

Construction of a PLC Test Bed for Network and Transport Layer Experiments

Brad Zarikoff* and David Malone†

Hamilton Institute, National University of Ireland Maynooth, Ireland

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

Abstract—With the recent ratification of the IEEE 1901 power line communications (PLC) standard, PLC is at a tipping point and may become ubiquitous as a home networking technology. 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. It is imperative that the network dynamics of PLC systems be fully understood in advance of any large-scale roll-out. We present this work as a first step in our development of a comprehensive PLC network test bed. It provides the details of our test bed together with some initial experiments and observations. In particular, we are able to confirm the basic operation of our test bed and highlight issues, such as the interaction of TCP and the PLC MAC, which require further investigation.

I. INTRODUCTION

With the recent acceptance of the IEEE P1901 Draft 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. The power line medium has drawn the research community to the problem of dealing with the hostile environment that PLC signals must contend with (impulsive noise, sudden network topology changes, hidden nodes, high signal loss, etc.). For this reason, higher layers have not received the same attention to this point, with some exceptions.

There has been substantial test bed development for the investigation of the physical layer [2], [3]. However, networking issues have largely been addressed solely in simulations. Recently, the NetLab group at the Hamilton Institute has begun to gather HomePlug AV compliant devices with the goal of creating a network-centric test bed. Currently, the test bed consists of a modest 5 devices, but the hope is to expand to 30 devices in the next few months. The immediate goal of creating such a test bed is two-fold: to verify existing analytical results for PLC networks and to search for previously unsuspected network issues in PLC networks. A side benefit of this work is to provide a helpful reference for other labs wishing to setup their own network test beds, and to add to the existing network test bed literature in [4]–[6].

The authors wish to acknowledge funding support from the Irish HEA PRTL1 Cycle 4 FutureComm (<http://futurecomm.tssg.org>) programme and SFI grant 07/IN.1/1901.

II. BACKGROUND

Our interest in power line networks is in network dynamics and in modeling phenomenon likely to be observed by real-world network users. Characterisation of our PLC test bed is essential to knowing if the phenomena we witness are due to experimental or implementation issues or due to MAC/PHY interactions that make up a PLC network. To this end, we will begin with a brief review of network research into PLC.

Aspects of the HomePlug and IEEE 1901 PLC standards are inspired by the 802.11 CSMA/CA (carrier sense multiple access/collision avoidance) MAC [7]. The 802.11 standard transmits on a shared and volatile medium (wireless at 2.4 GHz and 5 GHz), which is similar to the power line medium in that it suffers from multipath channel distortion, high signal attenuation, and communication link asymmetry (hidden nodes). Since the half-duplex nature of wireless channel access prohibits a radio from sensing a collision, it is in the networks best interest to eliminate collisions (the radio is half-duplex because of the relative signal levels at each receiver). To do this, collision avoidance in the form of a back-off mechanism is used. Upon receiving a packet, 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 synchronized 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 contention period of each beacon frame in the PLC standards IEEE 1901 [1] and HomePlug AV [8], and for the entire operation of TIA 1113 [9], 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 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 [10]. Since the power line MAC has the above extensions, and the analysis in [10] does not hold for them, a modified Markov chain analysis was proposed in [11], which accounts for the DC, the PR period, and the new QoS states enforced by PR (although not in a heterogenous case).

To enable the use of low latency and high throughput traffic in PLC networks, time division multiplexing and frame concatenation have been introduced [1], [8]. The contention-free period (CFP) in PLC has been analysed [12], [13], as has the MAC framing [14], [15]. Framing has been looked at as a means to preserve throughput as the payload symbols grow to 4096 QAM. To this end, the new standard uses a two-level MAC framing scheme, where MAC service data units (MSDUs) are encapsulated into MAC frames and segmented into fixed sized physical blocks (PBs) [15].

For power line network test beds, the majority of existing work has looked at the effect of the PHY layer. Early work in [2] dealt with physical phenomenon and recreating test bed channels that might mimic ones seen in an in-home scenario, while recent work in [3] uses recorded data from over 40,000 different users to evaluate adaptive transmit power techniques. Work in higher layer (TCP/IP/MAC) PLC network interactions has been limited to this point. Benchmark tests have been done on HomePlug 1.0 and AV networks in [4]. Some experimentation has been done with the TCP/IP layer in [5], which considers large-scale traffic limits of a very specific topology, while a recent paper on an in-home mesh network [6] deals with path selection and congestion issues.

III. TEST BED SETUP

Our power line test bed contains 5 nodes. Each node is built using a Soekris Net5501 single board computer and a Devolo dLAN 200 AVminiPCI card. The Devolo dLAN 200 is based on an Intellon INT6300 PLC chip on a miniPCI card. The INT6300 is connected to the PCI bus via a Realtek 8139C Ethernet chip, and PLC cards appear as a standard Ethernet card so no special drivers are required. The INT6300 provides an implementation of the HomePlug AV MAC and PHY.

The Soekris Net5501 is based on an AMD Geode processor and has enough RAM to allow standard installs of FreeBSD or Linux to be used. An install of FreeBSD on a compact flash card was used for the operating system, and standard tools such as iperf, tcpdump and mgen for network tests. The Soekris Net5501 also provides 4 ethernet ports, which were used for management of the test bed and baseline tests.




Fig. 1. Test bed logical schematic

Figure 1 shows a logical schematic of our test bed. It shows the 5 Soekris boxes with associated PLC cards, a connector block that serves as the communication medium, and a spectrum analyser for monitoring traffic on the channel. Since the PLC card is powered from the miniPCI bus, we have flexibility on what is used as the PLC communication medium. For example, we could use our connector block (simple channel, time non-selective), a hardware channel emulator (simple or complex channel, time selective or non-selective), or the mains (complex channel, time selective). We have used standard Devolo coupling devices to connect each of the miniPCI cards to a common wiring block. Our choice gives us control over the medium, and also allows us to monitor the physical channel directly.

For measuring the physical traffic on the test bed, a Rohde & Schwarz FSL6 spectrum analyser was connected via a wire to the wiring block. Zero-span mode was used to capture traffic traces. Our measurement settings were as follows: 20 MHz for the resolution bandwidth, 10 MHz video bandwidth, 5 MHz start frequency and 25 MHz stop frequency. Our tests indicate that the spectral mask of the Devolo miniPCI cards was set to restrict sub-carriers above 30 MHz, and since our spectrum analyser is limited to a 20 MHz resolution BW, we centered our measurements mid-band.

A. Test Bed Considerations

Upon setting up the AVminiPCI cards on each Soekris box, it was necessary to enable 100baseTX mode rather than autonegotiation mode on the Realtek Ethernet driver. This was required before any communications via the power line was possible. The use of commodity Soekris boxes has necessitated some modifications to our system settings, which are noted here to aid colleagues in future test bed construction. Ensuring that an embedded device that uses PLC has sufficient CPU resources is important, as some potential applications of PLC includes devices such as TVs or PVRs, which may have similar constraints to our Soekris devices. In comparison to the average home computer or dedicated network router, such devices operate with low-performance processors. When receiving multiple flows, this did not appear to be an issue for our Soekris boxes; however, for tests where a station generates many packets, such as when we generate multiple flows, a significant amount of interrupts are generated. Our hardware idle time was below 20% with only two downlink flows.

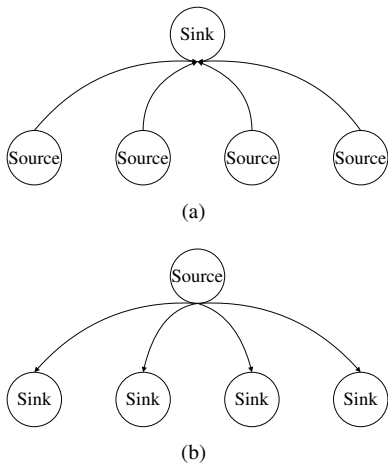


Fig. 2. (a) shows the uplink case, (b) shows the downlink case.

To remedy this problem, interrupt polling was enabled in FreeBSD and set to 1000 interrupts per second. This technique will effect the ping time of packets by adding at most 2 ms to each ping, but has allowed us to run 4 downlink flows with more than 50% idle processor time.

The Realtek driver supports hardware VLAN tagging, which can be used to set the Priority Code Point (PCP) in the VLAN tag. The PLC standards specify how the VLAN PCPs are mapped to various PLC access categories. We modified the Realtek driver to allow us to manually set the PCP for traffic generated by our nodes, effectively allowing us to assign test traffic to different access categories.

B. Measurement Methods

The measurement software iperf was used for TCP and UDP tests. For TCP, the window size was set to 3600 kB. This ensures that flows will not be window-limited and that congestion control mechanism will be active on the PLC network.

In many networks, we expect there to be a node acting as a gateway between the PLC network and the Internet. This node plays a similar role to an access point in a wireless network. With this in mind, our traffic tests are done with uplink (Figure 2(a)) and downlink (Figure 2(b)) flows. The distinction between the two tests are that for uplink tests, the access point receives multiple flows, while the downlink must source multiple flows. Potential problems can occur in the uplink for TCP flows if the ACKs cannot access the channel due to contention issues [16].

IV. EXPERIMENTAL OBSERVATIONS

In the following, some preliminary experimental results are presented on our 5 node power line test bed. Future work will include analysis for each of the individual problems or issues.

A. Approximate Throughput

To establish a base point for the expected throughput of a UDP flow in a PLC HomePlug AV network, a quick back-of-the-envelope calculation can be done. It is worth noting

that this calculation is considered as an upper bound to our experimental observations, as our transmission medium is not perfect. First, we assume that a single UDP flow will access the network in contention mode, and so the contention period can fill the entire interval between beacons. Each channel access consists of the following intervals: two priority resolution symbols of a total $71.68 \mu s$, contention period (variable duration), preamble of $51.2 \mu s$, frame control of $59.28 \mu s$, data (variable duration), response interframe space (RIFS) of $140 \mu s$, MAC ACK of $70.48 \mu s$ (includes the preamble and frame control solely as a delimiter) and a contention interframe space (CIFS) of $100 \mu s$. We further assume a worst case contention period for the highest priority (CA3), which means that the transmission will consistently back-off by 7 slot times (each $35.84 \mu s$) before transmission.

The frame length for an IEEE 1901 packet is limited to $2501.12 \mu s$ in CSMA mode. With a guard time interval of $7.56 \mu s$ and an OFDM symbol duration of $40.96 \mu s$, this will allow for approximately 50 OFDM symbols. If the channel is perfect, and we can use the maximum 4096 QAM symbol in each of the 917 sub-30 MHz OFDM sub-carriers, then each packet can carry 550,200 uncoded bits. For the active cycle of $2501.12 \mu s$ and a total cycle of $3244.6 \mu s$, we can expect a raw throughput of 169.56 Mbps. Coupled with MAC layer overheads of 28 bytes for a 1528 byte frame, PHY block overheads of 8 bytes for a 520 byte block, and a FEC rate of either $R_{FEC} = 1/2$ or $16/21$, we get an estimate for our maximum achievable throughput as

$$S_{est} = (169.56 \text{ Mbps}) \left(\frac{1500}{1528} \right) \left(\frac{512}{520} \right) R_{FEC}.$$

If we assume $R_{FEC} = 16/21$ [1], we get an approximate throughput of 124.87 Mbps. Measured results from our test bed indicate that a single UDP flow can achieve approximately 92.5 Mbps with a MSDU payload size of 1400 bytes. The gap between our calculation and our experimental results can be attributed primarily to sub-carriers using less than the optimal 4096 QAM constellation and zero-padding.

B. Impact of additional flows

To verify the operation of our test bed, we used iperf to simulate TCP and UDP traffic. Our tests included uplink and downlink tests, and served both for verification and to identify MAC and PHY limitations of our test equipment. Pure Ethernet tests were performed as well to ensure that any artifacts we observe are due to the Devolo PLC modem rather than the rest of the test setup. For our initial tests, all flows were set to a default priority of CA1.

Figure 3 shows the trace results for 1 to 4 uplink TCP flows. As expected, as flows are added to the system, the per-flow throughput is reduced by a factor of n . Interestingly, for 3 and 4 flows, timeouts occur for short bursts: this suggests that a similar TCP ACK starvation phenomenon to that reported by [16] may be an issue for HomePlug AV and IEEE 1901 systems as well. Future testing aims include verifying this result using a larger test set: as described in [16], the downlink




Fig. 3. Traffic trace (in bits per second, per second averages) for 300 seconds worth of TCP uplink traffic and 1, 2, 3, or 4 PLC flows. The blue line is the sum throughput of all active flows.

TABLE I
TCP TRAFFIC FLOW

Number of flows	Average sum throughput
1	71.1 Mbps
2	71.3 Mbps
3	69.9 Mbps
4	65.2 Mbps

node must transmit $n/2$ TCP ACKs for n uplink flows. While each source/uplink node has a single flow to buffer, the sink/downlink node has to deliver $n/2$ the number of channel accesses. The sum throughput hits a maximum of around 70 Mbps (see Table I for average throughputs). Apart from short (10 cm) unshielded wires, our communications medium is isolated from any outside noise or interference, and we can assume that 70 Mbps is close to the maximum TCP traffic throughput of our HomePlug AV test bed.

Figure 4 shows the trace for 1 to 4 uplink UDP flows, where each flow is set to transmit at 30 Mbps. For 1 and 2 flows, the per-flow throughput stays close to 30 Mbps. For 3 and 4 flows, the per-flow throughput drops accordingly and in a similar manner as for TCP, although the sum throughput is higher than TCP due to TCP's protocol overheads.

Figure 5 shows traces for 1 to 4 downlink TCP flows. This corresponds to a single node transmitting multiple flows to a number of data sinks, and thus the data source does not need to contend with any congestion on the power line medium. Transmission errors are due to buffer overflow on account of throughput saturation on the power line.

Figure 6 shows traces for 1 to 4 downlink UDP flows. Each flow is set to transmit at 30 Mbps. For 1, 2, and 3 flows, the medium is capable of delivering 30 Mbps without fault. Separately, the peak UDP throughput of our test bed was verified with a single flow to be around 92.25 Mbps.

C. Ping test results

A common method to test the delay characteristics of networks is to send ping packets and measure the round-




Fig. 4. Traffic trace (in bits per second, per second averages) for 300 seconds worth of UDP uplink traffic and 1, 2, 3, or 4 PLC flows. Each UDP flow is set to 30 Mbps.

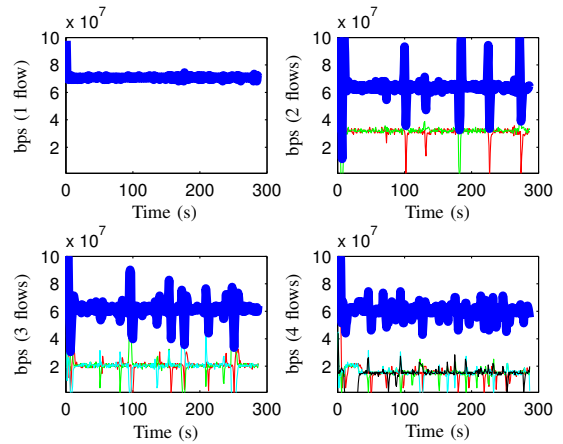


Fig. 5. Traffic trace (in bits per second, per second averages) for 300 seconds worth of TCP downlink traffic and 1, 2, 3, or 4 PLC flows. The blue line is the sum throughput of all active flows.




Fig. 6. Traffic trace (in bits per second, per second averages) for 300 seconds worth of UDP downlink traffic and 1, 2, 3, or 4 PLC flows. The blue line is the sum throughput of all active flows.




Fig. 7. Ping test packet sizes versus RTT. Error bars correspond to 1 standard deviation.




Fig. 8. Ping test packet capture via spectrum analyser in zero-span mode, with 460 byte packet sizes. Data symbol duration of $109.1 \mu s$ corresponds to 2 symbols.

trip return time (RTT) for each transmission. Ping packets take a special form, defined as Internet Control Messages Protocol (ICMP) packets. By controlling the size of the ping packets, we can get a measure for how a PLC network will handle traffic with different MAC service data unit (MSDU) sizes. An experiment was done between two PLC nodes with results displayed in Figure 7. Our parameters included variable sized ping packets, 50 ms spacing, averaged over 1000 trails, and default priority of CA1. Interrupt polling was disabled to prevent any distortion of the ping values. For verification, a similar experiment via the Ethernet interface was performed. In terms of the expected RTT duration from 200 Mbps PLC networks, the results correspond favourably to results shown in [4]. However, apart from the obvious overhead of using PLC compared to Ethernet (around a factor of 10 difference), there is a very abrupt drop in the RTT between 460 and 470 byte packets, with the lower curve showing the detail.

The reason for this drop has been narrowed down to the transition between transmitting 2 and 3 FFT symbols: it was verified by a spectrum analyser in zero-span mode that when the ping size was increased from 460 to 470 bytes, the PHY layer required an increase in the number of FFT symbols, and the frame increased from $109.1 \mu s$ to $158.8 \mu s$ (note that an OFDM symbol duration is $40.96 \mu s$ plus a guard interval). Example spectrum analyser screen captures of each ping process are given in Figures 8 and 9 for 460 bytes and 470 bytes, respectively. So far, the exact reason for this deviation has not been identified, and whether it is an implementation or PHY/MAC issue is still under investigation.




Fig. 9. Ping test packet capture via spectrum analyser in zero-span mode, with 470 byte packet sizes. Data symbol duration of $158.8 \mu s$ corresponds to 3 symbols.




Fig. 10. Traffic trace (in bits per second, per second averages) for 300 seconds worth of UDP downlink experiments. Out of 4 flows, there are 1, 2, 3 or 4 high-priority (hp) flows. Red trace is high-priority flow, blue trace is low-priority flow.

D. Throughput Comparison and Priority Resolution

The priority resolution mechanism in the HomePlug and IEEE 1901 standards provides a means to pre-empt all packets with lower priorities. During contention, only packets from the highest advertised priority level are permitted to transmit. This is contrary to the conventions in 802.11e, where flows with higher priority have a lower CW_{min} or $AIFS$, and thus while they have a higher probability of transmitting when the medium is busy, high priority packets are not certain to access the medium before other packets.

Figure 10 shows the results of UDP downlink experiments with a varying number of high priority flows. In this experiment, low-priority flows are marked as class CA1, while high-priority are marked as CA3 [11]. All UDP flows are set to transmit at 90 Mbps to ensure medium saturation. As expected, the high-priority flows severely restrict the traffic of the low-priority flows, with higher restriction as more flows are made high-priority. Table II shows the average throughput for each experiment.

Looking at the effect on TCP flows, Figure 11 shows similar experiments with low-priority TCP ACKs. The traces




Fig. 11. Traffic trace (in bits per second, per second averages) for 300 seconds worth of TCP downlink traffic. Out of 4 flows, there are 1, 2, 3 or 4 high-priority flows. Red trace is high-priority flow, blue trace is low-priority flow. TCP ACKs are not prioritised.




Fig. 12. Traffic trace (in bits per second, per second averages) for 300 seconds worth of TCP downlink traffic. Out of 4 flows, there are 1, 2, 3 or 4 high-priority flows. Red trace is high-priority flow, blue trace is low-priority flow. TCP ACKs are prioritised.

show TCP stalling frequently, indicated by abrupt drops in throughput. These are caused by unprioritised TCP ACKs, since the downstream TCP ACKs must contend for channel access with the upstream flows. This has been verified in further simulations (Figure 12) which use prioritised TCP ACKs. The number of discontinuities are greatly reduced, as has the throughput of the low-priority flows. This demonstrates that caution must be used to ensure that TCP ACKs are set to the same priority as the associated data flow.

V. CONCLUSION

In this paper, we have described the construction of a network PLC test bed using commodity equipment. Several issues that arose while constructing the test bed were highlighted, including setting the media type and ensuring that the PLC cards did not generate excessive interrupts. Results of some basic

TABLE II
UDP TRAFFIC FLOW WITH MULTIPLE PRIORITIES (ALL VALUES IN MBPS)

CA1 flows	CA3 flows	CA1 throughput	CA3 throughput
3	1	1.55	82.90
2	2	0.99	39.15
1	3	0.26	24.52
0	4	-	17.21

validation experiments for the test bed were presented, which indicate that the test bed is operating correctly in most regards. Some features of the test bed that require further investigation were also demonstrated, such as ping time anomalies and the interaction between MAC and TCP mechanisms. Such issues must be considered to allow network users to make effective use of PLC communication technologies.

REFERENCES

- [1] Broadband Over Power Lines PHY/MAC Working Group, "IEEE P1901/D4.01 Draft Standard for Broadband over Power Line Networks: Medium Access Control and Physical Layer Specifications," IEEE, Tech. Rep., 2010.
- [2] P. L. So and Y. H. Ma, "Development of a Test Bed for Power Line Communications," *IEEE Trans. on Consumer Elec.*, vol. 50, no. 4, pp. 1174–1182, 2004.
- [3] N. Weling, "Field Analysis of 40,000 PLC Channels to Evaluate the Potentials for Adaptive Transmit Power Management," in *IEEE Int. Symp. on Power Line Commun. and Its Applications*, 2010, pp. 201–206.
- [4] B. Jensen, H. Slavensky, and S. Kjaersgaard, "Benchmarking and QoS of In-House Powerline Equipment for AV Streaming Applications," in *IEEE Int. Symp. on Power Line Commun. and Its Applications*, 2006, pp. 160–165.
- [5] V. Gutierrez, J. A. Galache, R. Aguero, and L. Munoz, "Performance Evaluation of a PLC Field Trial: The Behavior of the TCP/IP Stack," in *IEEE Int. Symp. on Power Line Commun. and Its Applications*, 2007, pp. 378–384.
- [6] D.-E. Meddour, A. Kortebi, and R. Boutaba, "Mesh-based Broadband Home Network Solution: Setup and Experiments," in *IEEE Int. Conf. on Commun.*, 2010, pp. 1–5.
- [7] 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.
- [8] HomePlug Powerline Alliance, "HomePlug AV White Paper," <http://www.homeplug.org/tech/whitepapers/>, 2005.
- [9] Standards and Technology Department, "Medium-Speed (up to 14 Mbps) Power Line Communications (PLC) Modems using Windowed OFDM," Telecommunications Industry Association, Tech. Rep., 2008.
- [10] 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.
- [11] 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.
- [12] 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.
- [13] 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.
- [14] S. Katar, L. Yonge, R. Newman, and H. Latchman, "Efficient Framing and ARQ for High-Speed PLC Systems," in *IEEE Int. Symp. on Power Line Commun. and Its Applications*, 2005, pp. 27–31.
- [15] K. H. Afkhamie, S. Katar, L. Yonge, and R. Newman, "An Overview of the Upcoming HomePlug AV Standard," in *IEEE Int. Symp. on Power Line Commun. and Its Applications*, 2005, pp. 400–404.
- [16] D. Malone, D. J. Leith, A. Aggarwal, and I. Dangerfield, "Spurious TCP Timeouts in 802.11 Networks," in *Workshop on Wireless Network Measurement (WiNMee)*, 2008, pp. 43–49.