MAC layer, as an efficient MAC can further improve performance when many stations contend for the medium. Therefore, in this paper we focus on understanding the MAC layer dynamics, and on enhancing its performance building on this understanding.

Due to the shared nature of power lines, HomePlug devices employ a multiple-access scheme based on CSMA/CA that is specified by the IEEE 1901 standard² [7]. The 1901 CSMA/CA protocol bears some resemblance to the CSMA/CA mechanism employed by IEEE 802.11 that has been extensively studied in the literature (for instance, in [8, 14]). Nevertheless, 1901 differs from 802.11 in that its CSMA/CA mechanism is more complex, which makes its theoretical analysis challenging. In particular, in addition to using a backoff counter, it also uses a so-called *deferral counter*. The deferral counter significantly increases the state-space required to describe the backoff procedure, which contrasts

¹PLC is also employed for hybrid networks comprising multiple technologies such as Ethernet and Wireless LAN.

²All HomePlug specifications, including 1.0, AV, AV2, and GreenPhy, employ the same CSMA/CA process.

with the comparatively small state-space required to analyze 802.11 (see, e.g., the Markov chain used in [8]). As a result, despite the commercial success and massive adoption of PLC technologies, the analysis of the 1901 CSMA/CA protocol has received little attention.

From a general perspective, it turns out that 1901 implements an approach to contention resolution that is drastically different from the usual 802.11 CSMA/CA procedure. In particular, 802.11 can only react to contention (by doubling its contention window) after detecting a collision. In contrast, 1901 can already react when it senses the medium busy during a certain number of time slots (decided by the deferral counter). Such a protocol design has two distinct advantages over 802.11:

1. The contention window can be increased as many times as required to reach appropriate backoff durations without suffering any collision. In contrast, with 802.11 the contention window can only be doubled after a collision, and thus one or more collisions might occur before the contention window reaches the appropriate value. As a result, 1901 can reduce the channel time wasted in collisions, potentially leading to better performance.
2. By appropriately selecting the number of busy slots that trigger an increase of the contention window, we can adjust with fine granularity the level of contention that triggers a reaction. In contrast, this is not possible in 802.11, where contention is detected by the binary signal given by channel occupation: either the channel is busy upon a transmission attempt, which yields a collision, or it is not, and any finer refinements are not possible.

Following the above reasoning, it seems that 1901 can substantially outperform 802.11 if properly configured. However, as we will observe, the default configuration of 1901 does not achieve the level of efficiency that one would expect given these premises, and indeed performance can be largely improved. One important cause of the (relatively) poor performance of the protocol is the lack of an accurate and simple analysis that provides an insightful understanding of its dynamics and that can be used to configure the protocol appropriately.

Motivated by the above, in this paper we propose a general framework for modeling and enhancing the CSMA/CA process of 1901. The main contributions of this paper are as follows:

- First, we introduce a model that accurately captures 1901 performance while reducing very substantially the required space-state; this model comes in the form of a fixed-point equation which we show that admits a unique solution. Our model is accurate as long as the number of stations is not too small, due to the assumptions upon which we rely to derive a simple model. In [19], we analyze the impact of these assumptions and propose an alternative model that is more accurate; however, the model of [19] is more complex and does not provide insights on the optimality of the protocol, which is one of the main contributions of this paper.
- Second, we employ our model to compute a configuration that improves the throughput of the network. Our configuration consists in simply setting existing MAC parameters – such as the contention windows – to ap-

propriate values, and thus can be readily implemented by manufacturers of PLC devices. The proposed configuration provides drastic performance improvements.

- Third, we validate our analysis through simulation as well as by means of experimental results of a HomePlug AV testbed with 8 stations. Using our testbed, we unveil important details of 1901 implementations such as HomePlug AV, and we investigate the overhead of management messages exchanged between the PLC stations. This information can be useful for improving HomePlug AV2 implementations, as this specification aims at reducing various MAC layer overheads.

One of the most remarkable results of the paper is that, with the proposed configuration, 1901 provides a performance very close to that of an optimally configured MAC protocol without requiring to know the total number of contending stations N . This contrasts with similar methods for enhancing the CSMA/CA procedure of 802.11 networks that do require knowing N (see, e.g., [16]), which challenges their practicality in real deployments because N varies in time. Thus, with the configuration proposed here, the 1901 protocol represents an interesting step towards a practical and optimal MAC protocol.

The rest of this paper is organized as follows. In the next section, we present the 1901 backoff procedure. In Section 3, we review related work on models and enhancements for MAC protocols. We then present our model for 1901 in Section 4, and we propose enhancements for 1901 configuration in Section 5. Our model and proposed configurations are evaluated in Section 6. Finally, we give our concluding remarks in Section 7.

2. THE 1901 MAC MECHANISM

In this section, we present the relevant aspects of the 1901 CSMA/CA procedure [7]. We give insights on the requirements that established this intricate protocol, by comparing 1901 with 802.11.

The first HomePlug specification that included this CSMA/CA mechanism is HomePlug 1.0. HomePlug 1.0 employs a frame preamble comprising 7.5 OFDM symbols³. The slot duration was determined by the time required by a station to decide whether the medium is busy or idle (i.e., to detect a preamble transmission) [15], and it is equal to the duration of 7 HomePlug 1.0 symbols (i.e., 35.84 μ s). Although newer technologies have different symbol durations, the slot duration has remained the same for all HomePlug standards for backward compatibility. Observe that the slot duration is large compared to the one of 802.11 (which is 9 μ s for 802.11a/g/n/ac). In the next paragraphs, we explain the effect of the slot duration in the backoff process.

The backoff process of 1901 uses two counters: the backoff counter (BC) and the deferral counter (DC). In addition, there are four *backoff stages*⁴. We now discuss the common features of 1901 and 802.11, and we elaborate later on the deferral counter. When a new packet arrives for transmission, the station starts at backoff stage 0, and it draws the backoff counter BC uniformly at random in $\{0, \dots, CW_0 - 1\}$,

³HomePlug 1.0 provides rates up to 14 Mbps. Similar to all HomePlug standards, it employs an OFDM scheme at the PHY layer.

⁴In the standard [7], the backoff stage is determined by the so-called *backoff procedure counter* (BPC).

where CW_0 refers to the contention window used at backoff stage 0. Similarly to 802.11, BC is decreased by 1 at each time slot if the station senses the medium to be idle, and it is frozen when the medium is sensed busy. In case the medium is sensed busy, BC is also decreased by 1 once the medium is sensed idle again. When BC reaches 0, the station attempts to transmit the packet. Also similarly to 802.11, the station jumps to the next backoff stage if the transmission fails (unless it is already at the last backoff stage, in which case it re-enters this backoff stage). When entering backoff stage i , a station draws BC uniformly at random in $\{0, \dots, CW_i - 1\}$, where CW_i is the contention window at backoff stage i , and the process is repeated. For 802.11, the contention window is doubled between two successive backoff stages, and thus $CW_i = 2^i CW_0$. For 1901, CW_i depends on the priority level, and is given in Table 1. There are four priority classes in 1901, CA0 to CA3. CA0/CA1 priorities serve best-effort applications, and CA2/CA3 the delay-sensitive ones.

Now, when there are few contending stations (i.e., 1 or 2), or when the traffic load is very low, the time spent in backoff is a large overhead and increases as the contention window increases. Given the large slot duration of 1901, the average delay due to backoff ($\sim (CW_0 - 1)/2$ time slots) can be reduced when there are few contending stations – that is, small collision likelihood – by choosing a small minimum contention window, e.g., $CW_0 = 8$, as specified for 1901 (Table 1). However, as expected, small contention windows yield higher collision probabilities when the number of stations increases or when the traffic load rises. The deferral counter DC was introduced as a countermeasure in the CSMA/CA process of 1901, to reduce collisions induced by small contention windows. This is achieved by triggering a redraw of the backoff counter BC before the station attempts a transmission.

Priority class:	CA0/CA1		CA2/CA3	
backoff stage i	CW_i	d_i	CW_i	d_i
0	8	0	8	0
1	16	1	16	1
2	32	3	16	3
3	64	15	32	15

Table 1: IEEE 1901 parameters for the contention windows CW_i and the initial values d_i of deferral counter DC , for each backoff stage i .

The main difference between 1901 and 802.11 is the introduction of DC in 1901, which allows a station to enter a higher backoff stage even if it did not attempt a transmission. The mechanism to decide when this occurs works as follows. When entering backoff stage i , DC is set at an initial DC value d_i , where d_i is given in Table 1 for each i . After having sensed the medium busy, a station decreases DC by 1 (in addition to BC). If the medium is sensed busy and $DC = 0$, then the station jumps to the next backoff stage (or re-enters the last backoff stage, if it is already at this stage), and it re-draws BC without attempting a transmission. Figure 1 shows an example of such a backoff process.

We now discuss the effect of the deferral counter on jitter, because this justifies the choice of rest of the parameters in 1901. As we have seen above, the deferral counter manages to reduce collisions. It can however introduce high jitter. The transmitting station has an advantage against the other stations because its CW is maintained at its initial value CW_0 , whereas the other stations might increase their CW

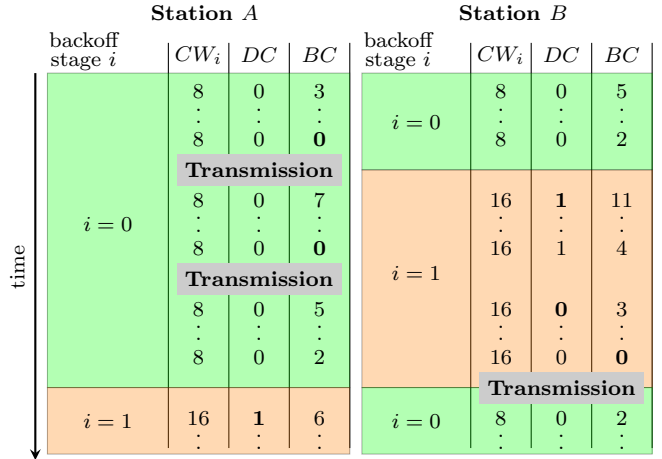


Figure 1: Evolution of the backoff process with 2 saturated stations A and B. Initially, both stations are at backoff stage 0. A transmits twice consecutively. Note the change in i when a station senses the medium busy and has $DC = 0$.

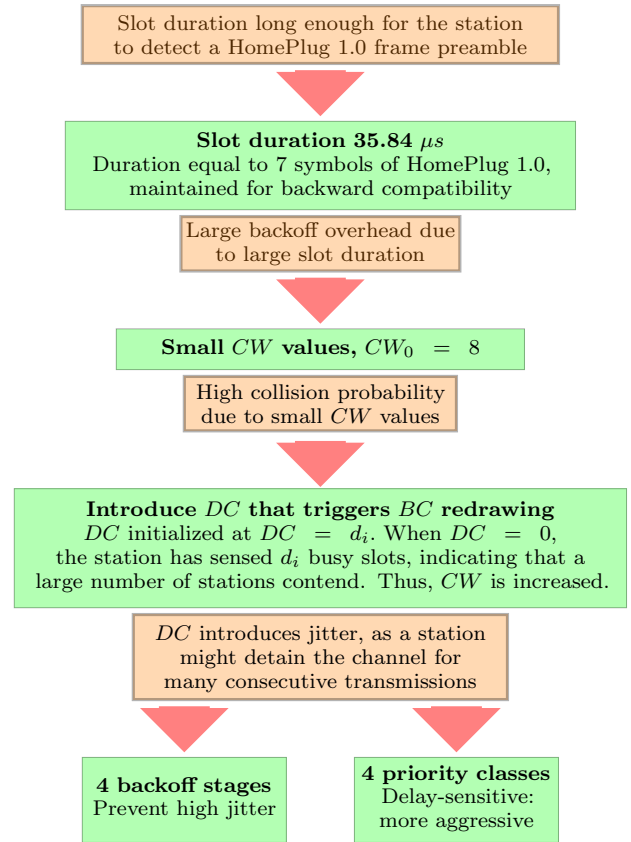


Figure 2: A discussion of the design choices of the 1901 MAC.

when $DC = 0$, as explained earlier. As a result, a station can grasp the channel for multiple consecutive transmissions until it releases the channel, and then wait for a long sequence of other station transmissions until obtaining the token again, which yields high jitter. To reduce the CW imbalance between the transmitting station and the others, 1901 employs a total of 4 backoff stages. Moreover, there are two

different configurations for best-effort and delay-sensitive applications (see Table 1). The delay-sensitive class employs smaller contention windows, and the contention window is not doubled between backoff stages 1 and 2. This improves jitter, but yields a higher collision probability, i.e., lower throughput, compared to CA0/CA1 class. In this paper, we focus on modeling and enhancing 1901 performance for best-effort classes (targeting high throughput), whereas in [19] and [21], we investigate the short-term fairness aspects of 1901, related to jitter. Figure 2 summarizes the insights concerning the MAC parameters described in this section.

Let us underline a major consequence of the difference between 1901 and 802.11, which is due to the introduction of the deferral counter in 1901. A 1901 station can redraw its backoff counter *without* experiencing a collision, if its deferral counter is 0 and it senses the medium busy. This results in 1901 having the two distinct features that we mentioned earlier, and which offer new opportunities for access control:

- With 1901, the contention window can be increased without suffering from a collision, hence a station can adapt its transmission behavior to the level of contention in the network without wasting channel time in collisions.
- By appropriately configuring its deferral counter, 1901 can optimally adjust its level of reaction to contention with a fine granularity, which is contrary to 802.11, where the level of reaction cannot be tuned.

Figure 3 provides an evidence of the 1901 potential. We investigate the time evolution of the protocol dynamics and compare 1901 with 802.11 for two different values for the number of stations in the network, $N = 2$ and $N = 15$. Observe that for $N = 2$, 802.11 has large average CW only after consecutive collisions, whereas 1901 adapts CW even after successful transmission attempts. The 1901 average CW increases or remains constant as a given station keeps transmitting, and decreases only when the other station successfully wins the medium. Finally, notice that for $N = 15$, 1901 adapts CW before a collision, thus it wastes less time in collisions and the average CW converges much faster to the steady state. Hence, under dynamically changing traffic, 1901 can adjust faster CW to the load demand.

In Section 5, we show that the two above 1901 features make it possible to configure 1901 so that it efficiently adjusts to the level of congestion in the network, *without knowing the number of active stations*. In contrast, although many approaches to efficiently configure and optimize 802.11 have been proposed in the literature, these are limited to a given (known) number of active stations and, to the best of our knowledge, there exists no such 802.11 configuration that works when the number of stations is unknown.

3. RELATED WORK

Analytical models of 802.11 can provide insights for analyzing 1901. Indeed, 802.11 can be viewed as a simplified version of 1901 in which the deferral counter never expires. We review here major studies on 802.11 and some of the few studies on 1901.

Models of IEEE 802.11.

A large number of performance evaluation models have been proposed for 802.11 (e.g., [8, 14]). In particular, the model proposed by Bianchi in [8] for single contention domain networks is very popular. This work models the backoff




Figure 3: Time evolution of the contention window CW averaged over all stations in the network (left axis) and the binary outcome of transmission attempts (right axis), i.e. success (S) or collision (C), in simulation. Our simulator is validated experimentally in Section 6.2. All stations are saturated and start at backoff stage 0. 1901 is simulated with CA1 class, and 802.11 with $CW_0 = 16$ and 7 backoff stages (i.e., 802.11a/g/n), and we use the timings specified in Section 6.

process of 802.11 using a discrete time Markov chain. Importantly, Bianchi assumes that the backoff processes of the stations are independent, which is known as the *decoupling assumption*. In this case, the collision probability γ experienced by all stations is time-invariant, and can be found by solving a fixed-point equation that depends on the protocol parameters.

Kumar et al. [14] study the backoff process of 802.11 under the decoupling assumption using renewal theory. We employ a similar method for finding a fixed-point equation for 1901. However, due to the increased complexity of 1901, proving the uniqueness of the solution of this fixed-point equation is significantly more challenging than for 802.11 and it is the object of Theorem 1.

Models of IEEE 1901.

To the best of our knowledge, the only works analyzing the backoff process of 1901 are [12, 11, 20, 19]. The authors in [12] propose a model similar to Bianchi's model

for 802.11 [8]. However, the paper does not provide the corresponding fixed-point equation for the collision probability, due to the increased complexity of the Markov chain introduced by the deferral counter. To compute the collision probability in this case, a costly system of more than a thousand non-linear equations has to be solved. Moreover, it has not been investigated whether this system of equations has a unique solution. The work in [11] discusses the validity of the decoupling assumption and simplifies the model in [12] in the case the stations are not saturated. Here again, the authors do not investigate the uniqueness of the solution(s) of their model. Our short paper [20] includes the model of Section 4, but without any proofs or enhancements.

We propose a model which is strictly equivalent to the model of [12] in terms of modeling performance. The key difference is that, in our case, the collision probability can be obtained by solving a single fixed-point equation. In this sense, our model can be seen as a drastic simplification of [12], and this simple form enables us to derive some efficient configuration parameters for 1901. In [19], we introduce an alternative model that does not rely on the decoupling assumption. However, this model is intricate and does not yield insights for enhancing performance.

Enhancements of IEEE 802.11.

There is a large body of work introducing enhancements for the 802.11 CSMA/CA procedure. In particular, Bianchi [8] computes τ_{opt} , the transmission probability that achieves maximum throughput. τ_{opt} is obtained by tuning the contention window, and it is a function of the number of contending stations N . Typically, N is unknown and varies, and practical implementations of such optimal configurations need to use some estimation techniques.

To apply Bianchi’s analysis, several attempts have been made to estimate the number of contending stations [9, 16]. These methods typically rely on measuring the collision probability or the channel activity and on estimating N periodically. The main disadvantage of such approaches is that they introduce more complexity at the MAC layer, making them difficult to implement in practice. The improvements that we propose for 1901 do not require the knowledge of N , nor do they require that the CSMA/CA process itself be changed. Instead, they can be simply implemented, as they only consist in judiciously choosing the backoff parameters CW_i and d_i (and, optionally, the number of backoff stages m). In Section 6, we observe that they provide drastic performance improvements compared to the default 1901 and offer similar efficiency levels as an optimally-configured 802.11 network, but without knowing N .

Enhancements of IEEE 1901.

There are not many studies on enhancements of 1901 MAC. The authors in [10] propose a constant d_i equal to 0 for all backoff stages i , which means that whenever a station senses the medium busy it doubles CW . This technique decreases the collision probability, but it yields the most extreme case of 1901 unfairness; for small N , always doubling CW leads idle stations to have fewer chances of accessing the channel, compared to a station that just transmitted successfully (and whose CW is minimal). This yields a high variance of delay, as explained in Section 2. Moreover, the authors do not use an analytical model and do not investigate the dif-

ferent tradeoffs involved in the parameter setting to find the optimal configurations.

Another modification of 1901 is proposed in [13]. The authors provide a mechanism that keeps d_i and CW constant, but where CW depends on the number of stations N . Requiring the knowledge of N is impractical.

4. ANALYSIS

In this section, we introduce our model for the 1901 CSMA/CA process. We analyze the protocol under the following assumptions. First, there are N saturated stations in the network. Second, all stations belong to a single contention domain. Third, there is no packet loss or errors due to the physical layer. Therefore, transmission failures are due only to collisions. Fourth, the stations never discard a packet until it is successfully transmitted, hence the retry limit is infinite⁵. Finally, the 1901 standard specifies that only the stations belonging to the highest contending priority class participate in the backoff process⁶. We follow this property of 1901, and we assume that all the contending stations are in the same class and use the same set of parameters.

We now turn our attention to the modeling assumptions, which assert that the backoff process of a station is independent of the aggregate attempt process of the other $N - 1$ stations, and are referred to as the *decoupling assumption*. In detail, this term refers to the following set of assumptions (see e.g., [8, 14]):

- [A1] The stations attempt a transmission in each time slot with a *constant* transmission probability referred as the average attempt rate τ .
- [A2] Given a tagged station, the probability that at least one of the other stations transmits is fixed (state independent) and is denoted by γ . Therefore, transmission attempts experience a fixed collision probability γ . Furthermore, a station with $BC \neq 0$ senses the medium busy at any time slot with probability γ , independently of the medium status at previous slots.

The above assumption, which considers that the collision probability of a station is independent of its state, has been proven to be accurate in 802.11. However, in [19, 21], it has been shown that a PLC network with few stations can be short-term unfair for some configurations, which implies that there is some coupling between the stations. Indeed, for such configurations, when a station is in a state with large backoff stage, the probability that some other station is at small backoff stage state is high, and as a result its collision probability is larger in that state, which contradicts the decoupling assumption. The simulation results of Section 6.2 confirm that, while the model has reduced accuracy for some configurations when there are few stations in the network, accuracy is high as long as the number of stations is not too small. To improve the 1901 modeling accuracy for small N , a theoretical framework without relying on the decoupling

⁵IEEE 1901 does not specify any retry limit, contrary to 802.11. However, there exists a vendor specific timeout on the frame transmission. This timeout is 2.5 s for the devices of our testbed presented in Section 6, and is very large compared to the maximum frame duration (2.5 ms), which is in line with our hypothesis.

⁶In practice, the contending class is decided using a simple system of busy tones called *priority resolution symbols*.

assumption is applied in [19]; however the resulting model is more complex and challenging. Here, our work invokes the decoupling assumption, which yields a simple model that is accurate for large N and can be employed for performance enhancements. As shown in the results of Section 6.3, our model indeed enables us to provide configurations that boost significantly the 1901 performance.

Our analysis requires to compute the expected number of time slots (a time slot can be idle or can contain a transmission) spent by a station at backoff stage i . Let k denote the value of BC drawn uniformly at random in $\{0, \dots, CW_i - 1\}$, when the station enters stage i . If the station is running 802.11, BC always expires and the station spends exactly $(k + 1)$ time slots in stage i , the $(k + 1)$ th slot being used for transmission, because a station leaves backoff stage i only when it attempts a transmission. In contrast, in 1901 a station might leave backoff stage i either because of a transmission attempt, when BC expires (like in 802.11), or because it has sensed the medium busy $d_i + 1$ times, before BC has expired. In this latter case, the station spends a number of slots at backoff stage i , which is at least $d_i + 1$ and no more than k . More precisely, the station spends $d_i + 1 \leq j \leq k$ slots when it senses the medium busy d_i times in any of the first $j - 1$ slots, and in the j th slot.

Let us write bc_i for the expected number of time slots spent by a station at backoff stage i . To compute bc_i we need to evaluate the probability of the events that (i) a station attempts a transmission or (ii) senses the medium busy $d_i + 1$ times within the k slots. Let T be the random variable describing the number of slots during which the medium is sensed busy among the k slots. Because of [A2], T follows the binomial distribution $\text{Bin}(k, \gamma)$. Now, let x_k^i be the probability that a station at backoff stage i jumps to the next stage $i + 1$ in k or fewer time slots due to event (ii). Then,

$$x_k^i = \mathbb{P}(T > d_i) = \sum_{j=d_i+1}^k \binom{k}{j} \gamma^j (1-\gamma)^{k-j}. \quad (1)$$

We can compute bc_i as a function of γ via x_k^i . We distinguish two cases on k . First, if $k > d_i$, then event (i) occurs with probability $(1 - x_k^i)$, in which case the station spends $(k + 1)$ slots in stage i . Now, event (ii) occurs with probability x_k^i . More precisely, event (ii) occurs at slot j , with $d_i + 1 \leq j \leq k$, with probability $(x_j^i - x_{j-1}^i)$ ⁷, in which case the station spends j slots in backoff stage i . Second, if $k \leq d_i$, then event (ii) cannot happen. Thus, the backoff counter expires, event (i) always takes place, and the station spends $(k + 1)$ time slots in stage i . By summing all the possible cases described above, it follows that bc_i is given by

$$bc_i = \frac{1}{CW_i} \sum_{k=d_i+1}^{CW_i-1} \left[(k+1)(1-x_k^i) + \sum_{j=d_i+1}^k j(x_j^i - x_{j-1}^i) \right] + \frac{(d_i+1)(d_i+2)}{2CW_i}. \quad (2)$$

Let R be the random variable describing the number of transmission attempts experienced by a successfully transmitted packet. Similarly, let X be the random variable describing the total number of slots spent in backoff for a successfully transmitted packet. Then from the renewal-reward

⁷Observe that $(x_j^i - x_{j-1}^i)$ is the difference of two CCDFs and denotes the probability that (ii) happens exactly at slot j .

theorem, with R being the reward and X being renewal lifetimes [14], the average attempt rate is given by

$$\tau = \frac{\mathbb{E}[R]}{\mathbb{E}[X]}. \quad (3)$$

Let t_i be the probability that a station at backoff stage i ends this backoff stage by attempting a transmission. Similarly to (2), t_i can be computed as

$$t_i = \frac{d_i + 1}{CW_i} + \sum_{k=d_i+1}^{CW_i-1} \frac{1}{CW_i} (1 - x_k^i). \quad (4)$$

Finally, let us introduce $s_i = (1 - \gamma)t_i$, which is the probability that a station at backoff stage i ends this stage after a successful transmission. We are now ready to compute $\mathbb{E}[R]$ and $\mathbb{E}[X]$, and thus τ by (3).

LEMMA 1. *The expected number of slots spent in backoff per successfully transmitted packet is*

$$\mathbb{E}[X] = \sum_{i=0}^{m-2} bc_i \prod_{j=0}^{i-1} (1 - s_j) + \prod_{i=0}^{m-2} (1 - s_i) \frac{bc_{m-1}}{s_{m-1}}.$$

PROOF. Let X_i be the random variable describing the number of slots that a station starting in stage i spends in backoff before transmitting its current packet successfully. Using the law of total probability, we have

$$\begin{aligned} \mathbb{E}[X] &= \mathbb{E}[X_0] = s_0 bc_0 + (1 - s_0)(bc_0 + \mathbb{E}[X_1]) \\ &= bc_0 + (1 - s_0)\mathbb{E}[X_1]. \end{aligned}$$

Repeating recursively for $\mathbb{E}[X_1]$, we have

$$\begin{aligned} \mathbb{E}[X_0] &= bc_0 + \\ &(1 - s_0)(bc_1 + (1 - s_1)(bc_2 + \dots (1 - s_{m-2})\mathbb{E}[X_{m-1}])). \end{aligned} \quad (5)$$

Now, $\mathbb{E}[X_{m-1}]$ is given by

$$\mathbb{E}[X_{m-1}] = bc_{m-1} + (1 - s_{m-1})\mathbb{E}[X_{m-1}].$$

Solving for $\mathbb{E}[X_{m-1}]$, we obtain $\mathbb{E}[X_{m-1}] = bc_{m-1}/s_{m-1}$. Plugging this expression into (5) concludes the proof. \square

LEMMA 2. *The expected number of transmission attempts per successfully transmitted packet is*

$$\mathbb{E}[R] = \sum_{i=0}^{m-2} t_i \prod_{j=0}^{i-1} (1 - s_j) + \prod_{i=0}^{m-2} (1 - s_i) \frac{t_{m-1}}{s_{m-1}} = \frac{1}{1 - \gamma}.$$

PROOF. Similar to the proof of Lemma 1. \square

Since γ is the probability that at least one other station transmits, it can be expressed as a function of τ :

$$\gamma = \Gamma(\tau) = 1 - (1 - \tau)^{N-1},$$

where $\tau = \mathbb{E}[R]/\mathbb{E}[X]$ is also a function of γ , which we denote by $G(\gamma)$. The composition of the two functions $\tau = G(\gamma)$ and $\gamma = \Gamma(\tau)$ yields the fixed-point equation for the collision probability

$$\gamma = \Gamma(G(\gamma)). \quad (6)$$

Theorem 1 below establishes the uniqueness of the solution of this fixed-point equation, for configurations satisfying $CW_{i+1} \geq 2CW_i - d_i - 1$, $0 \leq i < m - 1$. Note that, from

Table 1, these constraints on CW_i are compliant with the standard, except for the class CA2/CA3 at backoff stage $i = 1$. While our numerical experiments seem to indicate that uniqueness of the solution holds for a wider range of configurations, we leave the extension of this result to other configurations for future work.

THEOREM 1. $\Gamma(G(\gamma)) : [0, 1] \rightarrow [0, 1]$ has a unique fixed-point if $CW_{i+1} \geq 2CW_i - d_i - 1$, $0 \leq i < m - 1$.

PROOF. By Brouwer's fixed-point theorem, since $\Gamma(G(\gamma))$ is a continuous function, there exists a fixed-point in $[0, 1]$. Furthermore, if $\Gamma(G(\gamma))$ is monotone, this fixed-point is unique. As $\Gamma(\tau)$ is non-decreasing in γ , it is thus sufficient to show that $G(\gamma)$ is monotone in γ .

Let $Q(\gamma) = (1 - \gamma)\mathbb{E}[X]$. Then, using Lemma 2, we have $G(\gamma) = 1/Q(\gamma)$. Now, $G(\gamma)$ is non-increasing in γ if and only if $Q(\gamma)$ is non-decreasing in γ . In the following, we show that $Q(\gamma)$ is non-decreasing.

Let B_i denote the expected number of backoff slots between two transmission attempts of a station that always stays at backoff stage i . At backoff stage i , according to Lemma 3 in Appendix, B_i is given by $B_i = bc_i/t_i - 1$. Replacing bc_i with $t_i(B_i + 1)$ in the expression for $\mathbb{E}[X]$, $Q(\gamma)$ can be recast as

$$Q(\gamma) = \sum_{i=0}^{m-2} s_i (B_i + 1) \prod_{j=0}^{i-1} (1 - s_j) + \prod_{i=0}^{m-2} (1 - s_i) (B_{m-1} + 1).$$

The derivative of $Q(\gamma)$ with respect to γ is given by

$$\begin{aligned} \frac{dQ}{d\gamma} &= \sum_{i=0}^{m-2} s_i \frac{dB_i}{d\gamma} \prod_{j=0}^{i-1} (1 - s_j) + \prod_{i=0}^{m-2} (1 - s_i) \frac{dB_{m-1}}{d\gamma} \\ &\quad - \sum_{i=0}^{m-2} \frac{ds_i}{d\gamma} \left[\sum_{j=i+1}^{m-2} \left((B_j + 1) s_j \frac{\prod_{k=0}^{j-1} (1 - s_k)}{1 - s_i} \right) \right. \\ &\quad \left. - (B_i + 1) \prod_{j=0}^{i-1} (1 - s_j) + \frac{\prod_{j=0}^{m-2} (1 - s_j)}{1 - s_i} (B_{m-1} + 1) \right]. \end{aligned} \quad (7)$$

We now show that $Q(\gamma)$ is non-decreasing. First, in Lemma 3 in Appendix we prove that $dB_i/d\gamma > 0$. Thus, the first two terms in (7) are positive and it follows that

$$\begin{aligned} \frac{dQ}{d\gamma} &> - \sum_{i=0}^{m-2} \frac{ds_i}{d\gamma} \left[\sum_{j=i+1}^{m-2} \left((B_j + 1) s_j \frac{\prod_{k=0}^{j-1} (1 - s_k)}{1 - s_i} \right) \right. \\ &\quad \left. - (B_i + 1) \prod_{j=0}^{i-1} (1 - s_j) + \frac{\prod_{j=0}^{m-2} (1 - s_j)}{1 - s_i} (B_{m-1} + 1) \right]. \end{aligned} \quad (8)$$

Now, in Lemma 3, we show that x_k^i is increasing with γ (see (25)). Thus, s_i is decreasing with γ , because, using (4),

$$\frac{ds_i}{d\gamma} = -t_i - (1 - \gamma) \frac{\sum_{k=d_i+1}^{CW_i-1} dx_k^i/d\gamma}{CW_i} < 0.$$

Also, by Corollary 1 in Appendix, B_i is a strictly increasing sequence of i . Using these two properties (i.e., $ds_i/d\gamma < 0$, $B_{i+1} > B_i$) for $i = m - 2$ in (8), we have

$$\begin{aligned} \frac{dQ}{d\gamma} &> - \sum_{i=0}^{m-3} \frac{ds_i}{d\gamma} \frac{1}{1 - s_i} \left[\sum_{j=i+1}^{m-2} \left((B_j + 1) s_j \prod_{k=0}^{j-1} (1 - s_k) \right) \right. \\ &\quad \left. - (B_i + 1) \prod_{j=0}^i (1 - s_j) + \prod_{j=0}^{m-2} (1 - s_j) (B_{m-1} + 1) \right]. \end{aligned} \quad (9)$$

Using the same properties again in (9) and then, rearranging the factors in the products involving the s_i 's, we have

$$\begin{aligned} \frac{dQ}{d\gamma} &> - \sum_{i=0}^{m-3} \frac{ds_i}{d\gamma} \frac{1}{1 - s_i} \left[(B_i + 1) \prod_{j=i+1}^{m-2} \left(s_j \prod_{k=0}^{j-1} (1 - s_k) \right) \right. \\ &\quad \left. - (B_i + 1) \prod_{j=0}^i (1 - s_j) + \prod_{j=0}^{m-2} (1 - s_j) (B_{m-1} + 1) \right] \\ &= - \sum_{i=0}^{m-3} \frac{ds_i}{d\gamma} \frac{\prod_{j=0}^{m-2} (1 - s_j)}{1 - s_i} (B_{m-1} - B_i) \geq 0, \end{aligned}$$

with equality at $\gamma = 0$. This completes the proof. \square

We now explain how to obtain actual throughput figures from our model. Solving the fixed-point equation (6) gives the value for γ , which can be used to obtain τ . We can therefore compute p_s and p_e , the probability that a slot contains a successful transmission, or that it is empty, respectively. We have $p_s = N\tau(1 - \tau)^{N-1}$, and $p_e = (1 - \tau)^N$. Let p_c denote the probability that a slot contains a collision. We have $p_c = 1 - p_e - p_s$. We now have enough information to compute the normalized throughput S of the network as

$$S = \frac{p_s D}{p_s T_s + p_c T_c + p_e \sigma}, \quad (10)$$

where D is the frame duration, T_s is the duration of a successful transmission, T_c is the duration of a collision, and σ is the time slot duration.

5. ENHANCEMENTS OF THE IEEE 1901 MAC FOR HIGH THROUGHPUT

As highlighted in Section 2, 1901 has the advantage over 802.11 in that it reacts pro-actively to collisions. However, it can be configured more efficiently for high throughput. We employ our model to examine 1901 configurations that perform close to an optimal MAC protocol. Our model adopts the same assumptions as other models investigating optimal versions of 802.11 [8].

In the following, we propose a configuration of the 1901 MAC parameters that drives the system to the optimal collision probability γ_{opt} . Specifically, γ_{opt} is the collision probability that maximizes the achievable throughput of the network, as computed by (10). The optimal transmission probability τ_{opt} is the one that achieves γ_{opt} . It is computed in [8] as

$$\tau_{opt} \approx \frac{1}{N} \sqrt{\frac{2\sigma}{T_c}}. \quad (11)$$

Using the above, the optimal collision probability is

$$\gamma_{opt} = 1 - (1 - \tau_{opt})^{N-1} = 1 - \left(1 - \frac{1}{N} \sqrt{\frac{2\sigma}{T_c}} \right)^{N-1} \quad (12)$$

which, as N gets large, can be approximated⁸ by

$$\gamma_{opt} \approx 1 - e^{-\sqrt{\frac{2\sigma}{T_c}}}. \quad (13)$$

Following the analysis in Section 4, γ can be obtained from the following fixed-point equation:

$$1 - (1 - \gamma)^{1/(N-1)} = \frac{\mathbb{E}[R]}{\mathbb{E}[X]}. \quad (14)$$

⁸For the 1901 parameters (shown in Section 6), the approximation is good for $N \geq 3$.

We can re-write the right-hand side of (14) as

$$\frac{\mathbb{E}[R]}{\mathbb{E}[X]} = \frac{1}{(1-\gamma)^2 \sum_{k=0}^{\infty} \gamma^k BC_k}, \quad (15)$$

where BC_k is the average number of slots spent in backoff at the k -th transmission attempt. BC_k is the average number of slots in backoff stages where the station does not transmit plus the average number of slots in the backoff stage where the k -th transmission occurs. To compute BC_k , we make the following *key approximation*: We assume that if the expectation of the backoff counter at a given backoff stage is smaller than the average number of slots required to reach the deferral counter, then the backoff counter always expires before the deferral counter expires; conversely, we assume that the deferral counter expires otherwise.

The above approximation is necessary in order to come up with a configuration that is close to optimal for any number of stations. The key idea behind the proposed configuration is as follows. The expected value of the backoff counter does not depend on γ , whereas the average number of slots required for the deferral counter to expire is decreasing with γ . We set the expected value of the backoff counter to be equal to the average number of slots required for the deferral counter to expire at $\gamma = \gamma_{opt}$. Then, when $\gamma < \gamma_{opt}$, a station always transmits before the deferral counter expires (according to this approximation), thus increasing γ . Similarly, if $\gamma > \gamma_{opt}$, a station never transmits, and γ decreases. As a result, this configuration drives the system to operate at $\gamma = \gamma_{opt}$, which (as seen above) corresponds to the optimal operation independent of the number of stations. While there are other configurations that have been proposed to provide optimal performance [16, 13], these require to know the number of stations and thus, need to be adaptively adjusted as the number of active stations in the network changes.

At the k -th transmission attempt, there is a unique backoff stage where the transmission occurs, and there might be some backoff stages where the deferral counter expires. Let l be the backoff stage where the transmission occurs, at the k -th transmission attempt. Then, there is a set of backoff stages S_k for which the deferral counter expires. With our approximation mentioned above, BC_k is given by

$$BC_k \approx \sum_{i \in S_k} D_i + C_l, \quad (16)$$

where D_i is the average number of slots required for the deferral counter to expire at backoff stage i , C_l is the expectation of the backoff counter at stage l , and S_k is the set of backoff stages for which $D_i > C_i$ at the k^{th} transmission attempt. For simplicity, here we restrict ourselves to configurations where $CW_i = 2^i CW_{min}$ for each backoff stage $0 \leq i \leq m-1$. At stage i , with a configuration CW_i and d_i , we have $C_i = (2^i CW_{min} + 1)/2$, and $D_i \approx (d_i + 1)/\gamma$, because D_i can be expressed as the average number of attempts to achieve $d_i + 1$ successes of an event with success probability γ .

Now, we address the issue of finding the configuration of CW_{min} , m and d_i to ensure that the solution of (14) corresponds to the optimal collision probability γ_{opt} of (13). To do this, we proceed as follows. First, we set d_i such that $D_i = C_i$ for $\gamma = \gamma_{opt}$. The rationale behind this is as follows: to force the solution of (14) to be at γ_{opt} , we need that the right-hand side of (14) to sharply decrease from a large to

a small value at this point, so that the two curves cross at $\gamma = \gamma_{opt}$ for all N . Since increasing γ decreases D_i , and not C_i , if we set D_i such that with γ_{opt} , we have $D_i = C_i$, then for a smaller γ , D_i will be larger and the stations will be more likely to transmit. Otherwise, they will be more likely to defer. Indeed, if we force $D_i = C_i$ for $\gamma = \gamma_{opt}$, we obtain that $\mathbb{E}[R]/\mathbb{E}[X]$ sharply decreases from a maximum to a minimum value at $\gamma = \gamma_{opt}$, i.e.,

$$\frac{\mathbb{E}[R]}{\mathbb{E}[X]} \approx \begin{cases} \tau_{max}, & \gamma \leq \gamma_{opt}, \\ \tau_{min}, & \gamma > \gamma_{opt} \end{cases} \quad (17)$$

where τ_{min} is the value of $\mathbb{E}[R]/\mathbb{E}[X]$ given by our previous analysis of Section 4, which is minimized at $\gamma = 1$, and τ_{max} is the corresponding maximum value, at $\gamma = 0$. An illustration of our analysis is given in Figure 4.

Following the above, we need to enforce $D_i = C_i$, $0 \leq i \leq m-1$, which gives the configuration of d_i ,

$$d_i = \left\lceil \gamma_{opt} \frac{2^i CW_{min} + 1}{2} - 1 \right\rceil, \quad (18)$$

where the ceiling is used to avoid negative values of d_0 .

The remaining challenge is to configure CW_{min} , and m . To ensure that the curve $1 - (1-\gamma)^{1/(N-1)}$ crosses $\mathbb{E}[R]/\mathbb{E}[X]$ at $\gamma = \gamma_{opt}$, we need that τ_{max} be sufficiently large so that even with small N the curve is crossed at this point and, conversely, that τ_{min} be sufficiently small for large N .




Figure 4: The analysis used for our enhancements. τ is plotted as a function of γ . The fixed-point is the intersection of $1 - (1-\gamma)^{1/(N-1)}$ and $G(\gamma) = \mathbb{E}[R]/\mathbb{E}[X]$. We plot $1 - (1-\gamma)^{1/(N-1)}$ for various number of stations N , and show $G(\gamma)$ for both the 1901 CA1 configuration, and our proposed enhancement from (18), (19) for $m = 6$. Finally, we also show γ_{opt} (obtained from (13)), and the corresponding τ_{opt} (obtained from (11)). Observe that the default 1901 configurations yield fixed-points having large collision probabilities. By contrast, our proposed 1901 configurations yield fixed-points that are much closer to τ_{opt} .

To guarantee that τ_{max} is sufficiently large even for $N = 2$, the CW at backoff stage 0, CW_{min} , needs to be as small as the optimal CW for $N = 2$. Accordingly, we set CW_{min} equal to the optimal CW value for $N = 2$ as if CW was constant, which gives $\gamma_{opt} = \tau = 2/(CW_{min} + 1)$, i.e.,

$$CW_{min} = \left\lfloor \frac{2}{1 - e^{-\sqrt{\frac{2\sigma}{T_c}}}} - 1 \right\rfloor. \quad (19)$$

Finally, we need to ensure that τ_{min} is sufficiently small for a large number of stations. Therefore, m should be sufficiently large. Following this, we choose the configuration $m = 6$; with this setting for m , the resulting τ_{min} is sufficiently small to ensure that $\tau = 1 - (1 - \gamma)^{1/(N-1)}$ crosses $\mathbb{E}[R]/\mathbb{E}[X]$ close to $\gamma = \gamma_{opt}$ even for N as large as 30.

Figure 4 shows the point of operation resulting from our configuration as well as the default one, given by the intersection between the curves $1 - (1 - \gamma)^{1/(N-1)}$ and $\mathbb{E}[R]/\mathbb{E}[X]$. While according to our approximation $\mathbb{E}[R]/\mathbb{E}[X]$ should decrease sharply at $\gamma = \gamma_{opt}$, in reality this decrease is much smoother due to the randomness associated to the deferral and backoff counters. Despite this, our configuration is still much closer to γ_{opt} (and hence to the point of optimal performance) than the default configuration.

6. PERFORMANCE EVALUATION

In this section, we present our experimental and simulation settings, and evaluate the performance of our proposed 1901 configurations. Using experimental results from a HomePlug AV and WiFi testbed, we show how throughput is affected by the number of stations in practice for these two MAC layers. We then validate our simulator using testbed experiments. Finally, we evaluate our 1901 configurations proposed in Section 5 and we observe that they perform similar to an optimally-configured 802.11 (which would be aware of the number of stations).

6.1 Experimental Setup and MAC Overhead

We built a test-bed of 8 stations, each comprising a Homeplug AV interface (a miniPCI card with Intellon INT6300 chip), and a 802.11n interface (a miniPCI card Atheros DNMA-92). The stations are Alix boards running the OpenWrt Linux distribution [6]. We conduct our testbed experiments with the default configuration of 1901 (i.e., CA1). Unfortunately, the firmware of all HomePlug devices is encrypted, and parameters such as the CW_i 's and the d_i 's are firmware based and they cannot be modified. So far, vendors have not released any information on how to modify these values⁹.

Both WiFi and 1901 MAC layers introduce some overhead, due to backoff delays, headers, preambles, acknowledgments and inter-frame spaces. In addition, 1901 devices are subject to an additional overhead, due to the presence of management messages (MMEs). These MMEs are required for updating the modulation schemes between any two communicating stations. Both 802.11n and HomePlug AV use OFDM at the PHY layer. The difference between the two is that in 802.11n all carriers use the same modulation scheme, which is defined by the MCS index in the frame header [5], whereas in HomePlug AV each carrier can employ a different scheme, because different frequencies experience variable attenuation. As a result, the HomePlug AV stations have to exchange MMEs that indicate the modulation per carrier (917 carriers in total), and the receiver has to update the transmitter with this information each time the channel conditions (bit error rate) change [7]. The MMEs are sent with priority CA2/CA3 and they also use CSMA/CA and consume backoff time. The exact amount of messages ex-

changed is not specified in the standard and depends on the implementation.

We now conduct some experiments to measure MAC and MMEs overheads when the number of stations N varies. In particular, we consider N stations, with $1 \leq N \leq 7$, that send saturated UDP traffic using `iperf` to the same non-transmitting station. The stations are plugged to the same power-strip in order to avoid external interference and study the PLC MAC in isolation. We use the PLC tool `faifa` [2] to get the relevant statistics from the PLC interfaces. The guidelines to reproduce all the testbed measurements of this work are presented in [22]. Figure 5 shows the throughputs obtained with 802.11 and 1901, and the 1901 MME overhead (measured by the fraction of management frames among all data frames sent). Both MACs follow relatively




Figure 5: Total throughput of N saturated stations transmitting, for $1 \leq N \leq 7$, obtained by experiments with HomePlug AV and 802.11n devices (left), and MMEs overhead of the HomePlug AV tests (right). The measured physical rate for PLC is ~ 150 Mbps and for wireless ~ 300 Mbps. Both protocols include a large MAC overhead (about 40%), with HomePlug AV having the additional overhead of MMEs.

similar trends for throughput degradation when N grows. However, for $N \geq 3$, the MMEs overhead represents about 10% of all frames sent. Further, for $N = 1$, we can compute the MAC layer overhead by computing the ratio of received throughput at the MAC layer and the physical data rates used on the links. We observed similar overheads for WiFi (about 40%) and PLC (about 38%). The fact that the throughput degradation is similar for WiFi and PLC – even though PLC is subject to a significant MMEs overhead – suggests that the 1901 MAC is otherwise more efficient than 802.11. This efficiency is revealed by the introductory example of Figure 3, in which 1901 suffers from fewer collisions than 802.11. Our results also indicate that a significant fraction of the overhead in 1901 is specific to MMEs. As the channel conditions are ideal in our tests, we expect that this overhead is larger for lossy links or links with electrical devices creating impulsive noise.

6.2 Simulator Validation - Model Evaluation

In this subsection, we validate our simulator with testbed experiments, and we evaluate our model using simulations. We use the collision probability and not throughput to verify our simulator, because we ignore overheads such as MMEs, as their arrival rate depends on the channel conditions and on the 1901 implementation. We compare the value of γ measured in the testbed with the value γ computed in simulation, and γ obtained by our model. To measure γ in testbed, we use the Atheros PLC toolkit [1] as explained in [22]. To simulate 1901, we wrote a Matlab simulator,

⁹For these reasons, we also wrote a Matlab simulator for 1901, which will be validated against experiments in the next subsection.

which implements the full CSMA/CA process of 1901 and is provided in [22]. Figure 6 shows the average collision probability obtained from 10 tests run as described in Section 6.1, and 10 simulation runs. We observe an excellent fit between the experiments and simulations. In addition, our model accurately predicts the collision probability.




Figure 6: Collision probability obtained by simulation, our model of Section 4, and experiments with HomePlug AV devices.

We now present the normalized throughput obtained by our model and simulator. We ignore the bit loading of the frames, because in reality, the HomePlug AV physical rate varies in intervals of few *ms* [7]. We do not consider these rates as we propose later a MAC enhancement that can be used for different data-rate HomePlug specifications, since they all employ the same CSMA/CA process. To simulate the 1901 CSMA/CA method, we use the time slot duration and timing parameters as specified in the standard (see Table 2). A PLC frame transmission is preceded by two priority tone slots (*PRS*) and a preamble (*P*), and is followed by a response inter-frame space (*RIFS*), the ACK, and finally, the contention inter-frame space (*CIFS*). A PLC frame has a duration D . Thus, a successful transmission has a duration $T_s = 2PRS + P + D + RIFS + ACK + CIFS$. In case of a collision, the stations differ their transmission for *EIFS* μs , where *EIFS* is the extended inter-frame space used by 1901. Hence, a collision has a duration $T_c = EIFS$.

Parameter	Duration (μs)
Slot σ , Priority slot <i>PRS</i>	35.84
<i>CIFS</i> , <i>RIFS</i>	100.00
Preamble <i>P</i> , <i>ACK</i>	110.48
Frame duration D	2050.00
<i>EIFS</i>	2920.64

Table 2: Simulation parameters.

Figure 7 presents simulations of 1901 with the default parameters for the two priority classes CA1 and CA3 (CA0 and CA2 are equivalent). We show the normalized throughput of the network obtained by (10) and the collision probability predicted by our model.

We observe a good fit between analysis and simulation, especially when N is not too small, or when the CA2/CA3 priorities are used. The fact that the accuracy is somewhat reduced for small N in CA0/CA1 priorities is due to the decoupling assumption, which fails to capture the coupling introduced by the deferral counter. In [19], we study this coupling and we propose a different (and more complex) model that does not invoke the decoupling assumption and has a better accuracy.

6.3 Enhancement of the IEEE 1901 MAC

We now evaluate the enhancements proposed in Section 5. Given the parameters of Table 2 and using (18) and (19),

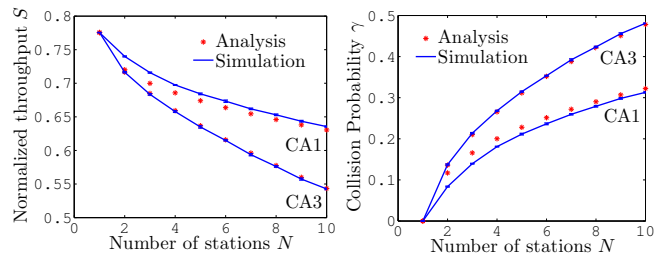


Figure 7: Performance metrics obtained by simulation and with our model, for the default configurations of 1901 (given in Table 1).

the parameters of the enhanced 1901 are $CW_{min} = 12$ and $d_i = \{0, 1, 3, 7, 14, 27\}$ for each backoff stage, with $m = 6$. Remember that our simulator is validated by testbed results and that the enhancements consist simply in modifying the parameters of the protocol.

Figure 8 compares the performance of (i) 1901 using default CA1 configuration, (ii) 1901 using our proposed enhanced configurations mentioned above, and (iii) “optimal” 802.11 for a varying number of stations. For the “optimal” 802.11, CW_{min} is computed from the fixed-point equation for the transmission probability of the 802.11 MAC (see [8, 16]) given τ_{opt} in (11) and γ_{opt} in (12). It appears that 1901 with our proposed configurations performs similarly – or better than – the optimally-configured 802.11, which requires knowing the number of stations N . Furthermore, the proposed configurations drastically boost the efficiency of 1901.

In Figure 9, we show similar results, but using $m = 4$ instead of $m = 6$ for the number of backoff stages. Here too, our proposed configurations yield substantial improvements. These improvements can offer tens of Mbps of throughput gain, given the data rate (up to 1Gbps) of HomePlug AV2.

In order to evaluate the sensitivity to the proposed configuration, we have conducted some experiments in which we perturbed the configuration of the various parameters around the proposed setting. Qualitatively, the results of these experiments showed that: (i) the proposed configuration is quite robust, as small perturbations of the parameters do not yield significant changes in the resulting throughput; (ii) the d_i parameter is particularly robust, as changing its configuration does not result in any noticeable variation in performance; (iii) the CW_i parameter is also quite robust, although in this case it can be noticed that, when increasing this parameter, throughput (slightly) decreases for small N and increases for large N .




Figure 8: Simulations of 1901, enhanced 1901 proposed here, and optimal 802.11 for $m = 6$.




Figure 9: Simulations of 1901, enhanced 1901 proposed here, and optimal 802.11 for $m = 4$.

In the following subsections we discuss the impact of our proposed configuration on multiple contention domains and non-saturated stations scenarios.

6.4 Multiple Contention Domains

The performance of the 1901 MAC under multiple contention domain scenarios is quite different from the results presented in previous subsections. 1901 specifies an optional RTS/CTS mechanism, which is not enabled by default. Thus, without this mechanism, the performance degrades in scenarios such as hidden nodes. Our proposed configuration does not solve such performance issues, but it attenuates their effects as it uses larger a contention window than the default one (see [18] for a similar analysis in an optimal 802.11 approach). Note that the deferral counter is not triggered if two interfering stations are hidden to each other.

6.5 Saturated and Non-Saturated Stations

We now evaluate the performance of our proposed configurations in a mixed scenario of both saturated and non-saturated stations. The non-saturated stations have Poisson traffic with arrival rate λ . Figure 10 presents the normalized throughput for a scenario where half of the stations are saturated and half are not. Also, it depicts the same metric under a scenario where only 3 stations are saturated, whereas the number of non-saturated stations varies. We observe that under these scenarios, our proposed configurations perform better than the optimal 802.11 and the default 1901.

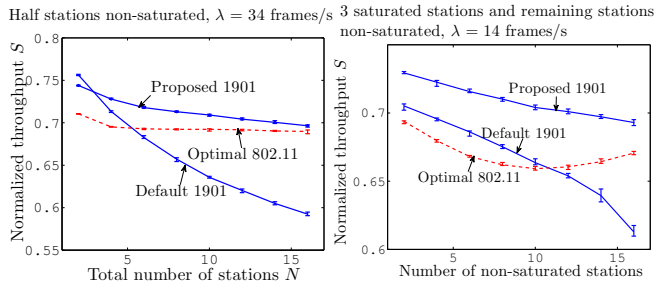


Figure 10: Simulations of 1901, enhanced 1901 proposed here, and optimal 802.11 for $m = 6$, in scenarios that include non-saturated stations.

7. CONCLUSIONS

The MAC layer of 1901 can react to contention with a fine granularity and without involving collisions, which offers the potential of strong gains in terms of throughput as compared to 802.11. Unfortunately, with the default parameter setting of 1901, the protocol operates far from optimality and does

not fully exploit its potential. We believe that the cause of this is the lack of a simple model of 1901 that can be leveraged to find an appropriate setting of its parameters. Indeed, despite the commercial success of IEEE 1901 and its wide adoption in home networks, this protocol has remained largely unexplored by the research community.

One of the main challenges to model 1901 is the complexity of the protocol, which has a very large state-space. To reduce the state-space and come up with a simple model, we make the assumption that the backoff processes of the stations are independent. Building on the resulting model, we derive a procedure to steer the network towards its optimal point of operation. With this, we obtain a protocol that provides performance close to optimal independently of the number of stations, which is a result that has been long pursued by the research community.

Our proposal consists in only modifying some existing parameters, and does not require to change the CSMA/CA 1901 algorithm itself. Therefore, it can be easily incorporated into practical deployments. Our simulations confirm the drastic performance improvements. Moreover, we employed testbed measurements to validate our results and investigate MAC overheads of 1901 implementations.

Appendix

LEMMA 3. *Let B_i be the expected number of backoff slots between two transmissions attempts of a station that always remains at backoff stage i . Then, B_i is given by $bc_i/t_i - 1$, and B_i is an increasing function of γ for any i .*

PROOF. By its definition, B_i is given recursively by

$$B_i = \frac{d_i(d_i + 1)}{2CW_i} + \sum_{j=d_i+1}^{CW_i-1} \frac{j(1 - x_j^i) + \sum_{k=d_i+1}^j (k + B_i)(x_k^i - x_{k-1}^i)}{CW_i} \quad (20)$$

Now observe that solving (20) over B_i , gives $B_i = bc_i/t_i - 1$, with bc_i and t_i given by (2) and (4).

To prove the second part of the lemma, we proceed as follows. (i) First, we compute $dB_i/d\gamma$. (ii) Second, we show that this derivative is positive at $\gamma = 1$. (iii) Third, we show that if the derivative is negative at some $0 < \gamma^* < 1$, it will also be negative at any value $\gamma > \gamma^*$. The proof then follows by contradiction: if the derivative was negative at some γ^* , it would also be negative at $\gamma = 1$, which would contradict our result (ii).

(i) After rearranging terms, (20) can be recast as

$$B_i = \frac{CW_i - 1}{2} + \frac{1}{CW_i} \sum_{j=d_i+1}^{CW_i-1} \left(B_i x_j^i - \sum_{k=d_i+1}^{j-1} x_k^i \right). \quad (21)$$

The derivative of B_i can be computed as

$$\frac{dB_i}{d\gamma} = \sum_{k=d_i+1}^{CW_i-1} \frac{\partial B_i}{\partial x_k^i} \frac{\partial x_k^i}{\partial \gamma}. \quad (22)$$

The partial derivative $\partial B_i/\partial x_k^i$ can be computed from (21) as

$$\frac{\partial B_i}{\partial x_k^i} = \frac{B_i - (CW_i - 1 - k)}{CW_i} + \frac{\partial B_i}{\partial x_k^i} \sum_{j=d_i+1}^{CW_i-1} \frac{x_j^i}{CW_i}, \quad (23)$$

which yields

$$\frac{dB_i}{d\gamma} = \frac{\sum_{k=d_i+1}^{CW_i-1} (B_i - (CW_i - 1 - k)) \frac{\partial x_k^i}{\partial \gamma}}{CW_i - \sum_{j=d_i+1}^{CW_i-1} x_j^i}. \quad (24)$$

To compute $\partial x_k^i / \partial \gamma$, we observe that x_k^i is the complementary cumulative function of a binomial distribution. By taking its partial derivative, we obtain

$$\frac{\partial x_k^i}{\partial \gamma} = \frac{k!}{(k - d_i - 1)! d_i!} \gamma^{d_i} (1 - \gamma)^{k - d_i - 1}. \quad (25)$$

(ii) Next, we show that $dB_i/d\gamma > 0$ at $\gamma = 1$. At $\gamma = 1$, we have $B_i = CW_i - d_i/2 - 1$ from (20). Substituting in (24) yields $dB_i/d\gamma = d_i/2 + 1$, i.e., $dB_i/d\gamma > 0$.

(iii) Next, we show that if $dB_i/d\gamma$ is negative at some value γ^* , then it is also negative for any $\gamma > \gamma^*$. Observe that in (24) some terms are negative for $k < CW_i - 1 - B_i$. Let us assume that the derivative is negative at γ^* . Then, for some $\gamma = \gamma^* + \epsilon$, $\epsilon > 0$ we have $B_i(\gamma) < B_i(\gamma^*)$. This implies that in (24) some of the terms that were positive with $\gamma = \gamma^*$ become negative for $\gamma = \gamma^* + \epsilon$. Furthermore, from (25) we have that

$$\frac{\partial x_k^i / \partial \gamma}{\partial x_{k-1}^i / \partial \gamma} = \frac{k}{k - d_i - 1} (1 - \gamma). \quad (26)$$

Thus, in (24) the relative weight of the negative terms increases and that of the positive terms decreases. From the above, if $dB_i/d\gamma$ was negative at any γ^* , it would also be negative for all $\gamma > \gamma^*$. Since this contradicts result (ii), we conclude that $dB_i/d\gamma \geq 0$ for $\gamma \in [0, 1]$. \square

COROLLARY 1. B_i is a strictly increasing sequence of i if $CW_{i+1} \geq 2CW_i - d_i - 1$, $0 \leq i < m - 1$.

PROOF. By Lemma 3, the minimum value of B_{i+1} is $B_{i+1}^{\min} = (CW_{i+1} - 1)/2$ at $\gamma = 0$, and the maximum value of B_i is $B_i^{\max} = CW_i - d_i/2 - 1$ at $\gamma = 1$. Setting $CW_{i+1} \geq 2CW_i - d_i - 1$, yields $B_{i+1}^{\min} \geq B_i^{\max}$, hence $B_{i+1} > B_i$ for all $\gamma \in [0, 1]$. \square

Acknowledgment

This work is financially supported by a grant of the Smart-World project of the Hasler Foundation, Bern, Switzerland.

8. REFERENCES

- [1] Atheros Open Powerline Toolkit. <https://github.com/qca/open-plc-utils>.
- [2] Faifa. <http://github.com/ffainelli/faifa>.
- [3] HomePlug Alliance (retrieved 9/2014). www.homeplug.org/media/filer_public/e4/cb/e4cbfc97-5e5e-4df1-9625-c226df16c93f/connected_home_summits_2013.pdf.
- [4] HomePlug Alliance (retrieved 9/2014). www.homeplug.org/news-events/member-pr/358/.
- [5] IEEE 802.11n-2009-Amendment 5: Enhancements for Higher Throughput.
- [6] OpenWrt. <https://openwrt.org/>.
- [7] IEEE Standard for Broadband over Power Line Networks: Medium Access Control and Physical Layer Specifications. *IEEE Std 1901-2010*, 2010.
- [8] G. Bianchi. Performance analysis of the IEEE 802.11 distributed coordination function. *Selected Areas in Communications, IEEE Journal on*, 2000.
- [9] G. Bianchi, L. Fratta, and M. Oliveri. Performance evaluation and enhancement of the CSMA/CA MAC protocol for 802.11 wireless LANs. In *Personal, Indoor and Mobile Radio Communications, 1996. PIMRC'96., Seventh IEEE International Symposium on*.
- [10] M. E. M. Campista, L. Costa, and O. Duarte. Improving the multiple access method of CSMA/CA home networks. In *Proc. CCNC*, 2006.
- [11] C. Cano and D. Malone. On Efficiency and Validity of Previous Homeplug MAC Performance Analysis. *arXiv preprint arXiv:1401.6803*, 2014.
- [12] M. Chung, M. Jung, T. Lee, and Y. Lee. Performance analysis of HomePlug 1.0 MAC with CSMA/CA. *Selected Areas in Communications, IEEE Journal on*, 24(7):1411–1420, 2006.
- [13] E. Kriminger and H. Latchman. Markov chain model of homeplug CSMA MAC for determining optimal fixed contention window size. In *Power Line Communications and Its Applications (ISPLC), 2011 IEEE International Symposium on*.
- [14] A. Kumar, E. Altman, D. Miorandi, and M. Goyal. New Insights From a Fixed-Point Analysis of Single Cell IEEE 802.11 WLANs. *Networking, IEEE/ACM Transactions on*, june 2007.
- [15] M. Lee, R. E. Newman, H. A. Latchman, S. Katar, and L. Yonge. HomePlug 1.0 powerline communication LANs—protocol description and performance results. *International Journal of Communication Systems*, 16(5):447–473, 2003.
- [16] P. P. Patras, A. Banchs, P. Serrano, and A. Azcorra. A control-theoretic approach to distributed optimal configuration of 802.11 WLANs. *Mobile Computing, IEEE Transactions on*, 10(6):897–910, 2011.
- [17] A. Schwager. An Overview of the HomePlug AV2 Technology. *Journal of Electrical and Computer Engineering*, 2013.
- [18] P. Serrano, P. Patras, A. Mannocci, V. Mancuso, and A. Banchs. Control theoretic optimization of 802.11 w lans: Implementation and experimental evaluation. *Computer Networks*, 57(1):258–272, 2013.
- [19] C. Vlachou, A. Banchs, J. Herzen, and P. Thiran. On the MAC for Power-Line Communications: Modeling Assumptions and Performance Tradeoffs. In *IEEE International Conference on Network Protocols (ICNP)*, 2014.
- [20] C. Vlachou, A. Banchs, J. Herzen, and P. Thiran. Performance Analysis of MAC for Power-line Communications. In *ACM Sigmetrics*, 2014.
- [21] C. Vlachou, J. Herzen, and P. Thiran. Fairness of MAC protocols: IEEE 1901 vs. 802.11. In *2013 17th IEEE International Symposium on Power Line Communications and Its Applications*.
- [22] C. Vlachou, J. Herzen, and P. Thiran. Simulator and Experimental Framework for the MAC of Power-Line Communications. Infoscience EPFL, Technical Report, 2014.