

Dynamics of Contention Free Period Reservation in IEEE 1901 Networks

Brad Zarikoff* and David Malone†

Hamilton Institute, National University of Ireland Maynooth, Ireland

{Brad.Zarikoff*, David.Malone†}@nuim.ie

Abstract—Power line communications (PLC) technology has become a competitor in the home networking arena. In an inter-networked home, PLC promises to provide an inexpensive, high throughput, and easy-to-install means of extending connectivity to areas of the home or small office that have poor wireless coverage. A key enabling technology within the IEEE 1901 PLC standard is that of the contention free period (CFP). Following a successful reservation, a station allocated to the CFP will not suffer any form of contention-based packet loss. The IEEE 1901 standard presents this as a way of accommodating flows with well defined delay, jitter and bandwidth requirements. However, to date there has been little research done on the dynamics of the CFP reservation procedure and its scalability. Since the procedure inherently relies on the collision-prone contention access period for reservations, a successful reservation bid is not guaranteed on the first attempt. Our work looks at the resulting delay characteristics of the IEEE 1901 reservation procedure. We present details for 1-persistent and VoIP traffic.

I. INTRODUCTION

With the recent ratification of the IEEE 1901 power line communications (PLC) standard [1], the networking community can expect to see an increase in the availability of commodity PLC hardware, and thus PLC-based home and small office networks. While there has been a huge push for research and analysis into PLC networks, the bulk of the focus has been on the physical layer. How home PLC networks will perform in scenarios that are more complicated than a single point-to-point link is still an open question.

The complications of PLC networks arise due to the challenging physical layer environment, and propagate into the media access layer (MAC) as a result. Phenomena such as impulsive noise and hidden nodes make the PLC channel look more similar to a wireless environment than a wired environment. As a result, the standardization groups have adopted a variant of the tried-and-tested IEEE 802.11 MAC CSMA/CA protocol [2]. The variations take the form of an additional deferral counter, which provides a rough measure of how many stations are sharing the channel, and of strict prioritization by way of priority resolution symbols (PRSS).

Another departure from 802.11 is that of the *scheduled* contention free period (CFP). The IEEE 1901 standard specifies contention free channel access via an explicit and advertised scheduling process, which is different to that of the IEEE




Fig. 1. PLC beacon structure. B represents beacon symbols. A beacon interval is fixed at 40 ms (50 Hz networks) and 33 ms (60 Hz networks). The contention access period (CAP), the multi-access channel where reservations are requests, is generally kept above a minimum of 4 ms.

802.11 Point Coordinated Function (PCF) [3]. In short, each 33.33/40 ms PLC frame is split into a contention access period (CAP) and a CFP. During the CAP, packets vying for the transmission medium must use the CSMA/CA protocol, but during the CFP, a packet flow can transmit without deferral or collision during its *scheduled* period. In order to gain access to the CFP, an IEEE 1901 station must make a reservation with the basic service set (BSS) manager (BM). This is different than the continuous polling done by the 802.11 PCF, which offers stations a contention free transmission opportunity in a round-robin fashion. While both techniques are rare in hardware implementations, we suspect that the IEEE 1901 CFP may prove to be an enabling technology for home high definition audio/video systems.

Our concern with the CFP mechanism is for a lack of analysis and investigation of the dynamics of the systems. To our knowledge, existing PLC modems restrict channel usage solely to the CAP. However, if significant gains can be realised by implementing the CFP, it may be more attractive for developers to spend time on it. This work here is a first step in an analysis of the CFP, and how the parameters involved control the overall throughput and delay characteristics of a PLC system.

II. BACKGROUND

In the IEEE 1901 standard [1], each segment of the CFP is allocated to an individual data flow, not to a station. As a result, a station has to contend for a new CFP slot by sending a request in the CAP. Specifically, a station has to send a *CC_LINK_NEW.request* message to the BM to request CFP setup. The CAP and CFP mechanisms involved in this process are described in the following.

A. Contention Access Period

Aspects of the HomePlug and IEEE 1901 PLC standards are inspired by the 802.11 CSMA/CA (carrier sense multiple access/collision avoidance) MAC [2]. Upon receiving a packet,

The authors wish to acknowledge funding support from SFI grant 08/SRC/I1403 and 07/SK/I1216a.

an 802.11 MAC will randomly pick a back-off value from between $[0, CW(i)]$, where i is the current back-off stage and CW is the contention window value. The back-off value will be decremented every slot interval (all radios in the system are synchronised based on previous channel usage). Once the back-off counter reaches zero, the radio will transmit. The carrier sense is invoked when a radio detects a transmission on the medium during back-off. If this happens, the radio will halt countdown until the transmission is complete. If a collision does occur after the back-off has reached zero and transmits, then the back-off stage is incremented (note that $CW(i)$ is a monotonically increasing function). This entire process is contingent on the fact that a slot time (the interval between back-off counter decrements) is much smaller than a packet transmission interval, and thus it is more efficient to spend time in back-off stages than in a collision state.

For the CAP of each frame in the IEEE 1901 [1] and HomePlug AV [4] standards, and for the entire operation of TIA 1113 [5], a slightly modified version of the 802.11 MAC is used. The modifications include:

- the deferral counter (DC), which is set at the beginning of a back-off period, and is decremented whenever carrier sense becomes active (this simulates collisions and is a way of quickly adapting the contention window without having to suffer multiple collisions on the medium),
- and priority resolution (PR), which defines hard QoS states such that only devices with the highest active PR level can transmit in the following contention slot.

The original distribution coordination function (DCF), which is the core technique that implements the CSMA/CA MAC in 802.11, has seen a massive amount of research since an analysable and quite accurate model was proposed by Bianchi [6]. Since the power line MAC has the above extensions, and the analysis in [6] does not hold for them, a modified Markov chain analysis was proposed in [7], which accounts for the DC, the PR period, and the new QoS states enforced by PR (although not in a heterogenous case). Similarly, our work does not consider different PR levels.

B. Contention Free Period

To enable the use of low latency and high throughput applications on PLC networks, time division multiplexing has been introduced into the standards [1], [4]. The inclusion of a CFP in PLC was analysed in [8], [9], with both papers championing a CFP. In [9], the idea of a periodic contention free multiple access (PCF/MA) technique is suggested, which is very close to the existing CFP in IEEE 1901.

Access to the CFP is controlled via the BM, which handles all scheduling tasks. The BM calculates and advertises the CFP schedule in the beacon frames that are transmitted every 33/40 ms. The CFP mechanism includes a level of persistence, which is meant to accommodate for periodic beacon losses at the stations [10]. The level of persistence is controlled by the current schedule countdown (CSCD) timer and the preview schedule countdown (PSCD) timer. Before a new schedule is introduced, the CSCD must countdown from its




Fig. 2. Reservation model.

current value to 0. This value is constantly transmitted in the beacon, and so if a station misses a beacon, it is aware of the minimum number of frames the current schedule will be valid for, e.g. if the CSCD is currently $M = 3$, and a new schedule has been chosen, the current schedule cannot be changed for 3 beacon frames. During this countdown period, a preview schedule is also transmitted. While this mechanism successfully protects stations against periodic beacon loss, it also introduces an additional delay burden on the reservation process. This preview schedule cannot be changed until it becomes the current schedule, and so any requests after the CSCD has begun countdown are put on hold. The worst case scenario for a successful reservation request in light of this persistence mechanism is if the request is made directly after the CSCD has begun countdown, and corresponds to a wait of $2M - 1$ frames.

Once a flow has successfully made a reservation and is in the current schedule, the reservation can be canceled by a) the scheduler, b) a station request, or c) by the station exceeded the inactivity time limit. The inactivity time limit (T_{il}) is defined as the maximum number of frames that a CFP reservation will be held in the absence of packet transmit attempts.

III. MODEL

A. Assumptions

Our goal in this work is to investigate the dynamics of the reservation process used in the IEEE 1901 standard. We restrict our investigation to delay dynamics, as they are of the most interest in a heterogeneous network of stations containing at least one reserved flow. Also, we only consider the contribution of the IEEE 1901 reservation MAC rather than network and transport layer contributions (which may be non-trivial). At the moment, our results serve to illustrate a lower bound on the delay.

We also do not consider any saturation conditions of the CFP. Note however that one OFDM symbol requires $T_{1,OFDM} = 25.604 \mu s$ of airtime, and so for a maximum CFP of 36 ms, and $RIFS$, T_{ack} and $CIFS$ from Table I, we can expect a maximum of 54 single slot reservations.

B. IEEE 1901

Our model consists of N_{sta} stations, with the i^{th} station requesting CFP service for n_i flows, for a total of $N = \sum_i n_i$ flow requests. Each request is transmitted on the same PR

CAP	Backoff Stage	CW	DC	
CA0 & CA1:	BPC = 0	7	0	
	BPC = 1	15	1	
	BPC = 2	15	3	
	BPC > 2	31	15	
CA2 & CA3:	BPC = 0	7	0	
	BPC = 1	15	1	
	BPC = 2	31	3	
	BPC > 2	63	15	
Timing	Parameter	Value	Parameter	Value
	T	40 ms	$RIFS$	140 μ s
	$T_{1,OFDM}$	25.604 μ s	$CIFS$	100 μ s
	T_{ack}	11.048 μ s	PRS	35.84 μ s
	T_{il}	(1 ms, 60 s)	σ_0	35.84 μ s

TABLE I

IEEE 1901 MAC PARAMETERS. FOR THE CAP, THE CONTENTION WINDOW (CW) AND DEFERRAL COUNTER (DC) PARAMETERS ARE LISTED PER EACH BACKOFF PROCEDURE EVENT COUNTER (BPC) VALUE AND CONTENTION ACCESS (CA) QoS STATE. FOR PACKET AIRTIMES THAT CORRESPOND TO BOTH THE CAP AND THE CFP, THE FRAME (T), OFDM SYMBOL ($T_{1,OFDM}$), RESPONSE INTERFRAME SPACE (RIFS), CONTENTION INTERFRAME SPACE (CIFS), PRIORITY RESOLUTION SYMBOL (PRS), BACKOFF SLOT (σ_0), AND INACTIVITY LIMIT (T_{il}) DURATIONS ARE PROVIDED.

level (highest priority, CA3), with a request taking the form of the $CC_LINK_NEW.request$ management message. We assume that each CFP service request requires one OFDM symbol ($T_{1,OFDM}$), which requires 25.604 μ s for transmission. The CSMA/CA parameters used in IEEE 1901 are shown in Table I.

C. VoIP Traffic

A primary example of the usage of the CFP in PLC networks is that of VoIP traffic. We model the traffic of a VoIP connection in the talk burst fashion from [11]. Each VoIP connection is modelled as a 64 kb/s on-off packet stream where the on and off periods are exponentially distributed with a mean of 1.5 seconds. The minimum talk-spurt is restricted to 240 ms. Given the data rate of each stream, during a 40 ms frame, each connection must transmit at least

$$64 \text{ kb/s} \times 0.04 \text{ s} = 2560 \text{ bits.}$$

The maximum payload of a single PLC OFDM symbol is 11004 bits, which would easily accommodate the VoIP traffic requirement. Thus, we assume that a VoIP CFP reservation would require airtime of

$$T_{1,OFDM} + RIFS + T_{ack} + CIFS \approx 508 \mu\text{s.}$$

This matches well with the reservation examples in Annex C of [1].

In [11], it is suggested that ‘‘mouth to ear’’ delay should be restricted to less than 100 ms, and so delivering a VoIP packet burst at every PLC frame interval will accommodate this (e.g. one packet every 40 ms). This interframe delay can be reduced with additional reservations during the same frame, but is done at the expense of additional CFP flows.

MOS (lower limit)	User satisfaction
4.34	Very satisfied
4.03	Satisfied
3.60	Some users dissatisfied
3.10	Many users dissatisfied
2.58	Nearly all users dissatisfied

TABLE II

THE RELATION BETWEEN THE E-MODEL MEAN OPINION SCORE (MOS) AND USER SATISFACTION (G.107/ANNEX B [12]).

D. Delay

Our main interest in this model is to measure the frame delay incurred by setting up and maintaining a CFP connection. We define a delay metric d as the number of CFP frames between the initial $CC_LINK_NEW.request$ and when the first packet can be transmitted at the reserved time.

In scenarios that consider bursty traffic, an average delay across all reservations will not properly capture the system dynamics on account of inactivity timeouts. This is because a lost reservation requires the flow to undergo the reservation process again, incurring additional delays. In order to capture this lost reservation dynamic, we opt for including a measure for the total number of inactivity timeouts, E_i , for each flow. Using $d_{e,i,k}$ to represent the delay of the e^{th} reservation request of the i^{th} flow during the k^{th} consecutive frame, we can then define the mean delay as

$$\mathbb{E}[d] = \frac{1}{KN} \sum_{i=1}^N \frac{1}{E_i} \sum_{e=1}^{E_i} \sum_{k=1}^K d_{e,i,k} \quad (1)$$

and the mean number of inactivity timeouts

$$\mathbb{E}[E] = \frac{1}{N} \sum_{i=1}^N E_i, \quad (2)$$

where $E_i = \sum_{k=1}^K E_{i,k}$ is the total number of inactivity timeouts for the i^{th} flow, $E_{i,k}$ is an inactivity timeout indicator function, and K is the total number of simulated frames. Note that $\mathbb{E}[E]$ grows unbounded with K . However, for our purposes, we are interested in the dependence of $\mathbb{E}[E]$ when parameters are changed, and so we keep $K = 2000$ fixed for all simulations unless otherwise noted.

Finally, in order to better capture the satisfaction that a user may experience using VoIP over PLC, we use the E-model mean opinion score (MOS) [12]. The E-model was developed by the International Telecommunications Union (ITU) in 2000 to help model the user experience of handset telephony. The relationship between the E-model generated MOS and subjective user experience is shown in Table II. In our usage of the E-model, we only focus on changing the one-way delay parameter Ta , leaving all other parameters to the default values. We use the mean delay $\mathbb{E}[d]$ for Ta in the E-model MOS algorithm during simulations.

IV. SIMULATION

The simulations of the IEEE 1901 reservation mechanism were done using a slotted Monte Carlo Matlab script. Fig.

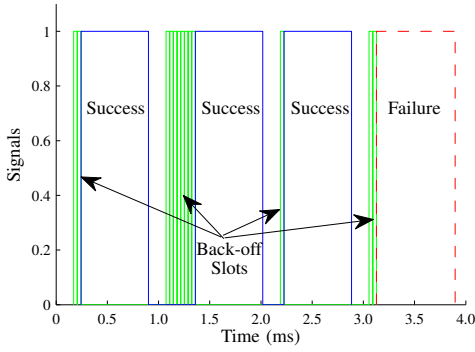


Fig. 3. A snapshot of the slotted simulator, with back-off slots, collisions and successes marked.

3 shows an example snapshot of the simulator, where back-off slots, successful reservation attempts, and collisions are denoted. The parameters used in the simulator are taken from Table I. Each run of the simulator consists of 2000 consecutive PLC frames (2000×0.04 seconds/frame = 80 seconds). Each curve in a figure is the result of averaging over 200 simulations. For the VoIP traffic simulations, each flow begins without a reservation and in a silence burst: with a silence burst having a mean of 1.5 s, and the simulation running 80 s, the results are well into steady state by the end of the simulation. The slotted time counter in each CAP starts with an immediate addition of $2PRS + CIFS$, followed by the back-off period. Each success is of size $T_{1,OFDM} + RIFS + T_{ack} + CIFS$, with an additional $2PRS + CIFS$ added afterwards to setup the next back-off period. Because of the NAV protection period, each collision is of size $T_{1,OFDM} + RIFS + 2T_{ack} + CIFS$, with an additional $2PRS + CIFS$ added afterwards to setup the next back-off period. Unless otherwise stated, we use the IEEE 1901 default of $T_{CAP} = 4$ ms and standard example value for CSCD of $M = 3$. The delay d is calculated as the number of frames, and then converted to seconds by $d \times T$, where $T = 40$ ms.

A. Saturated Traffic

To provide a baseline for our simulations, we first provide results for mean delay $\mathbb{E}[d]$ for saturated reservation traffic. This provides a snap-shot of the reservation setup time for N flows. Since the N flows are saturated, their reservations will never time-out. The results are shown for different values of CSCD M in Fig. 4, and demonstrate the effect of the CSCD on the minimum mean delay. Each curve starts with $N = 1$ at the point $(M + 1) \times T$. The effect of adding additional flows is clear, with an additional penalty incurred with an increase to the countdown schedule.

The recommended minimum value for T_{CAP} in the IEEE 1901 standard is set to 10% of the full beacon interval, so for a 50 Hz power signal, $T_{CAP} = 4$ ms. The effect of decreasing or increasing the CAP interval is shown in Fig. 5. The variation in $\mathbb{E}[d]$ is a feature of the relative size difference of T_{CAP} and a reservation packet. It is clear that the IEEE 1901 standard value of T_{CAP} provides a good trade-off between reservation




Fig. 4. Number of stations N versus the mean delay $\mathbb{E}[d]$ for saturated traffic. Variable countdown schedule length M . The error bars correspond to one standard deviation.




Fig. 5. Size of the contention access period T_{CAP} versus the mean delay $\mathbb{E}[d]$ for saturated traffic, for various N . The error bars correspond to one standard deviation.

delay and reservation period $T_{CFP} = T - T_{CAP}$. However, for small $N \leq 10$, reducing T_{CAP} to 3 or even 2 ms will not have a dramatic impact on $\mathbb{E}[d]$, e.g. 5 to 20% penalty on average for $N = 10$. Increasing to 5 or 6 ms will drop the mean delay by 5 to 10%.

B. VoIP Traffic

By modeling VoIP traffic, we get an idea of trade-off between delay d , the persistence level M , and the inactivity limit T_{il} . Fig. 6 shows the effect of changing M with a very large inactivity limit, e.g. the average number of timeouts $\mathbb{E}[E] = 0$. A trend similar to that in Fig. 4 is visible when varying M and N . Because the inactivity limit is set very high, the trend is similar to the system with saturated reservation traffic, but with a lesser slope. The reduced delay compared to Fig. 4 is attributed to the system having, on average, fewer reservation contentions (and therefore collisions).

Fig. 7 shows the location of an inactivity limit threshold, with Fig. 7(b) revealing that the $\mathbb{E}[E]$ threshold is insensitive to the number of reservation flows N . Our results only extend to $T_{il} = 400$ ms since further points along the $\mathbb{E}[d]$ curve are difficult to calculate due to an initial setup delay bias: with large T_{il} , there are fewer timeouts, and so the average delay becomes inflated with any collisions during the initial reservation process. To compensate, we have extended the




Fig. 6. Number of stations N versus the mean delay $\mathbb{E}[d]$ for VoIP traffic. Variable countdown schedule length $M = 1, 3$ and 7 ; inactivity limit T_{il} very large (1000 frames = 40 seconds). The error bars correspond to one standard deviation.

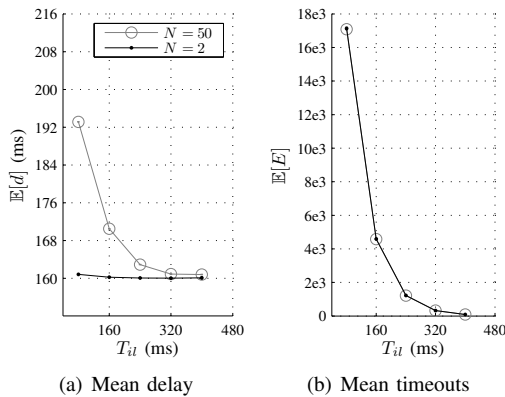


Fig. 7. Mean delay and timeouts versus T_{il} with variable N ; $M = 3$, $T_{CAP} = 40$ ms. To remove the initial settling bias, the simulator length was extended to 5×10^6 consecutive frames. The error bars correspond to one standard deviation.

simulation window to $N_{sim} = 5 \times 10^6$ frames (or 55 hours) to ensure that this bias was not present in Fig. 6. Both the delay and the number of timeouts have an inverse relationship with T_{il} . While the mean delay converges to the minimum delay of $T \times (M + 1)$, the mean number of timeouts drops to near zero above the threshold. For the VoIP traffic model considered, any $T_{il} > 400$ ms has little effect on the mean number of timeouts. In more concrete terms, a mean delay product $\mathbb{E}[d](\mathbb{E}[E] + 1)/(N_{sim})$ could be used as a simple metric to determine appropriate parameter values for QoS, where the additional 1 is added to $\mathbb{E}[E]$ to account for the initial reservation process. This product is an average delay per frame, over the entire simulation. For example, with $T_{il} = 160$ ms, the mean delay product from Fig. 7 is 0.1471 and 0.1563 ms for $N = 2$ and 50, respectively. By increasing T_{il} to 400 ms, the mean delay product drops to 0.0027 and 0.0027 ms, respectively. This threshold is a factor of the VoIP traffic on-off statistics, as further experiments have verified. As a result, we can conclude that it is difficult to tune T_{il} for general traffic, as it appears to be a feature of higher layer traffic statistics.




Fig. 8. Delay for VoIP flow reservations with saturated background traffic, $T_{il} = 10$, and $M = 3$. The mean delay $\mathbb{E}[d]$ increases with the number of VoIP calls.

C. VoIP Traffic Plus Background Saturated Traffic

A more realistic scenario for CFP reservations would include background CAP traffic. This traffic could be the result of, for example, web page loads, file downloads, and peer-to-peer file sharing. The result for the reservation process is increased congestion for the reservation requests. We look at the mean delay in the presence of background saturated traffic. Fig. 8 shows $\mathbb{E}[d]$ as the number of background flows increase. The effect on the VoIP reservation process is intuitive: as additional reservation flows and background flows are added, the mean delay increases. However, an unintuitive result is that the delay jitter is inversely impacted by the number of stations (not shown). With no background traffic, the delay jitter increases with additional reservation flows, from a variance of 3 ms with $N = 2$ to 28.7 ms with $N = 50$. However, this trend is reversed as background traffic is introduced: for $N_{bg} = 10$, the jitter is improved with additional reservation flows, from a variance of 470 ms with $N = 2$ down to 110 ms with $N = 50$. This is done at the expense of the mean delay. The mechanism behind this is thought to be the result of many reservations being queued during the current schedule countdown process, thus preventing any large delays due to CSCD countdowns.

D. Mean Opinion Score

To provide a look of how the mean delay in the PLC MAC will effect the VoIP call experience, we include some simulations of the MOS versus the number of active VoIP calls. These results consider otherwise optimum operating conditions, and represent the relative decay in the call experience as additional VoIP calls and background flows are added to the system. We consider the one-way trip delay T_a in [12] to be equivalent to $\mathbb{E}[d]$.

Fig. 9 shows that additional VoIP flows have significant effect on the MOS, particularly in the presence of saturated background flows. Indeed, this can be surmised from Fig. 4, where $\mathbb{E}[d]$ rises by nearly double as N nears 20 flows. More serious degradation is obvious as the number of saturated background CAP flows are increased. Referring to Table II and assuming $N_{bg} = 8$, the call quality falls in the ‘satisfied’




Fig. 9. Mean opinion score (MOS) versus the number of VoIP flows, with a variable number of saturated background CAP flows; $M = 3$, $T_{il} = 8$, $T_{CAP} = 4$ ms.

category with $N \leq 25$, while with $25 \leq N \leq 50$, the call quality falls in the ‘some users dissatisfied’ category. Additional network latency will only further degrade call quality. This shows the necessity to ensure that reservation flows must be prioritised over background flows. For $N_{bg} = 0$, the call quality falls in the ‘very satisfied’ category for $N \leq 25$. It is worth noting that this prioritisation is not a requirement for PLC systems, and will in fact be hard to ensure since it is trivial for the user to change the traffic priority of the flows [13].

V. DISCUSSION

As a result of the simulations in Section IV, we can make some suggestions for tuning IEEE 1901 system parameters for VoIP transmission.

1) *Schedule Persistence*: It is clear that M has a high impact on the mean delay, and so minimizing this variable is of some importance. Since schedule persistence is more important in channels where beacons are liable to be lost, a good starting point would be $M = 1$, only increasing M if the overall system performance is marginal.

2) *Inactivity Level*: A large value for inactivity level T_{il} will reduce the requirement of VoIP flows to re-reserve a CFP slot, thus reducing the mean delay $\mathbb{E}[d]$ and mean delay product $\mathbb{E}[d](\mathbb{E}[E] + 1)$. From our simulations, it appears that restricting $T_{il} > 400$ ms would be prudent, and that any further increase would not benefit VoIP performance. However, T_{il} also impacts the system throughput, as a higher setting for T_{il} will result in lost transmission opportunities for other stations. An optimisation between these two parameters is an interesting question for future work, keeping in mind that the optimal T_{il} is traffic dependent.

3) *Other Factors*: There are other factors as well that will impact the reservation delay characteristics. For example, prioritization of the reservation requests will permit the requests to take precedence when lower priority background traffic is sharing the channel. However, this should not be counted on, as there is currently no reason that background traffic can’t be prioritised as well.

VI. CONCLUSION

New standards in power line communication technology have included a powerful mechanism that ensures quality of service (QoS) for delay sensitive applications, such as voice and video. This work presents a model for the contention free period (CFP) reservation mechanism that serves as the backbone of this QoS mechanism. We present a simplified simulation model that enables us to investigate the reservation dynamics. We have included both realistic VoIP traffic dynamics and have shown the effect of including non-reservation based background traffic. As part of this work, we have presented suggestions as to how best to operate the CFP mechanism. For example, for standard VoIP traffic models, the inactivity level should be set to at least a value of 10 frames (400 ms) to minimize the mean delay, and that increasing the persistence level above 3 frames (120 ms) will have serious consequences for the delay. We also translate some of our results to a subjective mean opinion score, and demonstrate that in the absence of other network impairments, tuning of the CFP variables are essential to providing a high QoS.

REFERENCES

- [1] Broadband Over Power Lines PHY/MAC Working Group, “IEEE P1901/D4.01 Draft Standard for Broadband over Power Line Networks,” IEEE, Tech. Rep., 2010.
- [2] IEEE 802.11 Working Group, “IEEE Std 802.11TM-2007: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications,” IEEE, Tech. Rep., 2007.
- [3] B. Sikdar, “An Analytic Model for the Delay in IEEE 802.11 PCF MAC-Based Wireless Networks,” *IEEE Trans. on Wireless Commun.*, vol. 6, no. 4, pp. 1542–1550, April 2007.
- [4] HomePlug Powerline Alliance, “HomePlug AV White Paper,” 2005, <http://bit.ly/ptpeN1>.
- [5] Standards and Technology Department, “Medium-Speed (up to 14 Mbps) Power Line Communications (PLC) Modems using Windowed OFDM,” Telecommunications Industry Association, Tech. Rep., 2008.
- [6] G. Bianchi, “Performance Analysis of the IEEE 802.11 Distributed Coordination Function,” *IEEE Journal on selected areas in communications*, vol. 18, no. 3, pp. 535–547, 2000.
- [7] M. Y. Chung, M.-H. Jung, T. Lee, and Y. Lee, “Performance Analysis of HomePlug 1.0 MAC With CSMA/CA,” *IEEE Journal on Selected Areas in Commun.*, vol. 24, no. 7, pp. 1411–1420, 2006.
- [8] H. Hrasnica and R. Lehnert, “Reservation Domains in MAC Protocols for Broadband PLC Networks,” in *IEEE Int. Symp. on Power Line Commun. and Its Applications*, 2005, pp. 336–340.
- [9] Y.-J. Lin, H. A. Latchman, J. C. L. Liu, and R. Newman, “Periodic Contention-Free Multiple Access for Broadband Multimedia Powerline Communication Networks,” in *IEEE Int. Symp. on Power Line Commun. and Its Applications*, 2005, pp. 121–125.
- [10] S. Katar, M. Krishnan, B. Mashburn, K. Afkhamie, R. Newman, and H. Latchman, “Beacon Schedule Persistence to Mitigate Beacon Loss in HomePlug AV Networks,” in *IEEE Int. Symp. on Power Line Commun. and Its Applications*, 2006, pp. 184–188.
- [11] A. P. Markopoulou, F. A. Tobagi, and M. J. Karam, “Assessing the quality of voice communications over internet backbones,” *IEEE/ACM Trans. on Networking*, vol. 11, no. 5, pp. 747–760, Oct. 2003.
- [12] Telecommunication Standardization Sector of ITU, “The E-model, a computational model for use in transmission planning,” ITU, Tech. Rep., 2000.
- [13] B. Zarikoff and D. Malone, “Construction of a PLC test bed for network and transport layer experiments,” in *IEEE Int. Symp. on Power Line Commun. and Its Applications (ISPLC)*, Apr. 2011, pp. 135–140.