Electri-Fi Your Data: Measuring and Combining Power-Line Communications with WiFi

[Technical Report]

Christina Vlachou, Sébastien Henri, Patrick Thiran
EPFL, Switzerland
firstname.lastname@epfl.ch

ABSTRACT

Power-line communication (PLC) is widely used as it offers high data-rates and forms a network over electrical wiring, an existing and ubiquitous infrastructure. PLC is increasingly being deployed in hybrid networks that combine multiple technologies, the most popular among which is WiFi. However, so far, it is not clear to which extent PLC can boost network performance or how hybrid implementations can exploit to the fullest this technology. Motivated by questions such as which medium should an application use, we explore the spatial and temporal variations of WiFi and PLC.

Despite the potential of PLC and its vast deployment in commercial products, little is known about its performance. To route or load balance traffic in hybrid networks, a solid understanding of PLC and its link metrics is required. We conduct experiments in a testbed of more than 140 links. We introduce link metrics that are crucial for studying PLC and that are required for quality-aware algorithms by recent standardizations of hybrid networks. We explore the spatial and temporal variation of PLC channels, showing that they are highly asymmetric and that link quality and link-metric temporal variability are strongly correlated. Based on our variation study, we propose and validate a capacity estimation technique via a metric that only uses the frame header. We also focus on retransmissions due to channel errors or to contention, a metric related to delay, and examine the sensitivity of metrics to background traffic. Our performance evaluation provides insight into the implementation of hybrid networks; we ease the intricacies of understanding the performance characteristics of the PHY and MAC layers.

1. INTRODUCTION

Wireless technology is dominant in local networks; it offers mobility and attractive data-rates. Nevertheless, it often leaves "blind spots" in coverage, and the network becomes saturated because of the increasing demand for higher rates and of the explosion of network applications. Today's networks call for additional, simple technologies that can boost network performance, extend coverage, and improve quality of service. Several candidates, among which power-line and coaxial communications, are on the market. As the demand for combining diverse technologies increases, new specifications for hybrid networks are developed, such as the IEEE 1905 standard [2] which specifies abstraction layers for topology, link metrics, and forwarding rules.

Due to the growing demand of reliability in home networks, wireless and power-line communications (PLC) are combined by several vendors to deliver high rates and broad coverage without blind spots. PLC is at the forefront of home networking, as it provides easy and high data-rate connectivity. Its main advantage is coverage wider than WiFi and data-rates up to 1Gbps without requiring the wiring of a new network. It is obvious that PLC can be a lucrative backbone for WiFi. However, in the quest to provide *reliable* performance, some questions arise: Where and when does PLC perform better than WiFi? How fast does PLC channel quality change compared to WiFi? Which medium(s) should an application use? Such questions remain unanswered as of today and a goal of this work is to address them.

Despite its wide adoption, PLC has received far too little attention from the research community. Moreover, IEEE 1905 is technology agnostic and it does not provide any forwarding nor metric-estimation methods. To fully exploit the potential of each medium, hybrid networks require routing and load-balancing algorithms. In turn, these algorithms require accurate capacity estimation methods, and a solid understanding of the underlying layers of each network technology. To the best of our knowledge, there has not been any study on PLC; so far a very large body of work has only introduced theoretical channel models. In this work, we investigate PLC from an end-user perspective, and we explore link metrics and their variations with respect to space, time, and background traffic; this is our main contribution. We focus on two metrics required by IEEE 1905 [2]: the *PHY rate* (capacity) and the *packet errors* (loss rate).

The most popular specification for high data-rate PLC, employed by 95% of PLC devices [1], is HomePlug AV[1]. This specification was adopted by the IEEE

[1]HomePlug alliance is the leader in PLC standardization [1]. In addition to high data-rate PLC, there are low-rate specifi-

1901 standard [6]. In this work, we dig deeply into the 1901 performance and provide link-quality estimation techniques. We first present the key elements of the PHY and MAC layers in Section 2, and we detail on our measurement methodology for PLC in Section 3. In Section 4, we explore experimentally the gains of incorporating PLC in a WiFi network and explain why temporal variation studies are crucial for a reliable performance. We focus on WiFi blind spots and bad links and discuss how PLC can mitigate high-traffic scenarios.

We delve into both the PHY and MAC layers of PLC via a testbed of more than 140 links. In Section 5, we investigate the spatial variation of PLC and find that PLC links are highly asymmetric. This has two consequences: (*i*) Link metrics should be carefully estimated in both directions; (*ii*) Predicting which PLC links will be good is challenging. We study the temporal variation of the PLC channel in Section 6, and distinguish three different timescales for the link quality. Exploring temporal variation is important for exploiting to its fullest extent each medium and for efficiently updating link metrics (e.g., high-frequency probing yields good estimations but high overhead). In Section 7, we explore the accuracy of a capacity-estimation technique by designing a load-balancing algorithm and by employing our temporal-variation study. To explore the 1905 metric related to packet losses, we examine the retransmission procedure and how link metrics are affected by contention in Section 8. By employing our temporal variation study and our two link metrics, PLC performance can be fully characterized and simulated, thus reducing the overhead complexity of the exact representation of the channel model and the PHY layer mechanisms. We verify our findings by using devices from two vendors and HomePlug technologies. Our key findings and contributions are outlined in Table 1.

WiFi vs PLC	Section
In short distances, WiFi yields higher throughput, but with much higher variability, compared to PLC.	4.1
PLC usually offers high gains in quality of service enhancements, coverage extension and link aggregation.	4.1, 7.2
Capacity estimation methods and temporal variation studies are needed to fully exploit the mediums.	7.2
Channel Quality and Variation	**Section**
PLC links can exhibit severe asymmetry and spatial variation is difficult to predict.	5
Temporal variation of the PLC channel occurs over three different time-scales.	6
Variation on the long-term depends on the appliances and their power consumption.	6.3
Variation on the short-term depends on the noise produced by electrical appliances.	6.2
Link quality and link metric variation are strongly correlated and good links can be probed much less often than bad ones.	6.2, 8.1
Introduction of metrics and guidelines for accurate capacity estimation, which is required by IEEE 1905 [2].	7
Retransmissions Due to Errors or Contention	**Section**
Discussion on metrics that use broadcast probing.	8.1
Expected transmission count (ETX) in PLC.	8.1
Sensitivity of link metrics to background traffic.	8.3

Table 1: Main findings and contributions

2. BACKGROUND ON PLC

We now recall the main features of the PHY and MAC layers for the most popular PLC specification, which is HomePlug AV (HPAV) equivalently, IEEE 1901 [6].

2.1 PHY Layer

The physical layer of HPAV is based on an OFDM scheme with 917 carriers in the 1.8-30 MHz frequency band. Each OFDM carrier can employ a different modulation scheme among BPSK, QPSK, 8/16/64/256/1024-QAM. In contrast, in WiFi technologies, such as 802.11n, all carriers employ the same scheme and the *modulation and coding scheme* (MCS) index is used for decoding the frame [3]. Because each carrier employs different modulation schemes, PLC stations exchange messages with the modulation per carrier, the forward error correction code (FEC) rate, and other PHY layer parameters [6]. The entity that defines these PHY options is called the *tone map*, and it is estimated during the channel estimation process. To do so, the source initially sends *sound* frames to the destination by using a default, robust modulation scheme that employs QPSK for all carriers. This scheme is used for the initial channel estimation and communication between two stations, but also for broadcast and multicast transmissions. The destination estimates the channel quality using the sound frames, then it determines and sends the tone map with a unique identification – which is analogous to MCS for 802.11n – back to the source. The destination can choose up to 7 tone maps: 6 tone maps for different sub-intervals of the AC line cycle called *slots*, and one default tone map. PLC uses multiple tone maps for the different sub-intervals of the AC line cycle, because the noise and impulse response of the channel are varying along the AC line cycle. Tone maps are updated dynamically, either when they expire (after 30 s) or when the error rate exceeds a threshold [6].

cations for home automation, such as HomePlug GreenPhy.

2.2 MAC Layer

We now review the MAC layer and describe its most important sub-functions.

Physical Blocks (PB): The MAC layer employs two-level frame aggregation. First, the data are organized in *physical blocks* (PB) of 512 bytes, then the PBs are merged into PLC frames. The PLC frames themselves can be parts of *bursts*, with each burst containing up to 3 frames. A selective acknowledgment (SACK) of the PLC frame acknowledges each PB, so that only the corrupted PBs are retransmitted.

Access Methods: The MAC layer of HomePlug AV includes both TDMA and CSMA/CA protocols [6]. To the best of our knowledge, all current commercial devices implement only CSMA/CA. The CSMA/CA protocol is similar to 802.11 for wireless communications, but with important differences that are summarized in [18]. The main difference is that, contrary to WiFi, PLC stations increase their contention windows not only after a collision, but also after sensing the medium busy. One of the main consequences is short-term unfairness that might yield high jitter [18].

Management Messages (MMs): Management messages are a key feature of PLC. They are used for network management, tone-map establishment and updating. Stations must exchange MMs each time the tone map is updated, because the source has to be identified for the modulation scheme of each carrier.

Vendor-Specific Mechanisms: The IEEE 1901 standard leaves the implementation of some mechanisms, such as the channel estimation procedure described in Section 2.1, unspecified. Therefore, they are vendor-specific and so far, vendors have not released any detailed specification for their devices.

In addition to MMs specified by the standard [6], there are vendor-specific MMs. Vendor-specific MMs are employed to configure the devices, modify the firmware, or measure statistics. We use vendor-specific MMs to measure statistics or configure the devices, as described in the next section.

Start-of-Frame Delimiter (SoF): The frame control, or the start-of-frame (SoF) delimiter, of PLC contains information for both PHY and MAC layers. The bit loading estimate (BLE) is retrieved from the SoF delimiter and is an estimation for the capacity, as we observe in Section 7. The BLE is an estimation of the number of bits that can be carried on the channel per μs.

DEFINITION 1. *[6] Let T_{sym} be the OFDM symbol length in μs (including the guard interval), R be the FEC code rate, and PB_{err} be the PB error rate (chosen based on the expected PB error rate on the link when a new tone map is generated and shall remain fixed until the tone map becomes invalidated by a newer tone map). Let also B represent the sum of the number of bits per symbol over all carriers. Then, BLE is given by*

$$BLE = \frac{B \times R \times (1 - PB_{err})}{T_{sym}}. \quad (1)$$

Today's home networks, running 802.11n and/or 1901, contain fields in the frame header that help the receiver decode the frame and that accurately estimate capacity. We successfully employ these fields to aggregate bandwidth between the two mediums in Section 7. In the following, our PLC link metrics will be BLE and PB_{err}.

3. EXPERIMENTAL FRAMEWORK

We describe the experimental settings used to produce the measurements presented throughout this work. We provide guidelines for configuring and for obtaining various metrics from PLC devices.

3.1 Testbed and Setup

Our main testbed consists of 19 Alix 2D2 boards running the Openwrt Linux distribution [4]. The boards are equipped with a HomePlug AV miniPCI card (Intellon INT6300 chip), which interacts with the kernel through a Realtek Ethernet driver and with an Atheros AR9220 wireless interface. All our stations are placed on the same floor of a university building with offices. Figure 1 represents a map of the testbed along with the electrical map of the floor.

We next explain the PLC network structure. PLC uses a centralized authority called the central coordinator (CCo) to manage the network. To operate, each station must join a network with a CCo. Usually, the CCo is the first station plugged and it can change dynamically if another station has better channel capabilities than it does. Our floor has two electrical tables that are connected with each other at the basement of the building. This means that the cable distance between the two tables (more than 200m) makes the PLC communication between two stations at different tables challenging. Due to the two electrical tables, none of the stations can communicate with all stations and be the CCo. Hence, we create two different networks, shown with different colors in Figure 1. To avoid modifications in the network structure, we set the CCo statically in our testbed using [5], a tool described in the next subsection. These networks have different encryption keys (there is encryption on the MAC layer) and thus, only stations belonging to the same network can communicate with each other. In total, 144 links are formed.

In addition to using our main testbed, we experiment and validate our findings with HPAV500 devices, the Netgear XAVB5101 (Atheros QCA7400 chip)[2]. Due to space constraints, results are presented for our main testbed, unless otherwise stated.

3.2 Measurement and Traffic Tools

To retrieve the metrics for the PHY and MAC performance evaluation, we use a tool that interacts with the HomePlug AV chips, i.e., the Atheros Open Powerline Toolkit [5][3]. The tool uses vendor-specific management messages (MMs), as described in Section 2.2, to interact with, and to configure the devices. It also enables a sniffer mode with which we can capture the SoF delim-

[2]Note that, compared to HPAV described in Section 2.1, HPAV500 extends the bandwidth to 1.8-68 MHz.

[3]We have been equipped with devices from 6 vendors and have been able to retrieve statistics from all devices using [5].

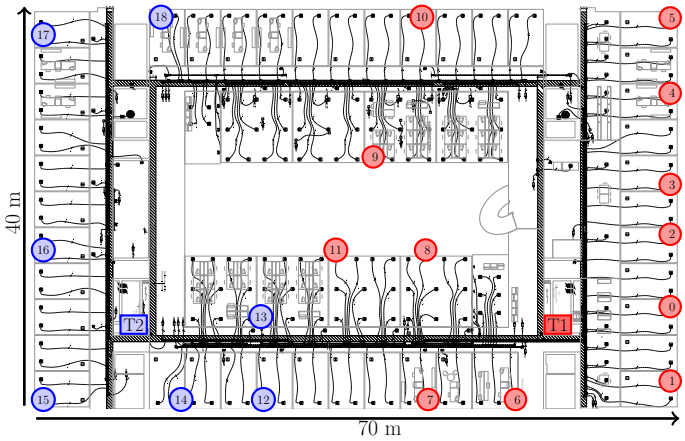

Figure 1: The electrical plan and the stations (0-18) of our testbed. There are two different PLC networks with CCo's at stations 11 and 15. Stations marked with the same color belong to the same network and are connected to the same electrical table (either T1 or T2).

iters of all received PLC frames. To generate traffic, we use `iperf`. For all the experiments, links are saturated with UDP traffic (unless otherwise stated), i.e., stations transmit at maximum available rates, so that we can measure metrics such as capacity. All the experiments of this work have been repeated multiple times over a period of one year to make sure that similar results are reproduced. Table 2 outlines the metrics used throughout this work, as well as the methods used to measure them.

Metric	Notation	Measured with
Arrival timestamp	t	SoF delimiter
Bit loading estimate	BLE	SoF delimiter
Burst length	BL	SoF delimiter
PB error probability	PB_{err}	MM (`ampstat` [5])
Average BLE	BLE	MM (`int6krate` [5])
Throughput	T	`iperf` or `ifstat`
MCS index (WiFi)	MCS	WiFi frame control

Table 2: Metrics and measurement methods

We are now ready to present our study on PLC.

4. WIFI VS PLC AND CHALLENGES

We first study the spatiotemporal variation of WiFi vs PLC in order to explore the possibilities of combining the two mediums towards quality of service improvement, coverage extension and bandwidth aggregation. We then discuss the challenges of hybrid implementations.

4.1 Spatial Variation: WiFi vs PLC

We first compare the spatial variation of Wifi and PLC in our testbed, with WiFi and PLC interfaces having similar nominal capacities[4]. This study quantifies the gains that PLC can yield in situations with wireless "blind spots" or bad links and also examines which medium an application should use. We conduct the following experiment: For each pair of stations, we measure the available throughput of both mediums back-to-back for 5 minutes, at 100ms intervals. These experiments are carried out during working hours to emulate a realistic home/office environment. We show the average and standard deviation of these measurements (for links with a non-zero throughput for at least one medium).

Let T_W and σ_W be, respectively, the average value and standard deviation of throughput for WiFi (T_P and σ_P, respectively, for PLC). Figure 2 illustrates the results of our experiment. Our key findings are as follows.

Connectivity: PLC yields a better connectivity than WiFi. 100% of station pairs that are connected with WiFi are also connected with PLC. In contrast, 81% of station pairs that are connected by PLC links, are also connected by WiFi links. At long distance (more than 35m), there is no wireless connectivity whereas PLC offers up to 41 Mbps. Thus, PLC can eliminate, to a large extent, blind spots.

Average performance: 52% of the station pairs exhibit throughput higher with PLC than with WiFi. PLC can achieve throughput up to 18 times higher than WiFi (40.1 vs 2.2Mbps). The maximum gain of WiFi vs PLC was similar, i.e., 12 times (46.3 vs 3.8Mbps).

Variability: At short distances (less than 15m), WiFi usually yields higher throughput, but PLC offers significantly lower variance. WiFi has higher variability with the maximum standard deviation of throughput being $\sigma_W = 19.2$ Mbps vs $\sigma_P = 3.8$ Mbps for PLC. The vast majority of PLC links yield a σ_P smaller than 4 Mbps.

Conclusion: At long distances, PLC eliminates wireless blind spots or bad links, yielding notable gains. At short distances, although WiFi provides higher throughput, PLC provides significantly lower variance, which can be beneficial for TCP or applications with demanding, constant-rate requirements, such as high-definition streaming. We explain this difference by the ability of PLC to adapt each carrier to a different modulation scheme, contrary to WiFi (see Section 2.1). PLC reacts more efficiently to bursty errors than WiFi, which has to lower the modulation scheme at all carriers.

The spatial variation of WiFi has been extensively studied (e.g., [14]). However, very few works exist on PLC; [12] focuses on a much older technology, and, due to the insufficient literature on specifications, [13] treats

[4]We use 802.11n, with 2 spatial streams, 20MHz bandwidth and 400 ns guard interval, yielding a maximum PHY rate of 130 Mbps. We selected a frequency that does not interfere with other wireless networks in our building. The highest PLC data-rate is 150 Mbps hence, both interfaces have similar nominal capacities. This is confirmed by the maximum throughputs exhibited by both mediums, shown in Figure 2.

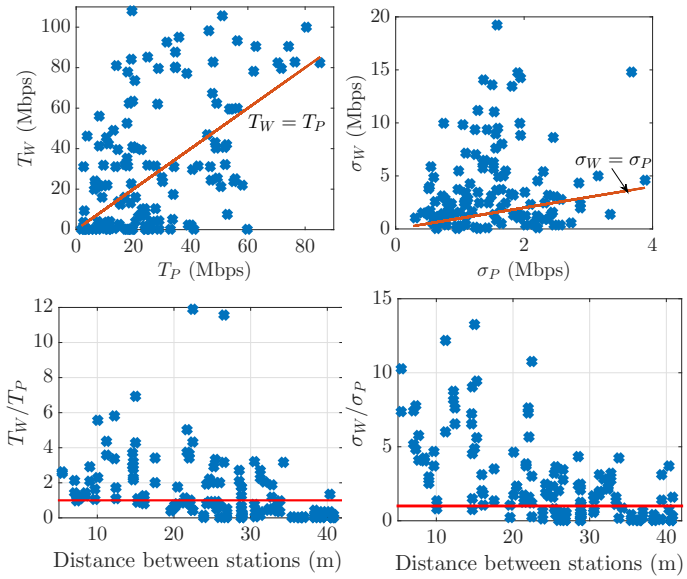

Figure 2: WiFi vs PLC performance for all links (top). Spatial variation of the performance ratio between WiFi and PLC (bottom).

PLC as a black-box and focuses on average performance and not on variability. Using two PLC technologies, we also investigate spatial variation in Section 5. Our study cross-validates the findings of [13], but in a network with more links, deployed on an area 4 times larger. The following subsection provides a first step towards the comparison of WiFi and PLC the temporal variations.

4.2 Temporal Variation: WiFi vs PLC

We now look at the concurrent temporal variation of WiFi and PLC during working hours for a much longer duration than before. We are interested in exploring the timescales at which the two mediums vary. Figure 3 shows the capacity for concurrent tests on WiFi and PLC, estimated by using MCS and BLE respectively, and averaged over 50 packets. We observe that link 3-8, which is a good link, exhibits a variation much higher with WiFi than with PLC. Although we would expect channel changes due to switching electrical appliances in the building, the PLC link is almost not affected by people leaving the premises (around 6pm). The average link 3-0 varies more for both mediums.

These preliminary results imply that PLC has low variability for good links and high for bad links. To the best of our knowledge, there are not any temporal-variation studies of the end-to-end performance of PLC. In contrast, many studies have focused on WiFi temporal variation. In Section 6, we study the PLC temporal variation and we observe that the variability is high in timescales of hours, because of the variations of the electrical load. We notice however, that this variation is

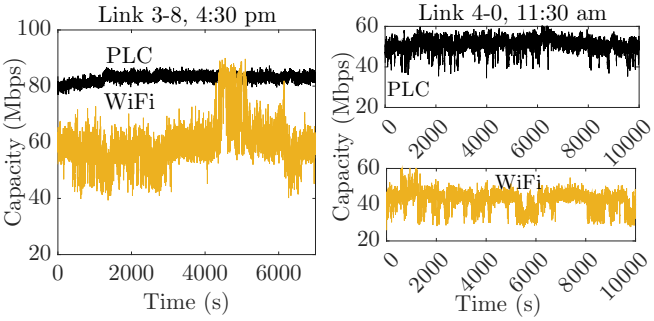

Figure 3: Temporal variation of capacity for PLC and WiFi for two links during working hours (started time written).

not significant, compared to the one of Wifi, and that it is high only for bad links.

4.3 Challenges in Hybrid Networks

As we observe in Section 4.1, although PLC boosts network performance, there are still a few links that perform poorly with both WiFi and PLC. As a result, mesh configurations, hence routing and load balancing algorithms, are needed for seamless connectivity in home or office environments. A challenge for these algorithms is that they have to deal with two different interference graphs with diverse spatiotemporal variation, and that, to fully exploit all mediums, they require accurate metrics for capacity and loss rates. To this end, unicast probes must be exchanged among the stations[5]. In a network of n stations, probing introduces an $\mathcal{O}(n^2)$ overhead that can be significantly reduced by employing temporal variation studies of each medium.

A significant challenge, highlighted by recent studies in 802.11n networks [16], is the accuracy of established quality metrics, such as the expected transmission count (ETX) or time (ETT) [8], in modern networks, i.e., 802.11n/ac. The authors in [16] show that due to the MAC/PHY enhancements introduced in 802.11n, these metrics perform poorly and that they should be revised, given that they have been evaluated only under 802.11a/b/g.

The above arguments raise a few questions: How often should the PLC link metrics, such as capacity, be updated in load-balancing or routing algorithms in order to achieve both small overhead and accurate estimation? How would ETX perform in PLC? We will answer these questions in Sections 5–8. In the rest of this work, we design link metrics for PLC and explore their variation with respect not only to time, but also to space and to background traffic.

5. SPATIAL VARIATION OF PLC

[5]Broadcast packets cannot be used to estimate capacity. See for example [7], [8].

We explore the spatial variation of PLC, as it is important for predicting coverage and good locations for PLC stations, and for implementing link metrics. We find that PLC is highly asymmetric, and this should be considered when estimating link metrics.

We first explain the main properties of the channel that affect both spatial and temporal variations. The two main components of PLC channel modeling are attenuation and noise. Consider an example of a simple electrical network with a transmitter (TX) and receiver (RX), as given in Figure 4. The main sources of attenuation and noise are the electrical appliances plugged in between. Modeled with dashed boxes in Figure 4, each connected appliance has an impedance and produces some noise that is usually non-Gaussian [9] and that depends on the device type.

The spatial variation of PLC is mainly affected by the position, the impedance, and the number of appliances connected to the network. When it comes to PLC, the electrical cable becomes a transmission line, with a characteristic impedance. The connection of appliances creates impedance mismatches to this transmission line, causing the transmitted signal to be reflected multiple times. For example, in Figure 4, at point M, we have an impedance mismatch and any signal s arriving at M is partly reflected (signal r) and partly propagates (signal t) towards the same direction as the original signal s. Reflections of signals at various impedance-mismatched points result in multiple versions of the initially transmitted signal arriving at different times at the receiver, thus establishing a multi-path channel for PLC. We will see in the next section that temporal variation is affected by multi-path effects, i.e., the appliances' impedance, in long-term timescales, whereas it is affected only by noise at short-term timescales.

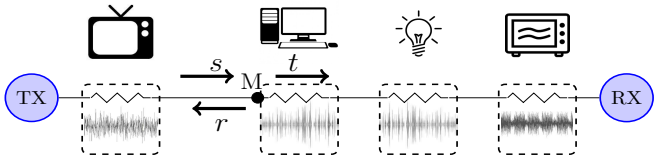

Figure 4: Multi-path and noise in PLC channels

A very important characteristic of power-line channels is that they exhibit performance asymmetry, i.e., capacity can differ significantly between the two directions of the link. In all the experiments we run (both with AV and AV500), we observe a performance asymmetry of more than 1.5x in approximately 30% of stations pairs in our testbed. Figure 5 presents typical examples of these links, for which the throughput in one direction is less than 60% of the throughput in the opposite direction. By re-conducting the experiments with AV500 devices, we verify that the asymmetry is not due to the hardware. Link asymmetry in PLC has been also observed in [13].

We attribute this asymmetry to a high electrical-load existing close to one of the two stations. In this case, the channel cannot be considered as symmetric and the two transmission directions in the link experience different attenuations.

In our tests with WiFi, presented in Section 4, we also observe that wireless channels can also exhibit asymmetry. However, compared to PLC, this occurs on a much smaller subset of links, and is much less severe (for instance, the asymmetry was up to 1.5x for good links and up to 3.5x for bad links). An asymmetry of loss rates has been found experimentally for residential WiFi networks in [14].

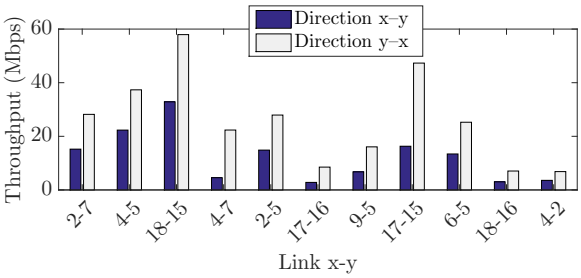

Figure 5: Throughput asymmetry in PLC links.

We now turn to our spatial variation study, where we use both AV and AV500. Figure 6 provides the available UDP throughput of *single* links as a function of the *cable* distance between the source and the destination of the traffic from a single experiment. There is a clear degradation of throughput as distance increases. However, because of the diversity in positions and types of connected appliances, there is a large range of possible throughputs at any specific distance. We observe that small distances (<30m) guarantee good links, but that large distances (30-100m) can yield either good or bad links. By comparing AV and AV500, we observe that AV500 enables some links with no AV connectivity to still enjoy a non-zero throughput, but with severe asymmetries (e.g., link 10-2 with 10x asymmetry).

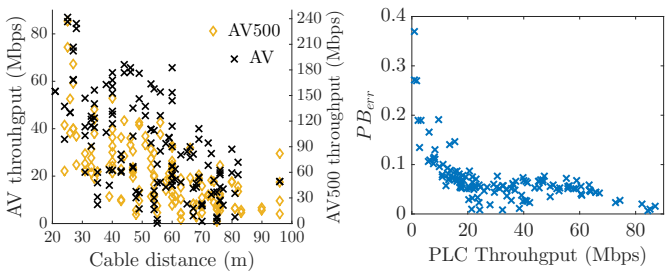

Figure 6: Throughput vs cable distance between source and destination for all links of the testbed (left). PB_{err} vs throughput for AV (right).

To further explore the causes that affect the attenuation, we run some experiments with two stations con-

nected by a long electrical cable and without any devices attached. We notice that the attenuation in an up to 70 m cable causes a throughput drop of at most 2 Mbps. The attenuation is therefore caused by the multi-path nature of the PLC channel. By plugging electrical appliances in this isolated experiment, we observe that asymmetry was introduced, as also found in [13].

Conclusion: The spatial variation of the PLC channel depends on two factors: (*i*) the structure of the electrical networks, i.e., the appliances attached and their position on the grid, (*ii*) the distance between the stations. PLC channels are very asymmetric and this is a key feature for their spatial variability.

To optimize performance not only in terms of throughput but also delay, hybrid networks need some estimation of the retransmissions a frame suffers due to channel errors. We now evaluate the relationship of the metric PB_{err} with the available throughput. Figure 6 illustrates PB_{err} vs the available throughput for all the links of our testbed. It shows that PB_{err} decreases as throughput increases, as expected. However, because the tone maps are updated based on this metric, some average links might have lower PB_{err} than the best links of the testbed. We further study the PB_{err} metric in Section 8, by delving into packet retransmissions. We show that PB_{err} can be used to predict the expected number of retransmissions due to errors.

6. TEMPORAL VARIATION OF PLC

Little is known about PLC temporal variation, and we observe in Section 4 that a temporal-variation study could improve the quality of service in hybrid networks and the accuracy of link metrics estimation. We now investigate the temporal variation of the PLC channel.

We examine separately the two main components of channel modeling, i.e., the variation of noise generated by the attached electrical appliances, and the variation of channel transfer function (or attenuation). We employ BLE to investigate the statistical properties of channels by using existing commercial devices. We show that BLE reflects the channel quality and the fundamental features of PLC channel modeling explained in Section 5. BLE captures the channel behavior because it is updated when the PB_{err} exceeds a threshold value.

We now discuss the timescales within which the channel varies. These timescales have been introduced for channel modeling and simulation in [15] (from which we borrow the terminology to name these timescales). We first focus on noise generated by electrical appliances. It has been shown by measurements, e.g., in [9], that the noise level varies across subintervals of the mains cycle, thus yielding the first scale governing PLC temporal variation (scale (*i*)). Due to the periodic nature of the mains, this noise also varies in a scale of multiples of the mains cycle and seconds, which results in another timescale for the temporal variation (scale (*ii*)).

We next focus on attenuation. As discussed in Section 5, attenuation is introduced due mainly to impedance mismatches in the transmission line (electrical cable) that are created by connected appliances. As expected, this attenuation changes when the structure of the electrical network changes hence, in scales of minutes or hours (scale (*iii*)). This variability strongly depends on peoples' usage of appliances and on switching the appliances, as this creates impulsive noise in the channel.

As hinted above, our study adopts an analysis of three timescales that is validated by our measurements in the following subsections. Our work differs from [15] in that we examine the channel quality from an end-user and practical perspective, exploring metrics affecting the end-to-end performance. The three timescales are as follows.

(*i*) **Invariance Scale:** subintervals of the mains cycle, such as the 6 tone-map slots of HPAV;

(*ii*) **Cycle Scale:** multiples of the mains cycles – depends on the noise produced by appliances;

(*iii*) **Random Scale:** minutes or hours – related to connection or switching of electrical appliances and depends on human activity.

We now introduce our variables, starting with some notations. For the invariance scale, we use the term tone-map slots for the subintervals of the mains cycle, as we can measure the channel quality with respect to tone-map slots by using PLC devices. Let L be the total number of tone-map slots of the mains cycle, with each slot s having a duration T_s, so that the total slots duration $\sum_{s=1}^{L} T_s$ is equal to half mains period (as specified in [6]). Let BLE_s, $1 \leq s \leq L$, denote the BLE of tone-map slot s. In order to study the channel with respect to the three scales defined above, we assume that time is discrete, with one time unit having real-time duration equal to the mains cycle. Let $\mu_s \in \mathbb{R}^+$ and $\sigma_s \in \mathbb{R}^+$, $1 \leq s \leq L$, represent the expected value and the standard deviation of BLE_s, and let ν_{σ_s} be a continuous random variable with 0 mean and variance equal to σ_s^2. In the cycle scale, the mean and variance of BLE_s, μ_s and σ_s, respectively, are considered to be constant, and the variation of BLE_s around its mean is described by ν_{σ_s}. In the random scale, μ_s and σ_s vary with time due to electrical load variability. Given the above, at any time step t, the channel quality is described as

$$BLE_s(t) = \mu_s(t) + \nu_{\sigma_s(t)}(t), 1 \leq s \leq L. \qquad (2)$$

The process ν_{σ_s} is different for each link and its distribution can be time-varying over the random scale for a specific link, due to the different types of operating appliances and to different channel transfer-functions.

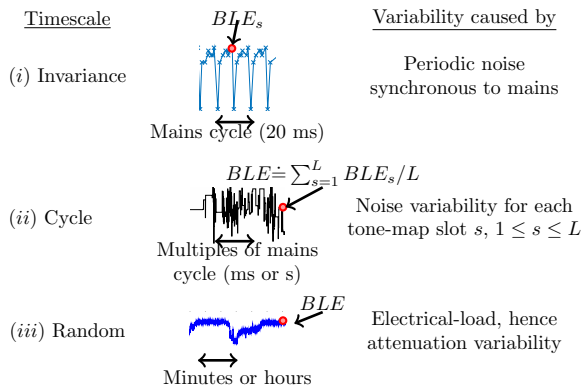

Figure 7: BLE temporal variation.

The distribution of ν_{σ_s} is link-dependent and its exact characterization is out of the scope of this work. In our study for cycle-scale variation, we will study how often the value of ν_{σ_s} changes and how σ_s behaves with respect to the link quality. Figure 7 illustrates the three timescales and the factors causing variability. We next examine each timescale.

6.1 Invariance Scale

The invariance scale of BLE is affected by the noise levels that the appliances produce at different subintervals of the mains cycle, and it has direct consequences on estimating link metrics. All our tests showed that noise has varying levels over the different tone-map slots. Figure 8 shows the instantaneous BLE_s from captured frames in typical examples of good and average links. We observe that in HPAV, the total duration of the 6 tone-map slots is equal to half of the mains cycle, thus BLE_s changes periodically, with a period of 10 ms. Each PLC frame uses a different BLE_s, depending on which tone-map slot its transmission takes place.

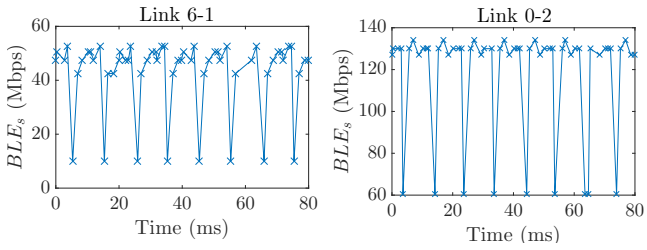

Figure 8: Invariance-scale variation of BLE from captured PLC frames of saturated traffic.

We highlight that this timescale is crucial for capacity estimation in PLC. With the examples of Figure 8, we observe that there might be significant variation along the mains cycle, even for good and average links. Thus, link metrics have to be estimated or averaged over all $L = 6$ tone-map slots. We next study the average BLE of all 6 slots to examine the variability of the average link-quality at longer time-scales, i.e., (ii), and (iii).

6.2 Cycle Scale

We now examine the average time during which the quality of the links is preserved in the cycle scale. This sheds light on the average length of probing intervals for link metrics, as there exists a tradeoff in probing: too large intervals might yield a non-accurate estimation, whereas too small intervals can generate high overhead.

We conduct experiments that last 4 minutes, over all links of the testbed. During each experiment, we request BLE_s, $1 \leq s \leq L$, every 50 ms, as this is the fastest rate at which we can currently send MMs to the PLC chip. As we need to avoid random changes in the channel due to switching electrical appliances, all the experiments of this subsection are conducted during nights or weekends (given the office environment). For the cycle scale variation of the channel, we assume that the electrical network structure is fixed.

Here, we evaluate the cycle-scale variation by using the average BLE over all tone-map slots, that is $BLE \doteq \sum_{s=1}^{6} BLE_s/6$. To optimize probing in hybrid networks, we compare the performance between good and bad links. Figure 9 presents the variation for typical good and bad links of our testbed. Observe that depending on their quality, links exhibit different behaviors. Our findings, validated not only by the representative examples shown here, but also by experiments over one year period in all the links of our testbed, are as follows.

Bad Links: Bad links, e.g., 11-4 and 6-5, tend to modify the tone maps much more often than good links do. Moreover, they yield a significantly higher standard deviation of BLE than good links.

Average Links: Average links, e.g., 18-15 and 1-2, vary less than bad links, and might preserve their tone maps for a few seconds. During periods when average links vary often, the standard deviation of BLE can be high, depending on the channel conditions.

Good Links: The tone maps of good links can be valid for several seconds, e.g., link 15-18. Good links that update often the tone maps, such as link 3-1, have insignificant increments or decrements, e.g., of up to 1%, or have impulsive drops of BLE, e.g., of up to 5%, with the channel estimation algorithm needing a few time-steps to converge back to the average BLE value.

Asymmetry in Temporal Variability: By observing links 15-18 and 18-15, we find that the asymmetry discussed in Section 5 translates not only in an average performance asymmetry, but also in a temporal-variation asymmetry.

Channel Estimation Algorithms: Temporal variation of link 15-18 is the same with HPAV and with HPAV500. By noticing the impulsive BLE drops in link 18-15 and by comparing HPAV with HPAV500, we detect a feature of the channel estimation algorithm that might be vendor-specific: The HPAV500 performance oscillation shows that the estimation algorithm returns

very low BLE values when bursty errors occur. This uncovers that temporal variation in PLC link quality might also depend on the channel estimation algorithm and future work should focus on comparing link-metric estimations for different vendors.

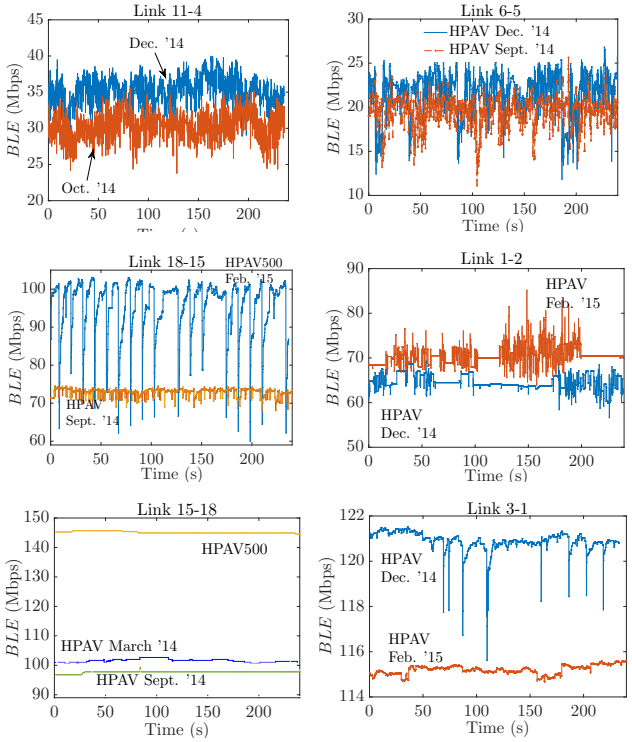

Figure 9: Examples of cycle-scale variation of BLE for links of various qualities

We next corroborate the above findings over all links of our testbed. Let α be the inter-arrival time of two consecutive BLE updates. Figure 10 shows the results of the average α values and the standard deviation of BLE for all links sorted by increasing BLE order, i.e., link quality. We observe that good links tend to update less often their tone maps, and also that BLE variability is smaller compared to bad links. Although some good links might update BLE at a similar frequency as bad links ($\sim$ 100 ms), as we discussed above, these links tend to have small increments and decrements of BLE, yielding a stable average performance over minutes and a low BLE standard deviation.

Conclusion: In cycle scales, that is seconds or minutes, good links should be probed less often than bad links to reduce overhead. The cycle-scale temporal variation unveils how link metrics should be updated depending on their quality.

6.3 Random Scale

In Sections 4.2 and 6.2, we observe that during timescales of minutes, PLC does not vary much, with a standard deviation of throughput up to 4 Mbps. We now look at longer timescales, i.e., in terms of minutes and hours,

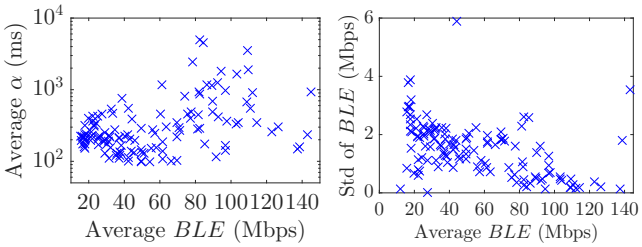

Figure 10: Cycle-scale variation of BLE with respect to the link quality (links are sorted with increasing average BLE order).

with two goals: (i) to examine whether some links could be probed at a slow rate, thus reducing overhead (ii) to characterize the variability of PLC performance in presence of high and low electrical loads. To study the channel quality variation over the random scale, we run tests over long periods, i.e., two days and two weeks, for various links. During these tests we measure the throughput, BLE, and PB_{error} every second. We now denote by μ the mean of $BLE = \sum_{s=1}^{6} BLE_s/6$, and by σ its standard deviation.

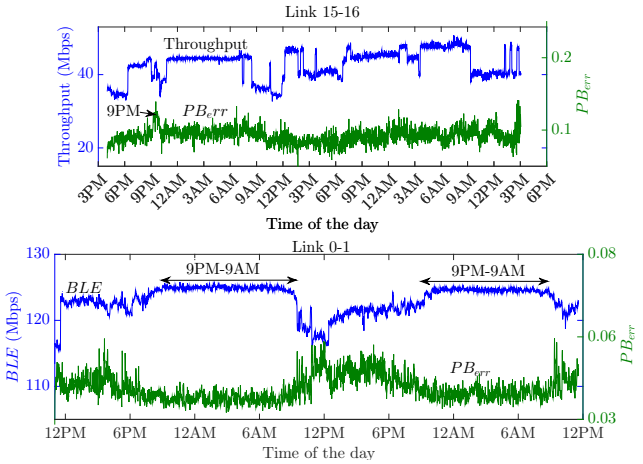

Figure 11: Random-scale variation of PLC over a total duration of 2 days. Metrics are averaged over 1 minute intervals. Every day at 9pm, all lights are turned off in our building, leading to a channel change for PLC.

Figures 11-13 show the results of our measurements. Our observations are as follows.

Link Quality vs Time: The variation of μ is governed by the electrical load. The larger the number of switched-on devices is (e.g., at working hours) the larger the attenuation is, and the lower μ is, as we have discussed in Section 5.

Link Quality vs Variability: Observe the differences in the y-axis scales in Figures 12 and 13 that represent a good and a bad link, respectively. For a given link, the random-scale variation of σ strongly de-

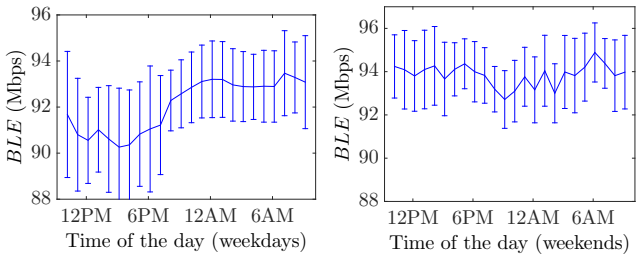

Figure 12: Random-scale variation of BLE for link 1-8 over 2 consecutive weeks. Lines represent the BLE averaged over the same hour of the day and error bars show standard deviation.

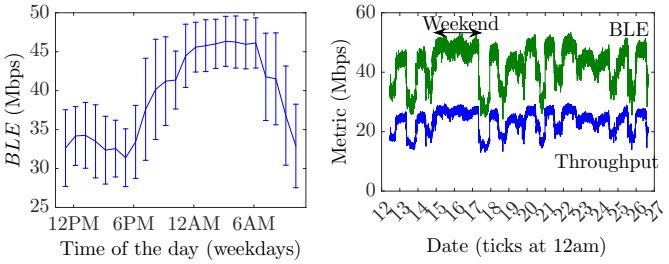

Figure 13: Random-scale variation of BLE for link 2-11 over 2 consecutive weeks in Nov. 2014.

pends on the noise of the electrical devices attached, and it is higher when μ is lower, as this implies that more devices are switched on and therefore, more noise is produced, or that devices are switched on/off more often, creating impulsive noise phenomena. σ is very small for good links; it increases as the link quality, i.e., μ, decreases.

Link Probing: Good links exhibit a negligible standard deviation, which implies that they can be probed every minute or hour, depending on the time of the day.

7. CAPACITY ESTIMATION PROCESS

We now explore a capacity estimation process for PLC. As mentioned in Section 2.1, stations estimate a tone map if and only if they have data to send. Thus, to estimate link metrics, a few unicast probe packets have to be sent[6]. In Section 6, we discuss the frequency of these probes given the link quality by sending saturated traffic. Here, we examine how capacity can be estimated by sending a few probe packets.

7.1 BLE as a Capacity Estimator

First, we show that BLE, which is included in the header of every PLC frame, accurately estimates the capacity of any PLC link. We repeat saturated tests for our 144 links and with a duration of 4 min. Figure 14 presents the measured throughput and BLE. We

[6]Traditional techniques for estimating bandwidth in WiFi include unicast packet-pair probing [7], [8].

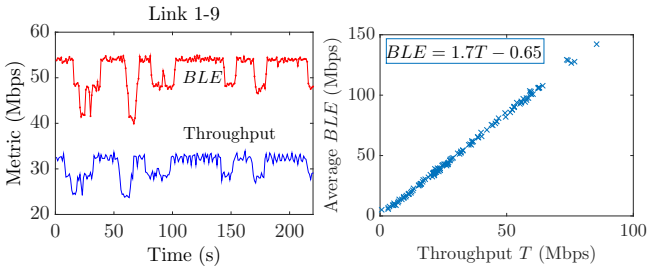

Figure 14: BLE and throughput averaged every 1 s, vs time, for link 1-9 (left). Average BLE vs throughput for all the links (right).

observe that BLE is an exact estimation of the actual throughput received by the application. Let T be the average throughput. Fitting a line to the data points, we get $BLE = 1.7T - 0.65$. We verified that the residuals are normally distributed.

We next discuss a capacity estimation technique that uses BLE and probe packets. To conduct a capacity estimation using BLE, a few packets per mains cycle and estimation interval should be captured, given our temporal variation study in Section 6.1. Here, we investigate an alternative technique that uses MMs to request the instantaneous BLE. The PLC devices provide statistics of the average BLE used over all 6 tone map slots. Probe packets need not to be sent at all sub-intervals of the mains cycle, because according to 1901 [6], the channel estimation process yields a BLE for all slots when at least 1 packet is sent.

We evaluate our capacity technique with probe packets and explore whether the number of the packets affects the estimation. For the purpose of this test, we reset the devices before every run. We perform experiments to estimate the capacity, by sending only a limited number of packets of size 1300B per second (1- 200 packets per second)[7]. Figure 15 shows that the estimated capacity converges to a value that does not depend on the number of packets sent; however, the number of packets sent per second affects the convergence time to the real estimation. We observe that the channel-estimation algorithm can have a large convergence time to the optimal allocation of bits per symbol for all the carriers, because it needs many samples from many PBs to estimate the error for every frequency, i.e., carrier. This convergence time depends on the (vendor-specific) channel-estimation algorithm and on the initial estimation (which was reset by us).

To evaluate the convergence time in realistic scenarios, we now perform a test in which we reset the devices at the beginning, but after 2000 s we pause the probing for approximately 7 minutes. Figure 16 shows the results of the experiments for various links. It turns out, that

[7]Note that the probe packets can be of any size. PLC always transmits at least a PB (500B), using padding.

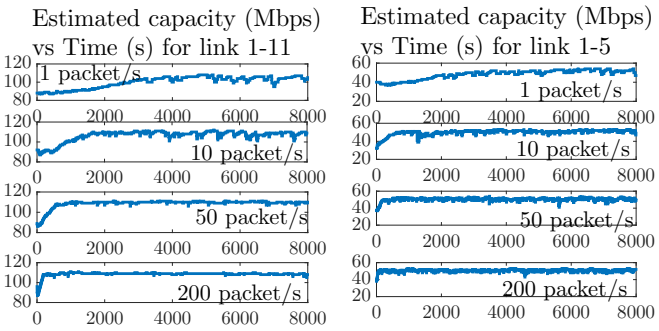

Figure 15: Estimated capacity for two links and different number of packet-probes per second.

the devices maintain the channel-estimation statistics, as the estimated capacity resumes from the previous value before stopping the probing process. Thus, the convergence time of the capacity estimation might not apply in realistic routing implementations for PLC.

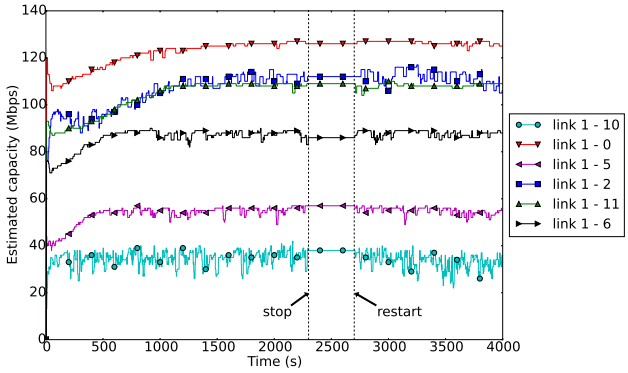

Figure 16: Estimated capacity for various links by probing with 20 packets per second. After 2000s, we stop the probing for 7 minutes.

Conclusion: Capacity should be estimated by sending probe packets and measuring BLE in PLC networks. To estimate capacity, given our study in Section 6.1, we have to take into account the invariance scale and to either compute the average $BLE = \sum_{s=1}^{6} BLE_s/6$, by capturing PLC frames or request it using MMs. In order to eliminate probing overhead, the probing frequency should be lower for good links and for periods of the day with low electrical load operating, as we observe in Sections 6.2 and 6.3. One of the remaining challenges in link-metric estimation is to take into account the technology-specific MAC mechanisms, such as frame aggregation. This remains a challenge also for the latest WiFi technologies, as highlighted in [16].

Next, we validate our capacity-estimation and temporal-variation studies by using a load-balancing algorithm.

7.2 Bandwidth Aggregation Using Capacity

To further validate our capacity estimation method, we employ a simple load-balancing algorithm that aggregates bandwidth between WiFi and PLC and operates between the IP and MAC layers. To implement our algorithm, we use the *Click Modular Router* [11]. We forward each IP packet to one of the mediums with a probability proportional to the capacity of the medium. At the destination, we reorder the packets according to a simple algorithm that checks the identification sequence of the IP header. We measure the jitter and compare it with the jitter when using only one interface, making sure that it does not worsen. Before presenting the performance results of our implementation, we explain the load-balancing and reordering algorithms.

Load-Balancing Algorithm: Our load-balancing algorithm assumes that has full information of the capacities is both mediums. We explain later how we estimate the capacities. Let C_W be the capacity of WiFi and C_P be the capacity of PLC. Then, for each packet that should be forwarded, the algorithm forwards the packet to WiFi (respectively PLC) with probability $C_W/(C_W + C_P)$ (respectively $C_P/(C_W + C_P)$).

Reordering Packets: Packets transmitted over from different mediums might experience different delays. Hence, while load-balancing, the original order of the packets is not preserved. We introduce a simple algorithm to reorder the out-of-order packets at the destination of a flow. To identify the packet order, we use the identification field of the IP header). According to the usual reordering approach, the out-of-order packets are saved in the destination's buffer for a time interval, i.e., there is a timeout for each packet. If this timeout expires, the packet is deleted and is considered as lost. Clearly, this timeout should depend on the link rates thus, it is not practical: If the timeout is very large, there are large delays when packets are lost (because the algorithm waits for these packets to arrive). Now, if the timeout is very small, the algorithm might consider some packets lost due to the timeout expiration, while the packets arrive later than this expiration.

For the reasons mentioned above, we design an algorithm that does not depend on timeouts for the packet reordering. The algorithm relies on the following fact: *Fact [A]: The packets that are transmitted over the same medium arrive at the destination in an increasing sequence-order. This assumption is valid if all interfaces use FIFO queues.*

Let us start with some necessary definitions. Our algorithm is executed each time a new packet arrives. Let p the packet just arrived in the input of the algorithm. We define as s the sequence of the p. Each packet can be forwarded, discarded, or stored for later forwarding. Let us define $last^{fwd}$ as the sequence of the last forwarded packet in the correct order. We also define $last_r^{seen}$ as the sequence of the last packet seen (forwarded, stored

or discarded) from medium r, $1 \leq r \leq n$, where we assume that we have n different mediums. Finally, we let R denote the index of the medium of packet p.

Now we are ready to present the algorithm. Packet p is forwarded if s is equal $last^{fwd} + 1$, i.e., if p arrives in the correct order with respect to the last forwarded packet. After forwarding p, the algorithm forwards all the stored packets with consecutive sequences, until it finds a sequence missing, where it terminates. Now, if $s < last^{fwd} + 1$, p is discarded, because it means that all the packets until $last^{fwd}$ have been processed (either discarded or forwarded). Finally, if $s > last^{fwd} + 1$, p has arrived earlier than packets with smaller sequences, and it should be stored. In this case, we have to check if we can forward some stored packets given the new information on the last sequence seen from medium R. To achieve this, the algorithm is based on [A]: If the sequence of the next packet that should be forwarded, i.e., $last^{fwd}+1$, is larger than at least one of the values of the sequence of the last packet seen per medium, then the packet $last^{fwd}+1$ can arrive from that medium thus, the algorithm waits for this packet. Conversely, if $last^{fwd}+1$ is smaller than each of the values of the sequence of the last packet seen per medium, then the packet $last^{fwd}+1$ is lost and cannot arrive from any medium. Therefore, the algorithm forwards all the stored packets (in increasing order) up to the minimum sequence of the last packet seen among all mediums, i.e., $\min_{1 \leq r \leq n} last_r^{seen}$. Algorithm 1 describes this procedure.

To estimate the capacities for our load balancing algorithm, we probe links with 1 packet per second and request BLE and MCS from the interfaces. The capacity for PLC is estimated using BLE, i.e., averaged over the 6 tone-map slots of the invariance scale, whereas for WiFi MCS capacity is averaged over the transmissions (data and probes) during every second, because, as we observe in Section 4.2, WiFi varies more than PLC within a second. For PLC, we take into account the relationship of BLE with respect to real throughput, thus the available throughput is estimated by dividing the value of BLE with 1.7. For WiFi, we perform a similar study to compute the relationship between maximum available throughput and estimated capacity. Figure 17, presents the results of estimated capacity vs available throughput for WiFi. We use this relationship to compute the available throughput i our implementation. Our load-balancing algorithm takes into account our temporal variation study on PLC: In Section 6.1, we uncover that the PLC channel quality is periodic, with every packet using a different BLE. Because an accurate synchronization at this time-scale is challenging for algorithms operating above the MAC layer (such as in IEEE 1905 standard), the capacity of PLC in hybrid networks has to be estimated by averaging over the invariance scale.

In Figure 18, we first present the throughput of exper-

Algorithm 1: Algorithm for reordering out-of-order packets. The algorithm is executed whenever a new packet arrives.

```
1  Input: Packet p, s (Sequence of p), R (Medium
     over which p was transmitted);
   Data: last^fwd (sequence of the last forwarded
         packet), n (number of mediums), last_r^seen
         (sequence of last packet seen from medium
         r, 1 ≤ r ≤ n);
2  Set last_R^seen ← s;
3  if s == last^fwd + 1 then
4  |  Forward p;
5  |  Forward stored packets until finding a missing
   |  (non-consecutive) sequence;
6  |  Update last^fwd;
7  end
8  else if s < last^fwd + 1 then
9  |  Discard p
10 end
11 else
12 |  Store p;
13 |  Set should_forward ← true;
14 |  for r from 1 to n do
15 |  |  if last_r^seen < last^fwd + 1 then
16 |  |  |  Set should_forward ← false;
17 |  |  |  break
18 |  |  end
19 |  end
20 |  if should_forward then
21 |  |  Forward all stored packets up to sequence
   |  |  min_{1≤r≤n} last_r^seen;
22 |  |  Update last^fwd;
23 |  end
24 end
```

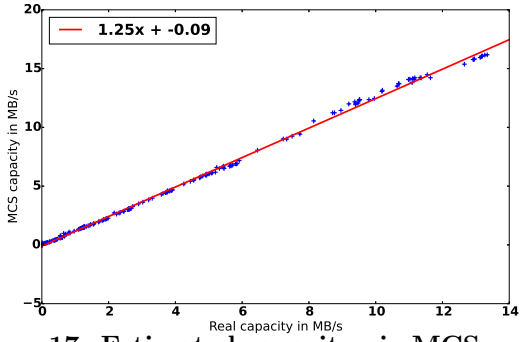

Figure 17: Estimated capacity via MCS vs available throughput measured.

iments on one link. We run four experiments back-to-back, using only one of the interfaces (WiFi, PLC) in two, using both interfaces and our load-balancing algorithm (Hybrid) in one, and using both interfaces and a round-robin scheduler for the packets (Round-robin) in the last one. We observe that by using simple load-balancing and reordering algorithms, and our capacity-estimation technique, we can achieve a throughput that is very close to the sum of the capacities of both mediums. In contrast, the throughput of a round-robin scheduler, which has no information on capacity, is limited to twice the minimum capacity of the two mediums (i.e., WiFi in this example), because it assigns the same number of packets to each medium and the slowest medium becomes a bottleneck. To evaluate our algorithm across our testbed, we also compare the completion times of a 600Mbyte file download using (i) only WiFi, and (ii) both mediums[8], observing in the same figure, a drastic decrease in completion times when using both mediums.

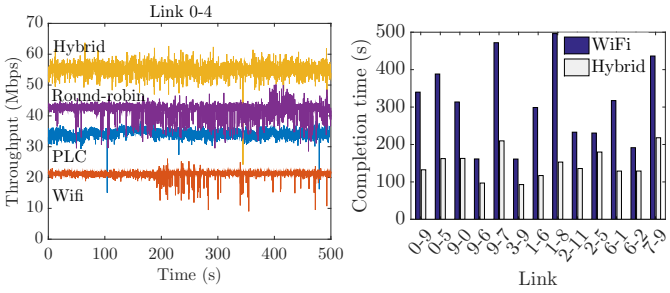

Figure 18: Performance boost by using hybrid Wifi/PLC, and our load-balancing and capacity-estimation techniques.

Our tests validate our capacity estimation methods. They also show that, to exploit each medium to the fullest extent, accurate link-quality metrics are required. However, an open question to be answered is; How should the link metrics be updated to take into account delay or contention? In the next section we investigate another link metric, i.e., the expected number of retransmissions, and the performance of link metrics with respect to background traffic.

[8]Contrary to WiFi, PLC uses queues that are non-blocking: the transport layer is not stopped from sending packets when the MAC queues are full. For these experiments, we omit PLC tests as dropped packets yield an unfair comparison.

8. RETRANSMITTING IN PLC CHANNELS

Capacity is a good metric for link quality. However, it does not take into account interference, which is very important for selecting links with high available bandwidth. Moreover, another metric could be useful for delay sensitive applications that do not saturate the medium but have low delay requirements. Delay is affected by retransmissions either due to bursty errors or to contention, and metrics, such as PB_{err} introduced in Section 5 (or *packet errors* [2]), are related to retransmissions. We explore the mechanism of retransmissions in PLC networks. We first study another link metric, which is the expected transmission count (ETX). Numerous works, e.g., [7], [8], study this metric (or its variations) in WiFi networks by sending broadcast probes. We examine how ETX performs in PLC and the relationship between broadcast and unicast probing.

After studying retransmissions due to errors, we evaluate the sensitivity of link metrics to background traffic. Link metrics in hybrid networks should estimate the amount of background traffic, or be insensitive to background traffic. Thus, a critical challenge for hybrid networks is to design link metrics achieving one of the aforementioned properties.

To this end, we represent the IEEE 1901 MAC layer in Figure 19. The Ethernet packets are organized in PBs (with size 512 bytes). Then, the PBs are forwarded to a queue and based on the BLE of the current tone-map slot s BLE_s, they are aggregated into a PLC frame. The frame duration is determined by BLE_s, the maximum frame duration (specified by [6]), and an aggregation timer that fires every few hundreds of ms (as concluded from our measurements in this section) after the arrival of the first PB[9]. The PLC frame is transmitted by a CSMA/CA protocol explained in [18]. The receiver decodes the frame and transmits a SACK that informs the transmitter which PBs were received with errors. We observe that the full retransmission and aggregation process, and, as a result, the MAC and PHY layers, can be evaluated or simulated by using only two metrics: PB_{err} and BLE_s.

8.1 Retransmission Due to Errors

We first explore how ETX would perform in PLC by sending broadcast packets. Because broadcast packets in PLC are transmitted with the most robust modulation

[9]Note that the frame duration is a multiple of the symbol duration, and that padding is used to fill these symbols.

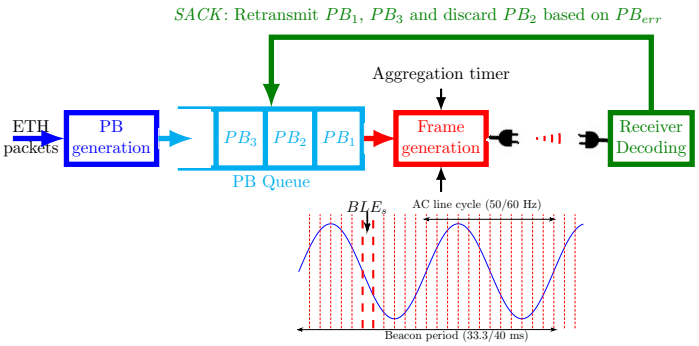

Figure 19: The PLC MAC layer

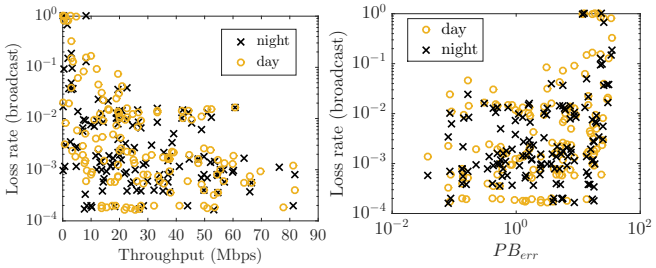

Figure 20: Loss rate for broadcast packets vs link throughput and PB_{err} for all station pairs.

and are acknowledged by some proxy station [6], we expect that this method yields very low loss rates.

For the purpose of this study, we set each station in turn to broadcast 1500 byte probe-packets (1 every 100ms) for 500 sec. The rest of the stations count the missed packets by using an identification in our packet header. We repeat the test for all stations of the testbed during night and working hours (day). Figure 20 shows the loss rate from these tests for all station pairs, as a function of throughput and PB_{err}. Each pair is represented with its link throughput (respectively, PB_{err}) during the night experiment.

Conclusion: Loss rate of broadcast packets in PLC is a very noisy metric for the following reasons:

(*i*) A wide range of links with diverse qualities (throughput from 10 to 80 Mbps) have very low loss rates ($\sim 10^{-4}$), and some links even have 0 loss rates. By observing high loss rates, e.g., larger than 10^{-1}, ETX can classify bad links in PLC, but nothing can be conjectured for link quality from low loss rates.

(*ii*) There is no obvious difference between experiments during the day, when the channel is worse, and night. A few bad links have much worse loss rates during the day, but at the same time, also a few average links yield extremely lower loss rates.

(*iii*) As PLC adapts the modulation scheme to channel conditions when data is transmitted, broadcast packets – sent at most robust modulation scheme – cannot reflect the real link quality. Moreover, given the low loss rates of a wide range of links, ETX appears to be 0 at short-time scales, which provides no or misleading information on link quality.

Due to the above observations, we further explore the mechanism of retransmissions with respect to link quality with unicast traffic. We now delve into the retransmissions of PBs by sending unicast, low data-rate traffic, i.e., 150Kbps, and by capturing the PLC frame headers. Under this scenario, an Ethernet packet of 1500 bytes is sent approximately every 75ms. The test has a duration of 5 min per link. As we have discussed above, broadcast packets might be missed by some stations when channel conditions are bad, because they are not retransmitted as soon as a proxy station acknowledges them. In contrast, unicast packets are being retransmitted until the receiver acknowledges them, hence are always received. For this reason, we look at the frame header SoF to study retransmissions. Because there is no indication on whether the frame is retransmitted in the PLC SoF, we employ the arrival time-stamp of the frame to characterize it as a retransmission or new transmission (if the frame arrives within an interval of less than 10ms compared to the previous frame, then it is a retransmission). We also measure PB_{err} every 500 ms.

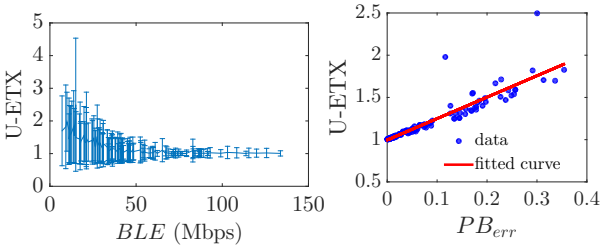

Figure 21: U-ETX vs BLE and U-ETX vs PB_{err}.

We conduct the experiment described above for all the links of our testbed. We compute the unicast ETX (U-ETX) for all the links of the testbed. We count the total number of retransmissions for a packet of 1500 bytes, which produces 3 PBs. A retransmission occurs if at least one of these PBs is received with errors. Figure 21 presents U-ETX as a function of average BLE (with links sorted in increasing BLE order) and PB_{err}. U-ETX is measured by averaging the number of PLC retransmissions for all packets transmitted during the experiment. We also plot error-bars with the standard deviation of the transmission count. It turns out that link quality is negatively correlated with link variability, a conclusion made also when exploring BLE in Section 6.2. The higher the U-ETX is, the higher the standard deviation of transmission count is. Links with high BLE are very likely to guarantee low delays, as U-ETX does not vary a lot. U-ETX and the averaged PB_{err} are highly correlated, with almost a linear relationship.

8.2 Retransmission Under Saturated Scenarios

Towards the goal of characterizing the performance of contending links, we first run saturated tests. We examine whether there exists rate anomaly, which occurs when the throughput of all contending links is bounded by the slowest transmission rate [10]. We discovered that rate anomaly does not exist in PLC networks. Figure 25 presents throughput measurements of two and three links contending simultaneously, along with the corresponding metrics when the links are isolated. The reason is that IEEE 1901 restricts the maximum frame size in μs and not in bytes. As a result, the frame length might be the same for saturated low-throughput and high-throughput links. The difference between the two links is the bit-loading (BLE) of the frames. In addition to this result, we observe that the best link might dominate and gain higher throughput than half its available throughput (when isolated).

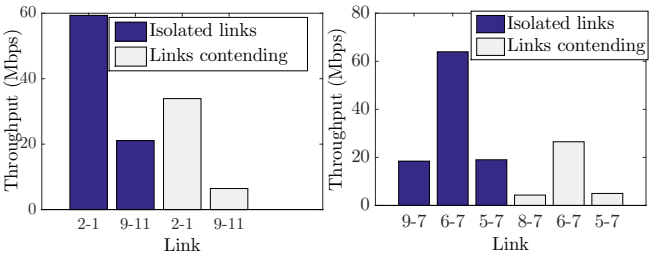

Figure 22: Performance of 2 contending links (left) and 3 contending links (right).

To understand and confirm the causes of the two observations above, we look at the average burst durations of saturated links. Figure 23 shows the average duration of the burst of frames vs average throughput for all links. We observe that this average is almost independent of the link quality except from bad links. This is one of the reasons why the best link dominates during contention of multiple flows: when a collision occurs, the link with the largest frame duration still manages to recover some of the PBs contrary to the other links that have all PBs collided. Also, since each symbol of the best link contains much more PBs than the bad links, even if the difference between their frame durations is a couple of symbols, the good link can recover tens of PBs. Another reason is a "capture effect" of the station with the highest SNR, which can decode a few PBs even if these overlap with those of a link with lowest SNR.

To explore the performance of our two link metrics BLE and PB_{err} under contention, we run the following three-phase experiment. We first activate one link, then a second one, and then deactivate the first link. Each phase has a duration of 4 min. Figure 24 presents such an experiment with four metrics: BLE, burst duration, UDP throughput, and PB_{err}.

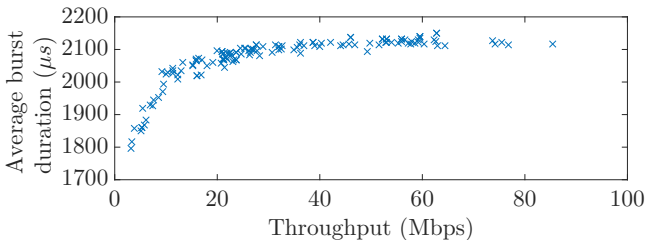

Figure 23: Burst duration vs throughput

Conclusion: We observe the following:

(i) BLE is not affected when multiple, saturated flows are contending.

(ii) Bad saturated links have a slightly smaller frame duration compared to good links.

(iii) PB_{err} has higher variation when there is high contention in the network. Thus, by using PB_{err} (both its mean and its variance), we could identify the presence of interference in the network. IEEE 1905 requires metrics to estimate interference [2]. Analyzing packet error metrics is significantly more challenging than capacity estimation. We leave this study for future work.

(iv) By using BLE, burst duration and the model of Figure 19, the data-rate demand can be computed.

Note that this data can be used to estimate throughput when multiple stations are contending using analytical models similar to [18]. Furthermore, it can be used for simulating the MAC layer of PLC.

8.3 Retransmission Due to Contention

To explore the sensitivity of link metrics to background traffic and to examine how interference can be considered in link metrics, we now experiment with two contending flows. We set a link to send unicast traffic at 150Kbps as in the previous subsection, emulating probe packets. After 200 seconds, we activate a second link sending "background" traffic at various rates. We measure both BLE and PB_{err}. In these experiments, we observe that BLE is insensitive to low data-rate background traffic for all pairs of links. However, BLE appears to be affected by high data-rate background traffic on a few pair of links. So far, we have not found any correlation between these pairs of links. We explain this phenomenon with the "capture effect", where the best link decodes a few PBs even during a collision due to very good channel conditions, yielding high PB_{err}. In this case, the channel-estimation algorithm cannot distinguish between errors due to PHY layer and errors due to collisions, hence it decreases BLE. Figure 25 presents two representative examples of link pairs for which BLE is sensitive and nonsensitive to high data-rate background traffic. Observe that PB_{err} explodes in link 6-11, which is sensitive to background traffic.

To tackle the sensitivity of BLE to high data-rate background traffic, we take advantage of the frame ag-

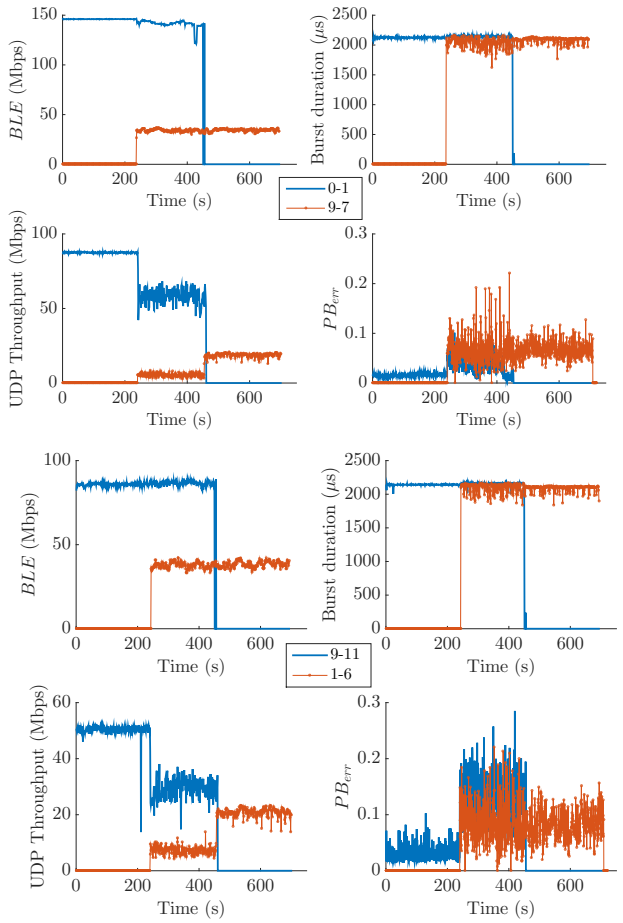

Figure 24: Performance of 2 contending links

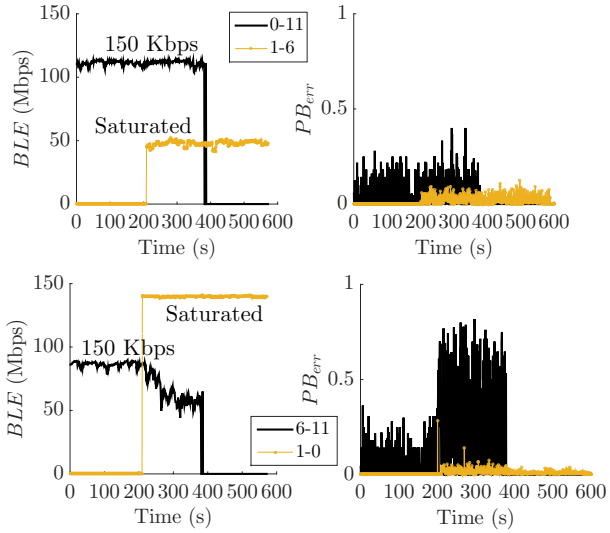

Figure 25: Link metrics of 2 sets of contending links with low data-rate and saturated traffic.

gregation procedure of the MAC layer, described in Section 2.2. We observe that transmitting a few PBs per 75ms (150Kbps ratw) yields a sensitivity of metrics to background traffic. However, when two saturated flows are activated, we never notice an effect on BLE (see Section 8.3). Due to frame aggregation, packets from different saturated flows have approximately the same frame length (i.e., maximum) and when they collide, the channel estimation algorithm works more efficiently than when short probe-packets collide with long ones. To emulate the long frame lengths of saturated traffic, we send bursts of 20 packets such that the traffic rate per second (i.e., the overhead) is kept the same (150Kbps). In Figure 26, we show another link for which BLE is sensitive to background traffic, and the results of our solution. By sending bursts of probe packets, BLE is no more affected by background traffic. This shows that by exploiting the MAC layer of PLC, we can tackle the sensitivity of link metrics to background traffic.

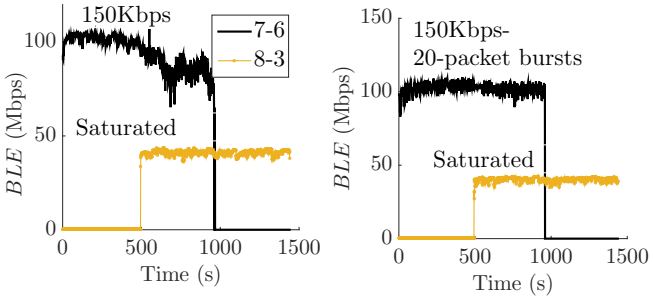

Figure 26: Tackling the link-metric sensitivity to background traffic by sending bursts of probes.

Conclusion: We have studied the mechanism of retransmissions in PLC. Although broadcast probe-packets yield significantly less overhead in link-quality estimation, they do not provide accurate estimations. In contrast, unicast probe-packets reflect the real link quality, but by producing more overhead. We observe that PB_{err} can be used to estimate U-ETX and to indicate interference in PLC. However, estimating the amount of interference is challenging and should be further investigated. We leave this extension for future work.

9. RELATED WORK

A large body of work (e.g., [9], [15]) focuses on channel modeling, and very little work, such as [13], investigates the PLC performance from an end-user perspective. The authors in [13] explore the performance of HomePlug AV when household devices operate in the network. They observe that switching the appliances affects significantly the performance and introduces asymmetry, and that different appliances create diverse noise levels. However, this study does not introduce any link metrics or temporal variation analysis.

Many previous experimental works focus on the PLC MAC layer under single contention domain scenarios and ideal channel conditions in order to model and evaluate MAC characteristics. To achieve these conditions, the

stations are plugged to the same power-strip and are isolated from the power-grid. Zarikoff and Malone [19] give the guidelines for a PLC testbed construction and perform measurements with both UDP and TCP traffic, and multiple contending flows. We also use a testbed setup of 7 stations to evaluate the performance of the HomePlug AV CSMA/CA process in [18].

A few works focus on comparing the wireless and PLC performance [12, 17]. [12] investigates older specifications of PLC and WiFi, i.e., HomePlug 1.0 and 802.11 a/b, respectively. The authors provide testbed measurements from 20 houses for metrics such as coverage, throughput, and connectivity. [17] introduces a comparison between hybrid PLC/WiFi networks and single-technology networks. The authors find that hybrid networks contribute to increase coverage in home networks; they also argue that using alternating technologies for multi-hop routes yields good performance. However, they do not study link metrics that can be used to optimize routing in such networks. To the best of our knowledge, there is no study on PLC link metrics.

10. CONCLUSION

We have shown that PLC can yield significant performance gains when combined with WiFi networks. Yet, there were open questions on how to exploit to the fullest the two mediums and PLC has received far too little attention from the research community; we introduce an experimental framework and investigate the performance of PLC. To this end, we explore its spatial and temporal variation, delving into the diverse time-scales of PLC channel variability.

We have studied PLC link metrics and their variation with respect to space, time, and background traffic. Similar metrics have been long pursued by the research community for WiFi and have been required by the recent standardization of hybrid networks. We have uncovered accurate link metrics and have given guidelines on metrics estimation in hybrid implementations. We have observed that there is a high correlation between link quality and its variability, a finding that has a direct impact on probing overhead and accurate estimations.

11. REFERENCES

[1] HomePlug Alliance. `http://www.homeplug.org/`.

[2] IEEE 1905.1-2013 Std for a Convergent Digital Home Network for Heterogeneous Technologies.

[3] IEEE 802.11n-2009-Amendment 5: Enhancements for Higher Throughput.

[4] OpenWrt. `http://https://openwrt.org/`.

[5] Qualcomm Atheros Open Powerline Toolkit. `https://github.com/qca/open-plc-utils`.

[6] IEEE Standard for Broadband over Power Line Networks: Medium Access Control and Physical Layer Specifications. *IEEE Std 1901-2010*, 2010.

[7] S. M. Das, H. Pucha, K. Papagiannaki, and Y. C. Hu. Studying wireless routing link metric dynamics. In *Proc. of the 7th ACM Conf. on Internet measurement*, pages 327–332, 2007.

[8] R. Draves, J. Padhye, and B. Zill. Routing in multi-radio, multi-hop wireless mesh networks. In *Proc. of ACM MobiCom, 2004*, pages 114–128.

[9] S. Guzelgoz, H. B. Celebi, T. Guzel, H. Arslan, and M. K. Mıhçak. Time frequency analysis of noise generated by electrical loads in PLC. In *17th International Conf. on Telecommunications (ICT)*, pages 864–871. IEEE, 2010.

[10] M. Heusse, F. Rousseau, G. Berger-Sabbatel, and A. Duda. Performance anomaly of 802.11 b. In *IEEE INFOCOM 2003*, pages 836–843.

[11] E. Kohler, R. Morris, B. Chen, J. Jannotti, and M. F. Kaashoek. The Click modular router. *ACM TOCS*, 18(3):263–297, 2000.

[12] Y.-J. Lin, H. A. Latchman, R. E. Newman, and S. Katar. A comparative performance study of wireless and power line networks. *Communications Magazine, IEEE*, 41(4):54–63, 2003.

[13] R. Murty, J. Padhye, R. Chandra, A. R. Chowdhury, and M. Welsh. Characterizing the end-to-end performance of indoor powerline networks. Technical report, Harvard University Microsoft Research Technical Report, 2008.

[14] K. Papagiannaki, M. D. Yarvis, and W. S. Conner. Experimental characterization of home wireless networks and design implications. In *IEEE INFOCOM 2006*.

[15] S. Sancha, F. Canete, L. Diez, and J. Entrambasaguas. A channel simulator for indoor power-line communications. In *IEEE International Symposium on Power Line Communications and Its Applications (ISPLC)*, pages 104–109, 2007.

[16] R. K. Sheshadri and D. Koutsonikolas. Comparison of routing metrics in 802.11n wireless mesh networks. In *IEEE INFOCOM 2013*, pages 1869–1877.

[17] P. Tinnakornsrisuphap, P. Purkayastha, and B. Mohanty. Coverage and capacity analysis of hybrid home networks. In *IEEE International Conf. on Computing, Networking and Communications (ICNC), 2014*, pages 117–123.

[18] C. Vlachou, A. Banchs, J. Herzen, and P. Thiran. On the MAC for Power-Line Communications: Modeling Assumptions and Performance Tradeoffs. In *IEEE International Conf. on Network Protocols (ICNP)*, 2014.

[19] B. Zarikoff and D. Malone. Construction of a PLC testbed for network and transport layer experiments. In *Proc. of IEEE ISPLC 2011*, pages 135–140.