WOLT: Auto-Configuration of Integrated Enterprise PLC-WiFi Networks

Hisham A. Alhulayyil*†, Kittipat Apicharttrisorn*, Jiasi Chen*
Karthikeyan Sundaresan‡, Samet Oymak*, Srikanth V. Krishnamurthy*
* *University of California Riverside, USA*, {halhu001, kapic001}@ucr.edu,{jiasi, krish}@cs.ucr.edu, oymak@ece.ucr.edu
†*Al Imam Mohammad Ibn Saud Islamic University, KSA*, halhulayyil@imamu.edu.sa
‡ *NEC Laboratories America, Inc*, karthiks@nec-labs.com

***Abstract*—Power Line Communication (PLC) based WiFi extenders can improve WiFi coverage in homes and enterprises. Unlike in traditional WiFi networks which use an underlying high data rate Ethernet backhaul, a PLC backhaul may not support high data rates. Specifically, our measurements show that arbitrarily affiliating users to PLC-WiFi extenders or based on their WiFi channel qualities alone may lead to poor network performance due to the differences in PLC link capacities. Thus, in this paper we build a framework, WOLT, to solve the problem of assigning users to the appropriate PLC-WiFi extenders to increase the aggregate network throughput in an enterprise setting, where one may expect a relatively large number of power outlets. WOLT accounts for both the qualities of the two concatenated links viz., the PLC and WiFi links. It hinges on estimating the best capacity offered by the PLC links, and accounting for these while assigning users. It incorporates a polynomial-time algorithm that assigns only a subset of the users to maximize the aggregate throughput on the PLC links, and then assigns the remaining users such that the degradation in the aggregate throughput is minimized. WOLT is evaluated through simulations and real testbed experiments with commodity PLC-WiFi extenders, and improves aggregate throughput by more than 2.5x compared to a greedy user association baseline.**

I. Introduction

A quick and easy way of improving indoor WiFi coverage (in homes and enterprises) is via the use of WiFi-capable PLC (Power Line Communication) extenders. These extenders are typically considered plug-and-play devices which, when plugged into a power line outlet, connect to a master router or gateway that is in turn connected to the Internet. Each of these extenders plays the role of an additional wireless access point (AP), to which client devices can attach themselves. An example setup is shown in Fig. 1. In the enterprise context for example, users in office spaces could potentially plug in extenders. Using such extenders is expected to boost the WiFi coverage [1] by improving the WiFi signal quality in the region of interest, especially in areas where previously (in the absence of these extenders) the coverage was poor. In fact, today several commodity WiFi-capable PLC extenders from various vendors are available on the market (e.g. TP-Link, Netgear, Zyxel, Linksys, and Amped) [1].

While the benefit of using PLC based WiFi extenders is potentially an improvement in throughput, a key observation to be made is that the PLC backhaul very much differs from an Ethernet backhaul. Almost all prior papers that try to manage WiFi networks (*e.g.*, [2]–[4]) assume that the link qualities and contention on the backhaul can be safely ignored, focusing on simply the wireless access portion of the end-to-end communications. However, this assumption does not hold true with a PLC backhaul as implicitly shown in [5], [6]. Specifically, first, unlike with Ethernet, PLC links often are much more constrained in terms of their capacity, which might in fact manifest as the bottleneck in a concatenated PLC-wireless link. In other words, the capacity of the WiFi link between a client and a PLC extender might exceed the capacity of the backhaul link between the extender and an Internet-connected master router. Second, a PLC link's capacity is shared not only between clients that attach to a specific extender, but also between extenders; this in turn could cause a client's achievable throughput on the PLC backhaul to be lower than that of its WiFi link, even if the PLC backhaul was of relatively high capacity.

Fig. 1: Powerline communications extend WiFi coverage through existing interior power lines in a home/enterprise.

Because of these properties of the PLC backhaul, if clients (aka users in this work) were to naïvely associate with the closest extender or the extender that offered the best received signal strength (RSS), they could (a) end up with low throughputs individually, and (b) also cause the aggregate network throughput to suffer (we show this later). Therefore, the quality of the associated backhaul PLC links and the contention on those links will need to be considered when associating users with the available extenders. In this paper, our objective then is to solve the problem of configuring the network in terms of assigning users to the available extenders towards maximizing the achievable network throughput.

Towards achieving our objective, we first perform several experiments using commodity PLC TP Link TL-WPA8630 PLC extenders, to not only showcase the issues alluded to above with naïve association, but to also understand how the capacity is shared using the 1901 MAC [7] on the PLC medium (*e.g.*, time fair or throughput fair?). Based on our experiments, we then formulate the problem of maximizing the aggregate end-to-end network throughputs over the possible associations of users with the different PLC extenders with which they can connect. We show that the problem is NP-hard. Based on this, we first then solve an unconstrained version of the problem wherein we automatically discover a subset of users (say $U_1 \subset U$, where U is the set of users) that can be assigned to achieve the maximum possible throughput; in other words, a subset of the users U_2 ($U_2 = U \setminus U_1$) whose associations cause the overall network throughput to degrade are ignored in this step. Subsequently, we assign the remaining U_2 users such that the degradation in the previously achieved throughput is minimized. Our algorithms are incorporated into a framework we call WOLT (the term replaces the letter V in the word Volt related to PLC with W to show the dependency on WiFi). We show via both real experiments on our PLC testbed and high-fidelity simulations that WOLT significantly improves throughput compared to both a naïve approach where users affiliate with the PLC extender that provides the best RSSI, and a greedy centralized algorithm that assigns each incoming user to the extender that yields the maximum aggregate throughput (other users are not reassigned). We also show that the reassignment load of WOLT incurs relatively minor overhead penalties.

In brief, a summary of our contributions in this paper are:

- We conduct extensive experiments on a WiFi network with a PLC backhaul to understand the interaction between the WiFi links and the PLC backhaul, and how this interaction affects the aggregate throughput of the whole network.
- We leverage the insights obtained from the measurement experiments to design WOLT for assocating users with the available WiFi-capable PLC extenders. WOLT runs in polynomial time and solves the user association problem, based on a relaxed version of the problem with guaranteed integer solutions.
- We evaluate WOLT with real downlink TCP traffic and we show that WOLT outperforms RSSI-based and a centralized online algorithm that performs greedy assignment, with WOLT achieving an aggregate throughput increase of up to 2.5x. To examine scalability, we also perform high-fidelity simulations (validated against the real world system at small scale) and show that WOLT performs well in enterprise setting with up to 15 extenders and 124 clients.

II. PLC Background in Brief

Access control in PLC networks is governed by the IEEE 1901 standard and can operate in two modes: CSMA or TDMA [8]. In brief, CSMA/CA mode arbitrates the communication channel among transmitting nodes within the same contention domain. As discussed in [7], the 1901 CSMA/CA is similar to what is used in 802.11, with some differences in terms of performance, complexity and fairness [9]. PLC also supports QoS classes by providing a TDMA-based medium sharing functionality. In TDMA mode, the PLC backhaul will be time-shared between clients [10]. TDMA and CSMA modes are supported by most major PLC devices on the market today that follow the IEEE 1901 standard, including major vendors such as Netgear, TP-Link, and TRENDnet. Today, PLC extenders support up to 2024 Mbits/sec [11]. This makes PLC suitable even for bandwidth intensive applications such as video streaming [12].

Modern PLC extenders are capable of providing WiFi connectivity to associated users. The WiFi access between extenders and other APs, is governed by the 802.11 standard. The WiFi medium is shared in a throughput-fair manner (as is the case with 802.11 [13]–[16]). This means that all users that are connected to a PLC extender will receive the same long-term average WiFi throughput (share of throughput on the wireless link to the extender). PLC extenders share the wireless medium if they operate on the same wireless frequency.

III. Measurement Study

In this section, we perform a measurement study to showcase the interactions between the PLC and the WiFi domains. These measurements form the basis for the models and algorithms that in turn are the underpinnings of WOLT. In addition, we also showcase via case studies as to how the interactions propagate from the WiFi domain to the PLC backhaul, or vice versa. Note that in the remainder of this paper, we make a distinction between the "throughput" and the "rate". By throughput, we mean the achieved bit rate on a WiFi/PLC link, which depends on the other users/extenders sharing that link. By rate, we mean the PHY bit rate of a WiFi user or a PLC extender, which depends on the current channel conditions and the selected modulation and coding scheme.

A. Medium sharing in the PLC and WiFi domains

Sharing in the WiFi domain: First, we only consider the WiFi part (although this is well studied). *We find that the well-studied performance anomaly with 802.11 [15] surprisingly still persists when we use currently available commodity PLC extenders.* Specifically, we connect one laptop to a TP Link TL-WPA8630 PLC extender (released in 2017) with an Ethernet cable, and two additional laptops to the extender via WiFi.

The Ethernet-connected laptop runs an `iperf3` [17] server and the other two laptops are the WiFi clients; all other extenders are unplugged i.e., both the clients are connected to the target extender. We transmit saturated downlink TCP traffic to the two clients simultaneously, and plot the throughputs of the two clients when placed at different locations in Fig. 2a. Initially, we position the two clients at the same location i.e., at equal distance from the extender (*location 1*). The WiFi channel qualities (rates) for both clients are similar (similar distance and RSSI), and the throughputs obtained by the clients are similar. Subsequently, we move one of the clients (User 2) to a location further away from the extender (*location 2*) to degrade its WiFi channel quality (rate), while keeping the other client (User 1) stationary at location 1. Not only does the

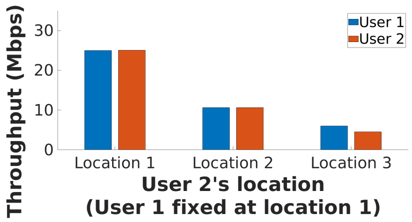

(a) WiFi only: Throughput-fair medium sharing behavior of WiFi clients.

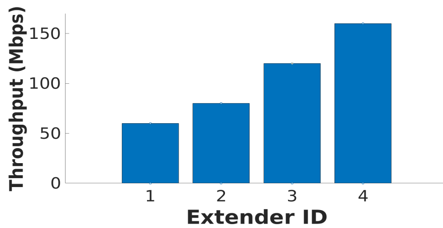

(b) PLC only: Different PLC extenders yield in different PLC throughputs.

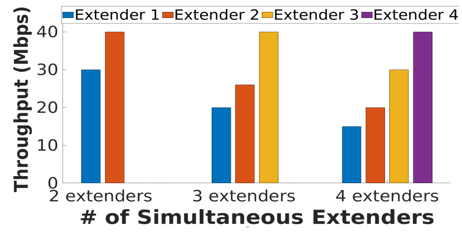

(c) PLC only: Time-fair medium sharing between active PLC extenders.

Fig. 2: Medium sharing in the PLC and WiFi domains.

further client see a throughput degradation, but the stationary client's throughput decreases as well, in accordance with the reported performance anomaly [15]. The further we move User 2 (to *location 3*), the greater the throughput loss experienced by both clients. This demonstrates that the sharing in the WiFi domain when PLC extenders are used is "throughput fair," and consistent with studies such as [15].

PLC backhaul sharing: Next, our goal is to examine the medium sharing on the PLC backhaul in isolation (*i.e.*, without WiFi). *We find that the the default operation of the PLC extenders adheres to time-fair medium sharing, with each extender receiving throughput commensurate with its PLC link quality (rate).* Specifically, we connect one laptop to a Netgear R7000-100NAR Nighthawk master router via Ethernet. The master router interfaces with the PLC backhaul via a PLC central unit. Four PLC extenders are plugged into different power outlets in our lab, with one client laptop connected to each extender through Ethernet (note that the Ethernet capacity is very high at 1 Gbps so any throughput degradation is caused by the PLC). A laptop connected to the master router serves as an `iperf3` server and transmits saturated downlink TCP traffic to the four client laptops. We experimented with four PLC links of varying link qualities (rates), with the maximum achievable throughput of each PLC extender in isolation ranging from 60-160 Mbps, as shown in Fig. 2b. To see the impact of medium sharing on PLC backhaul throughput, we activated multiple extenders simultaneously, and plot the results in Fig. 2c. With two extenders actively receiving `iperf3` traffic, we observe that each PLC link now delivers half of what it could in isolation (with higher throughput for the extender with better rate); with three extenders active, each PLC link delivered one third of what it could in isolation, and with four extenders active, each PLC link delivered one quarter of what it could in isolation. This suggests that the PLC backhaul is time-shared, and other researchers have made similar observations [18]. We note that the time-fair-like behavior we observed is the default behavior of the off-the-shelf TP-Link PLC equipment that are popular with consumers, and we have observed similar default behavior with other brands of PLC equipment as well.

WiFi with PLC backhaul sharing: As discussed above, with our experimental apparatus, the WiFi and PLC parts seem to adhere to different medium sharing schemes. Next, we consider two types of PLC-WiFi concatenated links: (a) a concatenated link wherein the WiFi link segment is of better quality (yields higher throughputs) than the PLC link segment, and (b) a concatenated link where the PLC link segment is of better quality than the WiFi link segment.

The total throughput of a group of users connected to the same extender follows the WiFi-only throughput-fair sharing behavior discussed before and can be expressed as:

$$\mathcal{T}_j^{\text{WiFi}} = \sum_{i \in N_j} \frac{1}{\sum_{i' \in N_j} \frac{1}{r_{i'j}}} \tag{1}$$

where $\mathcal{T}_j^{\text{WiFi}}$ is the aggregate WiFi throughput at extender j, i is the user index, N_j is the set of users connected to extender j, and r_{ij} is the WiFi rate between User i and extender j (similar to [2], [16]). Note here that the WiFi throughput is taken to be the sum throughput achieved across all users connected to the same extender (similar to the PLC link throughput, which is the total throughput across the clients sharing that link).

The PLC link segment throughput adheres to the time-fair sharing discussed above and can be expressed as:

$$\mathcal{T}_j^{\text{PLC}} = \frac{c_j}{A} \tag{2}$$

where $\mathcal{T}_j^{\text{PLC}}$ is the aggregate PLC throughput at extender j, c_j is the PLC rate and A is the number of active extenders.

Given the above, the achievable throughput that an extender (say j) can obtain is the minimum of the throughputs on the two concatenated link segments, *i.e.*, it is given by $\min(\mathcal{T}_j^{\text{WiFi}}, \mathcal{T}_j^{\text{PLC}})$. Furthermore, in some cases, a PLC link segment can yield a higher throughput than its aggregate WiFi links, while in other cases it might not. If the WiFi link segment that is part of a concatenated link yields lower throughput than what the associated PLC link can support (*i.e.*, $\mathcal{T}_j^{\text{WiFi}} < \mathcal{T}_j^{\text{PLC}}$), then that PLC link will have unused capacity. This leftover capacity (time share) can be then exploited by other extenders which might have a higher demand, because their connected WiFi users have good channel conditions and thus demand more traffic than their extender's time-fair share of the PLC medium could have supported. Experimental results from such a scenario are provided in § III-B (when discussing greedy association).

B. Showcasing the need for informed user association

Next, our goal is to showcase a simple exemplar case study that demonstrates that blindly connecting users to extenders

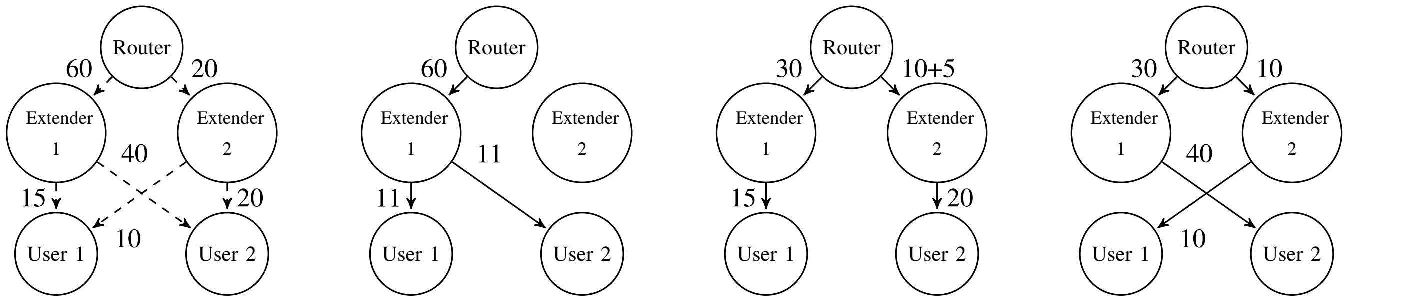

(a) Maximum achievable throughput (rate) of each link in isolation.

(b) RSSI-based assignment. Total throughput = 11+11 = 22 Mbps.

(c) Greedy assignment. Total throughput = 15+15 = 30 Mbps.

(d) Optimal assignment. Total throughput = 10+30 = 40 Mbps.

Fig. 3: Experiments in a scenario observed on our testbed with user association policies.

that offer them the best RSSI can lead to undesired effects in terms of throughput performance. We then consider what happens when users that connect to the network sequentially online (one after the other), are associated so as to maximize their own throughputs greedily. Finally, we do a brute force search that shows the throughputs with optimal association. In a nutshell, we find that the first two baseline association policies drastically underperform the last. The different user association strategies in this case study are shown in Figure 3. Fig. 3a shows our experimental setup with two PLC extenders and two WiFi clients, with labels indicating PLC or WiFi rates in the absence of contention.

Strongest RSSI-based association: In Fig. 3b, we depict this method of user association with an AP; this method has been commonly considered [2] previously, when the network had an associated Ethernet backhaul. The two users shown contend on extender 1's PLC link while extender 2's PLC link is interference-free. Because of the association policy, there is WiFi contention and this causes the aggregate throughput of this assignment to be 22 Mbit/sec (11 Mbps each, according to the throughput fair sharing on the WiFi). As shown later, the optimal strategy yields a much higher throughput.

Greedy association: A greedy association policy is shown in Fig. 3c. User 1 arrives and chooses extender 1 since this maximizes its own throughput. Next, User 2 arrives and chooses extender 2 so as to maximize its own end-to-end throughput (given that User 1 is fixed). Note that even though User 2 has a worse WiFi channel quality to extender 2 compared to extender 1, and despite the fact that extender 2 has poor PLC link capacity compared to extender 1, User 2 still prefers extender 2 because its end-to-end throughput is higher than if it had connected to extender 1 (which would result in the same scenario as RSSI-based association above). Extender 2 time-shares the medium equally with extender 1, causing extender 2's PLC time-shared capacity (10 Mbps) to become the bottleneck of User 2's end-to-end throughput. However, we notice that because extender 1 does not use all of its capacity of 30 Mbps (because User 1's WiFi capacity is only 15 Mbps), half of extender 1's leftover time (*i.e.*, one quarter of the total time) is re-allocated to extender 2, causing User 2's end-to-end throughput to increase to 15 Mbps.

Optimal user association: In Fig. 3d, we show the optimal extender associations for the two users. User 1 connects to extender 2 and receives 10 Mbps. User 2 connects to extender 1 and receives 30 Mbps, despite its WiFi capacity being 40 Mbps; this is because its end-to-end throughput is bottlenecked by the backhaul capacity of the extender 1.

IV. Problem Statement & Solutions

Our overarching problem is to maximize the total end-to-end-throughput of all users. To do this, in § IV-A we develop a system model based on the take aways from § III, and formulate the problem of maximizing the total throughput. We show that this problem is NP-hard; then, we propose solutions in § IV-B based on certain intuitive properties of the model.

A. Problem Statement and Hardness

In our network model, each link is a concatenation of PLC backhaul and WiFi wireless links. We seek to maximize the total end-to-end network throughput. The PLC and WiFi have different contention mechanisms at the MAC layer as discussed in § III, resulting in different throughput sharing functions. The concatenation of the PLC and WiFi links make the problem different and more challenging than the standard single-hop user association problem, which has been well-studied (*e.g.*, [19]–[21]). The model of the scenario is a single contention domain across the PLC extenders. This is found to be the case with the current standards regardless of whether the deployment is in homes or enterprises [10].

We formulate the throughput maximization problem in Problem 1 below (notation in Table I):

Problem 1. *PLC-WiFi User Assignment*

$$max_{x_{ij}} \quad \sum_{j=1}^{|A|} \min\left(\mathcal{T}_j^{WiFi}, \mathcal{T}_j^{PLC}\right) \tag{3}$$

$$s.t. \quad \mathcal{T}_j^{PLC} = \frac{c_j}{A}, \quad \forall\, j \in A \tag{4}$$

$$\mathcal{T}_j^{WiFi} = \sum_{i=1}^{|U|} t_{ij}, \quad \forall\, j \in A \tag{5}$$

$$t_{ij} = \left(\frac{1}{\sum_{i' \in N_j} \frac{1}{r_{i'j}}}\right) x_{ij}, \quad \forall\, j \in A, \forall\, i \in U \tag{6}$$

$$\sum_{j=1}^{|A|} x_{ij} = 1, \quad \forall\, i \in U \tag{7}$$

$$\sum_{i=1}^{|U|} x_{ij} \le B_j, \quad \forall\, j \in A \tag{8}$$

$$N_j = \{i : x_{ij} > 0\}, \quad j \in A \tag{9}$$

$$x_{ij} \in \{0, 1\}, \quad \forall\, i \in U, \forall\, j \in A \tag{10}$$

The objective (3) is to maximize the total end-to-end throughput across all extenders and users (*i.e.*, the minimum of the throughputs achieved on the PLC and WiFi links). Constraint (4) specifies the throughput of the PLC link connecting the master router to PLC-WiFi extender j, based on time-fair sharing of the PLC backhaul. Constraint (5) specifies the WiFi throughput at extender j, summed across all users. Constraint (6) specifies the throughput of user i connected to extender j, based on throughput-fair sharing. Constraint (7) postulates that each user must be connected to one extender. Constraint (8) postulates that each extender j should have no more than B_j connected users (this constraint will be relaxed later). Constraint (9) defines the set of users N_j connected to extender j as those users i who have a non-zero assignment to that extender. Finally, constraint (10) says that x_{ij} is a binary decision of the extender to which a user is assigned. Our model assumes that the users have saturated throughput demands (since we are interested in the worst case scenarios and use TCP). Since TCP shares capacity across flows in a fair manner [22], i.e., flows get fair long-term end-to-end throughput, we do not model TCP behavior but rather just focus on the long-term throughput. We target the user association problem and hence focus on maximizing the aggregate throughput; however, since each user must be connected (7), the overall fairness is similar to what WiFi would offer after association.

Complexity matters: In the enterprise scenario of interest, a brute force approach to determine the optimal user assignment will incur prohibitively high complexity. For example, in our university setting, within an enclosure of office spaces there are more than 30 outlets into which extenders can be plugged in. The number of smartphones and laptops exceed this number. Even if one were to conservatively assume that there are 10 extenders plugged in and 30 devices, the complexity would be of the order of 30^{10} if a brute force approach were to be applied. More formally, our analysis of Problem 1 shows that it

Variable	Description
A	Set of PLC-WiFi extenders.
B_j	The maximum number of users that can be connected to extender j
c_j	The PLC PHY rate between the master PLC router and the extender j
N_j	Set of users associated with extender j. The $\{N_j\}$ form a partition of U.
r_{ij}	The WiFi PHY rate of user i when connected to extender j
t_{ij}	The WiFi throughput of user i when connected to extender j
$\mathcal{T}_j^{\text{WiFi}}$	The WiFi throughput across all users connected to extender j
$\mathcal{T}_j^{\text{PLC}}$	The backhaul PLC throughput of adaptor j
U	Set of users.
U_1, U_2	Set of users assigned in Phase I and Phase II of Alg. 1, respectively. U_1 and U_2 are a partition of U.
u_{ij}	Utility of assigning user i to extender j (see Phase I of Alg. 1.)
x_{ij}	binary variable indicating whether user i is connected to extender j

TABLE I: Table of Notations

is NP-hard, as proved in theorem 1 below. The key idea in the proof is to show a reduction from the partition problem [23] to a simple, particular instance of Problem 1 with two extenders and very high PLC rates. Since the partition problem is known to be NP-hard, then even this simple instance of Problem 1 is NP-hard, and hence the general case of Prob. 1 is NP-hard.

Theorem 1. *Prob. 1 is NP-hard.*

Proof: Let $S = \{w_1, w_2, \ldots, w_M\}$ be the inputs to the partition problem. Let $W \equiv \sum_{\ell=1}^{M} w_\ell$. Then we propose the following polynomial time transformation of the partition problem. If M is even: for $k = 0 : 2 : M-2$, solve Prob. 1 with $N = M + k$ users, where there are M "regular" users and k "dummy" users. The WiFi rates of the regular users are $r_{ij} = -\frac{1}{w_i}$ $\forall i = 1, 2, \ldots, M$, the WiFi rates of the dummy users are set as $r_{ij} = -\infty$ $\forall i = M+1, M+2, \ldots, M+k$. Also let there be two extenders $|A| = 2$, all with very good PLC rates $c_j = \infty$ $\forall j$, with at most $B_1 = B_2 = \frac{M+k}{2}$ users connected to each extender. We claim that this particular instance of Prob. 1 returns the optimal solution to the partition problem where one partition has at most $\frac{M+k}{2}$ elements (proved below). Then for each iteration of k, we solve this instance of Prob. 1, and pick the best solution across all iterations to solve the partition problem. Hence we have found a polynomial-time reduction from the partition problem to a particular instance of Prob. 1. If M is odd, we perform the above procedure but with $k = 1 : 2 : M-2$.

To show the claim above for each iteration of k, note that Prob. 1 maximizes $\frac{\frac{M+k}{2}}{-\sum_{i \in N_1} w_i} + \frac{\frac{M+k}{2}}{-\sum_{i \in N_2} w_i}$, which is equivalent to minimizing $\frac{\frac{M+k}{2}}{\sum_{i \in N_1} w_i} + \frac{\frac{M+k}{2}}{\sum_{i \in N_2} w_i} = \frac{M+k}{2W_1} + \frac{M+k}{2(W-W_1)}$, where $W_1 \equiv \sum_{i \in N_1} w_i$. This objective is minimized for $W_1^* = \frac{W}{2}$. Up to $\frac{M+k}{2}$ of the users connected to extender 1 could be regular users, corresponding to elements from one partition of S. Hence this particular instance of Prob. 1 (with two extenders and WiFi and PLC rates defined above) solves the partition problem with partition sizes of up to $\frac{M+k}{2}$. ■

B. Solutions for PLC-WiFi User Assignment

In this section, we describe our proposed solutions towards solving Prob. 1. We first provide intuition for our method

Algorithm 1 PLC-WiFi User Assignment

Inputs: Set of users U, set of extenders A, WiFi rate r_{ij}, PLC rate c_j
Output: user assignment x_{ij}
Variables: user index i, extender index j, task utilities u_{ij}, set of users assigned in phase 1 U_1, set of users assigned in phase 2 U_2

```
1: for i ← 1 to |U| do
2:     for j ← 1 to |A| do
3:         u_ij ← min(c_j/A, r_ij)
4: (U_1, {x*_ij}_{i∈U_1}) ← ASSIGNMENT_SOLVER({u_ij}, U, A)
5: {x*_ij}_{i∈U_2} ← PROBLEM_2({r_ij}, {c_j}, {x*_ij}_{i∈U_1}, U, A)
```

before describing it more formally. Because Prob. 1 is NP-hard (from Theorem 1), we *first* propose solving a modified version of Prob. 1 in what we call Phase I. These modifications involve (a) relaxing constraint (7), so that not every user has to be connected to an extender, and (b) modifying constraint (8) so that each extender has at least one connected user, *i.e.*, $\sum_{i \in N_j} x_{ij} \geq 1, \forall\ j$. The intuition behind relaxation (a) is that the aggregate system throughput can be maximized if not all users need to be assigned, as assigning more users causes contention on the WiFi/PLC links, and decreases aggregate system throughput (exactly how many users should be assigned are given by Theorem 2). The intuition behind modification (b) is to utilize all possible PLC backhaul links to increase the amount of throughput that can be provided by the system, by distributing the users across the possible extenders, potentially decreasing contention on the WiFi links and increasing aggregate throughput. Overall, making these modifications to Prob. 1 enables us to transform the problem exactly into an assignment problem (Theorem 2), and use standard polynomial-time algorithms [24] to solve the transformed assignment problem.

Then, in Phase II of our algorithm, we add back constraint (7) and assign the remaining users. Adding in these remaining users may lower the aggregate throughput compared to Phase I, but we try to do this in a way that minimizes the throughput degradation (*i.e.*, maximizes the aggregate throughput with the Phase I users fixed). We formulate this as a nonlinear program, and prove that this nonlinear program has integer solutions (Theorem 3) and so, no rounding mechanism is needed. We next provide further details on each phase.

Phase I: Under the modifications to constraints (7) and (8), we first characterize the solution to determine exactly how many users should be assigned to the extenders to maximize the aggregate throughput. On the one hand, connecting more users increases the number of flows, potentially increasing aggregate throughput. On the other hand, having more flows could cause contention and decrease aggregate throughput. In Lemma 2, we prove that exactly $|A|$ users should be assigned to the $|A|$ extenders to solve our modified Prob. 1. The key idea is to show that any candidate solution with more than one user assigned to an extender can be improved on by disconnecting an appropriate user from the extender to increase the aggregate throughput, until only one user remains per extender. Which user should be selected to disconnect and increase the aggregate throughput is given by Lemma 1, which states that a sufficient condition to increase (or maintain) the aggregate throughput is to disconnect the user with worse reciprical of the WiFi rate than the average of its peers connected to the same extender.

Lemma 1. *If user i connects to extender j with $\frac{1}{r_{ij}} \leq \frac{1}{|N_j|}\sum_{i' \in N_j} \frac{1}{r_{i'j}}$, the objective function (3) of Prob. 1 increases or stays the same.*

If user i currently connected to extender j with $\frac{1}{r_{ij}} \geq \frac{1}{|N_j|}\sum_{i' \in N_j} \frac{1}{r_{i'j}}$ is disconnected, the objective function (3) of Prob. 1 increases or stays the same.

Proof: $\frac{1}{r_{ij}} \leq \frac{1}{|N_j|}\sum_{i' \in N_j} \frac{1}{r_{i'j}}$ can be re-arranged to:

$$\frac{|N_j|}{\sum_{i' \in N_j} \frac{1}{r_{i'j}}} \geq \frac{|N_j|+1}{\sum_{i' \in N_j} \frac{1}{r_{i'j}} + \frac{1}{r_{ij}}} \tag{11}$$

The left and right hand sides correspond to the extender throughput $\mathcal{T}_j^{\text{WiFi}}$ before and after user i joined, respectively, implying that $\mathcal{T}_j^{\text{WiFi}}$ increased or stayed the same compared to without user i. $\mathcal{T}_j^{\text{WiFi}}$ increasing or staying the same implies that the objective function (3) increased or stayed the same, proving the claim. A similar analysis follows for the second claim. ∎

Lemma 2. *There exists an optimal solution for the modified Prob. 1 (with (7) relaxed, and (8) modified to $\sum_{i \in N_j} x_{ij} \geq 1, \forall\ j$), where exactly one user is connected to each extender.*

Proof: Proof by contradiction. Assume there does not exist any optimal solution where $\sum_i x^*_{ij} = 1\ \forall\ j$, *i.e.*, each optimal solution $\{x^*_{ij}\}$ has a non-empty set of extenders $\{j' : \sum_i x^*_{ij'} > 1\}$. Then you could construct a new solution by disconnecting a user chosen according to Lemma 1, which would cause the objective function to increase or stay the same. If the objective function increases, this contradicts the assumption that the x^*_{ij} were optimal. If the objective function stays the same, we can continue disconnecting users from extenders in the set $\{j'\}$ using Lemma 1, without decreasing the objective function, until we have removed all users except one, *i.e.*, $\sum_i x^*_{ij'} = 1$, contradicting the assumption that no such solutions existed. ∎

Having established that exactly $|A|$ users should be assigned to the $|A|$ extenders (because $\sum_i x^*_{ij} = 1$) in the modified Prob. 1, we next consider which particular $|A|$ users should be assigned, and to which extenders. Our main idea is to map the modified Prob. 1 into an assignment problem, which can then be solved using standard methods. The assignment problem takes as inputs a set of users, a set of tasks, and a set of utilities corresponding to each (task, utility) pair, and assigns users to tasks to maximize the aggregate utility. In our version of the assignment problem, we map each extender j to a task, and set the the utility u_{ij} of user i to task j as:

$$u_{ij} \equiv \min(\frac{c_j}{A}, r_{ij}) \tag{12}$$

This definition of task utility is crucial to ensure that the modified Prob. 1 can be mapped exactly into an assignment problem. In Theorem 2, we show that this task utility definition (12) results in an exact mapping between the modified Prob. 1 and the standard assignment problem. The key idea in the proof is to show that the modified Prob. 1 under Lemma 2 and (12) simplifies to the assignment problem.

Theorem 2. *The modified Prob. 1 (with (7) relaxed, and (8) modified to $\sum_{i \in N_j} x_{ij} \geq 1, \forall\ j$) is exactly an assignment problem with task utilities $u_{ij} = \min(\frac{c_j}{A}, r_{ij})$.*

Proof: We know from Lemma 2 that there is at least one optimal solution $\{x^*_{ij}\}$ such that:

$$\sum_i x^*_{ij} = 1 \ \forall\ j \tag{13}$$

This means that we can transform Prob. 1 into an equivalent assignment problem as follows. We can reduce the size of the feasible region without affecting optimality by adding (13) as a constraint to Prob. 1. Then, since (13) $\implies$ $|N_j| = 1$, we can simplify (6) as $t_{ij} = r_{ij} x_{ij}$, and (5) as $\mathcal{T}^{\text{WiFi}}_j = \sum_i r_{ij} x_{ij}$. Then the objective function (3) is $\sum_j \min(\sum_i r_{ij} x_{ij}, \frac{c_j}{A}) = \sum_j \sum_i \min(r_{ij}, \frac{c_j}{A}) x_{ij}$, where the last equality only holds because of (13). Thus, Prob. 1 without constraint (7) is transformed into an equivalent problem of maximizing $\sum_j \sum_i \min(r_{ij}, \frac{c_j}{A}) x_{ij}$, with constraints (13) and (10), which is the standard assignment problem. ■

Phase II: After assigning $|A|$ users in Phase I (we call this set of users U_1), we next turn our attention to assigning the remaining $|U| - |A|$ users (we call this set of users U_2). We seek to do this in a way that minimizes the impact of the U_2 users on the aggregate throughput, assuming that the U_1 users are fixed. We formulate this as Prob. 2 below:

Problem 2. *WiFi User Assignment Only*

$$\begin{aligned}
&\mathit{max}_{x_{ij}, i \in U_2} && \sum_{j=1}^{|A|} \mathcal{T}^{\mathit{WiFi}}_j && (14)\\
&\mathit{s.t.} && (5), (6), (7), (9) && (15)\\
& && 0 \leq x_{ij} \leq 1, \quad \forall\ i \in U, j \in A && (16)
\end{aligned}$$

The formulation has several differences from Prob. 1. First, we only need to make decisions for users $i \in U_2$, *i.e.*, those users who were not assigned in Phase I. Second, the objective function (14) maximizes the WiFi throughput $\mathcal{T}^{\text{WiFi}}_j$ only, compared to the original objective function (3) which included the PLC backhaul throughput $\mathcal{T}^{\text{PLC}}_j$. The intuition behind this is that the user assignments in Phase I already saturated the PLC backhaul (to maximize aggregate throughput); thus, any additional user assignments in Phase II will not change the PLC throughputs by much. Finally, Prob. 2 allows fractional solutions of x_{ij}; however, we prove next in Theorem 3 that Prob. 2 has integral optimal solutions. The key idea in the proof is to show that any user with a fractional assignment can shift to an integral assignment, increasing the aggregate throughput. Empirically, we find that numerically solving Prob. 2 results in these integral solutions.

Theorem 3. *There exists an integer solution to Prob. 2.*

Proof: Consider the contribution of user i to the total throughput. If user i is not connected to extender j, the extender has throughput $\mathcal{T}^{\text{WiFi}}_j[\text{before}] = \frac{|N_j|}{\sum_{i' \in N_j} \frac{1}{r_{i'j}}}$, and if user i is connected, the extender has throughput $\mathcal{T}^{\text{WiFi}}_j[\text{after}] = \frac{|N_j| + x_{ij}}{\sum_{i' \in N_j} \frac{1}{r_{i'j}} + r_{ij}}$. User i has a net contribution to the total throughput as follows:

$$\sum_{j: x_{ij} > 0} \mathcal{T}^{\text{WiFi}}_j[\text{after}] - \mathcal{T}^{\text{WiFi}}_j[\text{before}] \tag{17}$$

$$= \sum_{j: x_{ij} > 0} \left(\frac{x_{ij}}{\sum_{i' \in N_j} \frac{1}{r_{i'j}} + r_{ij}} - \left(\frac{|N_j|}{\sum_{i' \in N_j} \frac{1}{r_{i'j}}} - \frac{|N_j|}{\sum_{i' \in N_j} \frac{1}{r_{i'j}} + r_{ij}} \right) \right) \tag{18}$$

where the equality happens after re-arranging the $\mathcal{T}^{\text{WiFi}}_j$ terms. The first term represents the throughput contribution for a particular allocation of x_{ij}, and the second term is a constant for a given user i and extender j. If user i has a fractional assignment to extender k, *i.e.*, $0 < x_{ik} < 1$, it can increase the total throughput by shifting x_{ik} to another extender j' with minimum $\sum_{i' \in N_{j'}} \frac{1}{r_{i'j'}} + r_{ij'}$ (the denominator of the first term), thus increasing the first term in (18) for extender j', and eliminating the second term (which is strictly positive) for extender k (since $x_{ik} = 0$, so extender k will no longer be included in the summation), thus creating a net throughput increase. ■

Summary of Algorithm 1 Phases: Putting it all together, the complete algorithm is written in Alg. 1. In lines 1-3, we compute the task utilities to input to the assignment problem. In Line 4, we optimally solve the Phase I assignment problem using known techniques (*e.g.*, the Hungarian algorithm), to decide the which users should be associated to which extenders, for a subset of users $i \in U_1$ In Line 5, we numerically solve a nonlinear program to decide the assignments for the remaining users $i \in U_1$. Note that the re-distribution of PLC capacity allocations when certain PLC links are underutilized is implicitly handled by this approach.

Algorithm Complexity: The first phase of our algorithm runs in $\mathcal{O}(|A|^3)$, where $|A|$ is the number of the PLC extenders and $|A|^3$ is the runtime of the *Hungarian algorithm* [24], [25]. The runtime of the second phase of our algorithm depends on the stopping criterion of our numerical solver which uses the interior point method; the solver stops when the improvement in the aggregate throughput is less than e^{-5}.

V. Implementation & Evaluation

In this section, we briefly describe WOLT's implementation and detailed evaluations via both small scale real experiments and larger scale high-fidelity simulations.

A. Implementation Details

PLC Testbed configuration and equipment: Our testbed consists of seven laptops from four different vendors (two Lenovo Ideapads 300S-14ISK, two Dell Inspiron 15s, one Acer Aspire E15, one Apple MacBook A1278 and one Apple MacBook Air A1370), and one central server running Windows 10, 64-bit. The testbed is equipped with three TP-Link

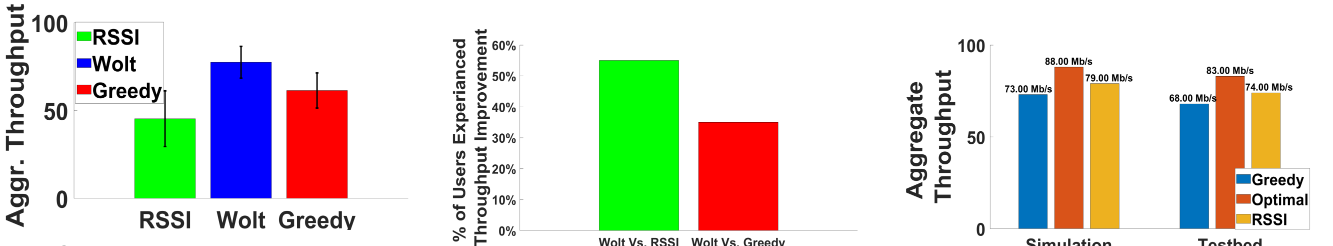

(a) WOLT, Greedy and RSSI Comparison on Testbed

(b) Per User Effects of WOLT

(c) Validating the Fidelity of our Simulations

Fig. 4: Experimental results

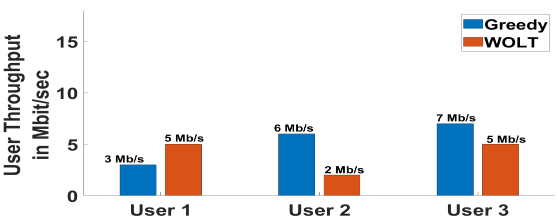

(a) WOLT vs. Greedy per-user throughput comparison for the poorer three users

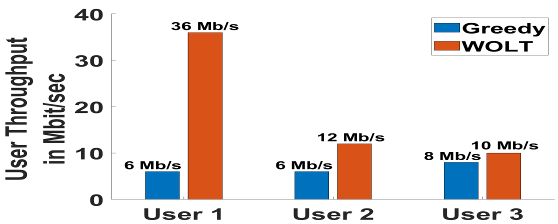

(b) WOLT vs. Greedy per-user Throughput Comparison for the best three users

Fig. 5: WOLT's Effects on Users' Throughputs

TL-WPA8630 extenders and one TP-Link TL-PA8010 central unit. These extenders support up to 1200 Mbit/sec at the PHY layer. The three extenders interface with the central unit via the PLC backhaul and with the users via WiFi. The central unit's role is to connect the three extenders to the master router through Ethernet. The central unit is a gateway for all the communications flowing between the extenders and the server.

Software implementation and WiFi details: We implement WOLT in Java as a user-space utility that runs on users' devices as well as the server. We name the server the Central Controller (*CC*). When a user arrives (needs association), it scans all available networks and estimate the WiFi channel quality of each extender. The network interface card (NIC) driver provides information on the modulation and coding scheme used for each WiFi channel, which is used to estimate the transmission bit-rate between the user and the extender. As verified in prior work such as [2], when a small number of APs are used, each operates on a non-overlapping 802.11 channel, and thus is able to operate interference free; thus, we assume that each extender operates on an non-overlapping channel relative to its neighbor extenders on the WiFi domain. The users (clients) gather this information on the reachable extenders and sends it to the *CC*. Note that a new user initially connects to the extender with the highest RSSI to communicate with the server and later switches extenders if needed, based on the new assignment from the *CC*.

Offline PLC backhaul link capacity estimation: The PLC backhaul link capacities are measured using `iperf3` [17]. We connect a machine to the PLC extender by an Ethernet cable and saturate the PLC link between that extender and the *CC*. The maximum amount of traffic the PLC link can deliver is then considered to be the capacity (rate in isolation) of the link. One can also potentially use Qualcomm Atheros Open Powerline Toolkit [26] to measure the PHY rate between PLC extender; unfortunately however, this tool is not compatible with the more recent AV2 PLC extenders we are using.

Simulation of large-scale WiFi networks with PLC backhaul: To consider larger scales than what our experiments support, we simulate a WiFi network with ten extenders, each connected to the *CC* via a PLC backhaul. We calibrate our simulator with PLC link capacities measured from different outlets in a university building. The user association requests arrive and depart the network according to Poisson distribution [27] with arrival rate of 3 and departure rate of 1. We use a simple model to simulate the WiFi channel qualities where the channel quality is a function of the distance between the extender and the user [28]. A 100 m $\times$ 100m 2D plane with 15 extenders and two hundred users is created. The users are geographically randomly distributed in the plane. The distance between every user and extender is computed and the corresponding WiFi channel is estimated.

B. *Greedy baseline for comparison (called Greedy)*

We compare WOLT against a greedy algorithm with which, each newly associating user is assigned such that the aggregate throughput after assignment is maximized. If there is no room for improving the aggregate throughput, the greedy algorithm will assign the user to the extender with the least impact on the aggregate throughput. The greedy algorithm computes all possible aggregate throughputs of the network when the new user is connected to different extenders and assigns the user to the extender that gives the highest aggregate throughput.

When a user arrives, it estimates the RSSIs of all the available WiFi APs and connects conventionally to the one with highest RSSI. The user communicates its WiFi channel estimations to the *CC* and waits for the response. Once the *CC* receives a new user message, it computes the greedy assignment that maximizes the aggregate throughput and sends an association directive back to the user. Upon receipt, the user

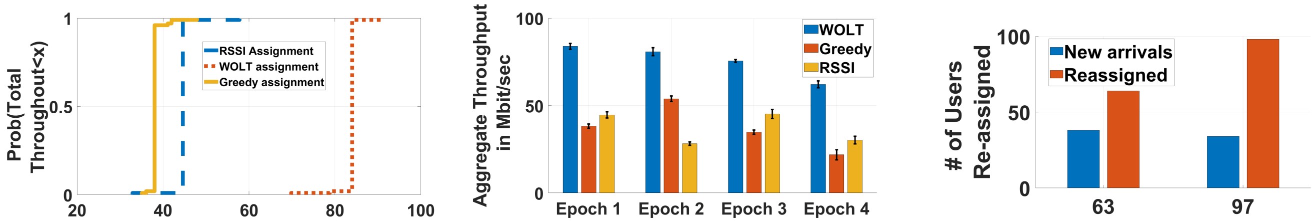

(a) CDF of aggregate throughput. (b) Average aggregate throughput over epochs. (c) Number of user re-assignments.

Fig. 6: Simulation results.

associates itself to the corresponding extender. Note that no reassignment of the other users is done (as done with WOLT).

C. RSSI baseline for comparison (called RSSI)

With the *RSSI* baseline, users are associated to the extender that yields the strongest received signal regardless of (a) the quality of the extender's PLC link segment, (b) how many users are contending in the WiFi cell for that extender. Once the user is connected, it provides an estimate of its WiFi capacity (throughput) to the *CC*. The *CC* has the knowledge about the capacity of each PLC link as well as what users connected to which extender. Unlike *Greedy*, users do not expect association directives to be received from the *CC* and remain associated with the extender with the highest RSS. It is worth mentioning that this assignment policy is the default on PLC-WiFi extenders today.

D. Experimental evaluations

Improvement in aggregate throughput: We perform experiments on our testbed with three extenders and 7 laptops in a university laboratory of 2408 m^2 area with several tables and chairs, computer equipment and two cubicles. We randomly picked three power outlets (among 10 outlets that are available) and moved the laptops around to create 25 different topologies. The results are shown in Figures 4a and 4b. In the first, we show the average throughputs when each algorithm is used. We see that WOLT outperforms both *Greedy* and *RSSI*. Average aggregate throughput improvements of 26% and 70% are observed over *Greedy* and *RSSI* respectively.

Per user effects: In the second figure, we show the percentage of users that enjoy an increase or suffer a degradation with WOLT as compared to *Greedy* and *RSSI*. We see that 35% of the users have a better throughput when using WOLT as compared to *Greedy* (65% experience a degradation). As compared to *RSSI*, 55% of the cases enjoy better throughputs with WOLT (45% experience a degradation). These changes occur because the objective with WOLT is to improve network throughput; while doing so WOLT's configurations benefit some users as compared to the baselines while disadvantaging others as one might expect.

Fairness: WOLT's objective is to maximize the network-wide throughput as discussed earlier. Thus, while formulating the problem for optimal user assocation, we focused on efficiency rather than fairness, so it can be expected that the fairness with WOLT will be penalized. Given this, we perform experiments to evaluate its fairness. Before we present our results, we point out that WOLT will not leave users unassociated (constraint (7) in Prob. 1).

Towards maximizing throughput, WOLT tries to ensure that the users with the best end-to-end channel qualities (i.e., both on the PLC and WiFi components) achieve their maximum throughputs that they can get; while doing so it could disadvantage users with poor channel qualities. To show that this effect is not significant[1], we consider the three users with the highest throughputs and the three users with the lowest throughputs in a randomly chosen topology in our experiment (we find that the results are very similar with all our scenarios).

In Figures 5a and 5b, we depict the individual *Greedy* and WOLT throughput for the three worst and best users in WOLT respectively. Note that with *Greedy*, all users (good and bad) try to get the best throughputs they can and thus, one can use this as a performance baseline for how well they can do. The first figure shows two out of the three poorest users (User 2 and 3) receive a better throughout with *Greedy* than they do with WOLT while one user (User 1) still has a better throughput with WOLT over *Greedy*. However, the loss of aggregate throughput of the worst three users when using WOLT compared to *Greedy* in Fig. 5a is only (in total) about 6 Mbps. On the other hand, the best three users (depicted in Fig. 5b) improve their throughouts to a total of about 38 Mbit/sec (30 Mbit/sec for User 1, 6 Mbit/sec for User 2 and 2 Mbit/sec for User 3). This shows that the modest hit taken by the poor users (a relatively low penalty in fairness) results in a significant throughput boost for the good users. In other words, our experiments show that WOLT offers its throughput improvements while only taking a modest hit in terms of fairness.

Fidelity of our simulations via comparison with experimental results from our testbed: We perform a few experiments where we mimic our experimental scenario in our simulation. Our objective is to compare the results across the two towards getting confidence that our simulations yield realistic results in larger scale settings. We show one such result (for a single topology since we need to make sure that the results hold for all topologies considered) in Figure 6c. We show the results for both from experiments from our testbed and our simulations (we have three extenders and seven users in the latter with the same channel qualities). We see that the results are very consistent with what we obtain in our experiments showing the fidelity of our simulations. Given this, we next present some larger scale scenarios that we

[1]Since the set up is small here, we do not consider a fairness metric such as the Jain's fairness index [29] here; we do so later when discussing our simulation results.

simulate to demonstrate that WOLT performs well even in such cases.

E. Simulation results

Total throughputs: We simulate the performance of the WOLT and the greedy algorithm with the simulation settings discussed in §V-A. We run 100 trials when there are $|U| = 36$ users in the area of interest, and plot the CDF of the aggregate (total) network throughputs across trials in Fig. 6a. We see that WOLT outperforms the greedy algorithm in all trials, with WOLT providing an average improvement (in terms of aggregate throughput) of 2.5x over the greedy approach. Compared to the experimental results in Fig. 4b, we see the relative improvement of WOLT over greedy is larger; we posit this is because the simulation contains a larger number of users with more uniform distribution of users with good and poor WiFi channel qualities; this fully exploits WOLT's potential i.e., it can properly assign users with poor channel qualities to maximize the aggregate throughput.

Online behavior of WOLT: Next, we examine the temporal dynamics of WOLT. As explained in §V-A, users arrive and depart from the system according to the Poisson distribution, with a net average increase of 33 users per epoch. In Fig. 6b, we plot the aggregate throughput of WOLT after each epoch. As more users join the system ($|U|$ increases from 36 in epoch 1, to 66 in epoch 2, to 102 in epoch 3), the aggregate throughput of the network gradually increases and saturates (not shown). At the same time we compare WOLT's performance with that of *Greedy* (recall that *Greedy* assigns the user one by one as they arrive in the current epoch). Our results show that WOLT outperforms *Greedy* even as the number of users increases to over 100.

Fairness: To evaluate the fairness we obtain with WOLT, here we consider the Jain's fairness index, comparing the metric with WOLT with that achieved with *Greedy* and *RSSI*. The results are consistent across our simulation experiments and we find that they are on average, 0.66, 0.52 and 0.65, respectively for WOLT, *Greedy* and *RSSI*, with minor deviations across experiments. This demonstrates that even though WOLT does not explicitly consider fairness among users, it has even better (or at least comparable) fairness than the other baseline policies that are considered.

Finally, we wish to examine how the user associations change over time as users arrive and depart from the system. In Fig. 6c, we plot the number of users who are re-associated by WOLT at the end of every epoch due to these user dynamics. WOLT re-assigns up to twice the number of arriving users (i.e., one user is swapped for every new user who arrives, on average), which intuitively makes sense as WOLT needs to re-assign some existing users to form a more optimal solution. The key observation is that the number of reassignments for each newly associating user is relatively low on average.

VI. Related Work

WiFi User Association: There are several papers that try to automatically configure a WiFi network in terms of appropriate associations of users with APs, towards optimizing a performance metric (mostly throughput with various fairness requirements); examples include [30]–[32] among others. These efforts are different from our work for the following reasons. First, they ignore the impact of the backhaul network, which usually is Ethernet, and only consider the wireless links. However, when we consider PLC as the backhaul which a set of WiFi extenders share, we need to account for the contention on the power line medium. Stated otherwise, to the best of our knowledge all past efforts assume that the WiFi networks last link is the bottleneck in end-to-end connections. With plug and play extenders, PLC can become a bottleneck if it provides throughput less than WiFi links.

PLC: Atya et al., [6] propose BOLT a learning-based algorithm to orchestrate flows in a PLC network. Vlachou et al. [33] propose a model to improve throughput of IEEE 1901 by modifying existing MAC parameters. However, both papers do not consider WiFi extenders which are today the most common means of utilizing PLC capacity. In [9], [7], and [33], the authors assume that PLC links support the same physical rates and do not perform experiments with differing PLC link qualities as done in this paper.

Hybrid WiFi-PLC: Vidyut [34] studies the use of PLC as a medium for delivering reference signals for wireless communications to enhance the throughputs of multi-cell MIMO systems. In [35], the authors study the performance of power line communications in terms of throughput and its potential for being used a backhaul network for WiFi front ends. Apicharttrisorn et al., [5] perform a measurement study of HomePlug AV2-compliant WiFi extenders. These studies however, do not consider user association problems.

EMPoWER [36] proposes congestion control algorithms sitting between MAC and IP layers of hybrid WiFi-PLC networks where each node is capable of WiFi or PLC or both. However, they fix the connectivity between nodes and do not target improving aggregate netowrk throughput via intelligent user association. Electri-Fi [37] studies the characteristics of PLC and WiFi networks in terms of thier spacial and temporal variations and reports analysis of different causes of retransmissions in PLC; again this work does not consider network throughput interactions between WiFi and PLC that affect how users must be associated towards optimizing throughput.

Hybrid Cellular-Adhoc Networks: In [38], [39], extending cellular network coverage with wireless adhoc connectivity is considered. While such networks contain concatenated links, cellular acess typically allows users to reserve capacity unlike in PLC, where the share of the capacity that a user obtains is dictated by the extender with which it associates and which other users share the two parts of the concatenated link.

VII. Conclusions

In this paper, we develop WOLT, which tries to maximize the network throughput in a hybrid PLC-WiFi network by optimally assigning the client devices to the available WiFi enabled PLC extenders. The challenge that we address is that unlike in WiFi networks with an Ethernet backhaul, the PLC links could be of lower capacity than their WiFi counterparts. We therefore need to account for the "bottleneck" capacity provided by an extender to a given client when making association decisions. We show that the optimal allocation of users to WiFi based extenders is NP-hard and solve a relaxed version wherein we make several constraints less stringent. The algorithms that we design towards this form the basis for WOLT. We show via experiments on our testbed and high-fidelity simulations that WOLT outperforms a baseline central greedy approach by as much as 2.5x in terms of the average network throughputs that are achieved.

VIII. Acknowledgment

This work was partially supported by the NSF NeTS grant 1528095.

References

[1] "Best wi-fi extenders of 2019: top devices for boosting your wifi network," https://www.techradar.com/news/the-best-wi-fi-extenders.

[2] B. Kauffmann, F. Baccelli, A. Chaintreau, V. Mhatre, K. Papagiannaki, and C. Diot, "Measurement-based self organization of interfering 802.11 wireless access networks." in *Infocom*, vol. 7. Citeseer, 2007, pp. 1451–1459.

[3] V. Shrivastava, S. Rayanchu, S. Banerjee, and K. Papagiannaki, "Pie in the sky: Online passive interference estimation for enterprise wlans," in *Proceedings of the 8th USENIX Conference on Networked Systems Design and Implementation*, ser. NSDI'11. Berkeley, CA, USA: USENIX Association, 2011, pp. 337–350. [Online]. Available: http://dl.acm.org/citation.cfm?id=1972457.1972492

[4] I. Broustis, K. Papagiannaki, S. V. Krishnamurthy, M. Faloutsos, and V. P. Mhatre, "Measurement-driven guidelines for 802.11 wlan design," *IEEE/ACM Transactions on Networking*, June 2010.

[5] K. Apicharttrisorn, A. O. F. Atya, J. Chen, K. Sundaresan, and S. V. Krishnamurthy, "Enhancing wifi throughput with plc extenders: A measurement study," in *Passive and Active Measurement*, 2017.

[6] A. O. F. Atya, K. Sundaresan, S. V. Krishnamurthy, M. A. Khojastepour, and S. Rangarajan, "Bolt: Realizing high throughput power line communication networks," in *Proceedings of the 11th ACM Conference on Emerging Networking Experiments and Technologies*, ser. CoNEXT '15. ACM.

[7] C. Vlachou, A. Banchs, J. Herzen, and P. Thiran, "On the mac for power-line communications: Modeling assumptions and performance tradeoffs," in *Network Protocols (ICNP), 2014 IEEE 22nd International Conference on*. IEEE, 2014, pp. 456–467.

[8] S. Goldfisher and S. Tanabe, "Ieee 1901 access system: An overview of its uniqueness and motivation," *IEEE Communications Magazine*, vol. 48, no. 10, pp. 150–157, 2010.

[9] C. Vlachou, J. Herzen, and P. Thiran, "Fairness of mac protocols: Ieee 1901 vs. 802.11," in *2013 IEEE 17th International Symposium on Power Line Communications and Its Applications*. IEEE, 2013, pp. 58–63.

[10] I. S. Association *et al.*, "Ieee standard for broadband over power line networks: Medium access control and physical layer specifications," *IEEE Std*, vol. 2010, no. 2010, pp. 1–1586, 1901.

[11] L. Yonge, J. Abad, K. Afkhamie, L. Guerrieri, S. Katar, H. Lioe, P. Pagani, R. Riva, D. M. Schneider, and A. Schwager, "An overview of the homeplug av2 technology," *Journal of Electrical and Computer Engineering*, vol. 2013, 2013.

[12] S. Katar, B. Mashburn, K. Afkhamie, H. Latchman, and R. Newman, "Channel adaptation based on cyclo-stationary noise characteristics in plc systems," in *2006 IEEE International Symposium on Power Line Communications and Its Applications*. IEEE, 2006, pp. 16–21.

[13] "Tp-link powrline ac wifi kit," https://www.tp-link.com/us/home-networking/powerline/tl-wpa8630-kit/.

[14] "Tp-link 300mbps wi-fi range extender, av600 powerline edition," https://www.tp-link.com/us/home-networking/powerline/tl-wpa4220kit/.

[15] M. Heusse, F. Rousseau, G. Berger-Sabbatel, and A. Duda, "Performance anomaly of 802.11 b," in *IEEE INFOCOM 2003. Twenty-second Annual Joint Conference of the IEEE Computer and Communications Societies (IEEE Cat. No. 03CH37428)*, vol. 2. IEEE, 2003, pp. 836–843.

[16] E. Aryafar, A. Keshavarz-Haddad, M. Wang, and M. Chiang, "Rat selection games in hetnets," in *2013 Proceedings IEEE INFOCOM*. IEEE, 2013, pp. 998–1006.

[17] "Tcp/udp bandwidth measurement tool," https://iperf.fr/.

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

[19] Q. Ye, B. Rong, Y. Chen, M. Al-Shalash, C. Caramanis, and J. G. Andrews, "User association for load balancing in heterogeneous cellular networks," *IEEE Transactions on Wireless Communications*, vol. 12, no. 6, pp. 2706–2716, 2013.

[20] R. Mahindra, H. Viswanathan, K. Sundaresan, M. Y. Arslan, and S. Rangarajan, "A practical traffic management system for integrated lte-wifi networks," in *Proceedings of the 20th annual international conference on Mobile computing and networking*. ACM, 2014, pp. 189–200.

[21] W. Lingzhi, H. Cunqing, Z. Rong, and N. Rui, "Online channel selection and user association in high-density wifi networks," in *2015 IEEE International Conference on Communications (ICC)*. IEEE, 2015, pp. 1571–1576.

[22] D. J. Leith and P. Clifford, "Tcp fairness in 802.11 e wlans," in *2005 International Conference on Wireless Networks, Communications and Mobile Computing*, vol. 1. IEEE, 2005, pp. 649–654.

[23] S. Chopra and M. R. Rao, "The partition problem," *Mathematical Programming*, vol. 59, no. 1-3, pp. 87–115, 1993.

[24] T. H. Cormen, C. E. Leiserson, R. L. Rivest, and C. Stein, *Introduction to algorithms*. MIT press, 2009.

[25] G. A. Mills-Tettey, A. Stentz, and M. B. Dias, "The dynamic hungarian algorithm for the assignment problem with changing costs," 2007.

[26] "Qualcomm atheros open powerline toolkit," https://github.com/qca/open-plc-utils.

[27] A. Balachandran, G. M. Voelker, P. Bahl, and P. V. Rangan, "Characterizing user behavior and network performance in a public wireless lan," in *ACM SIGMETRICS Performance Evaluation Review*, vol. 30, no. 1. ACM, 2002, pp. 195–205.

[28] "Cisco aironet 1200 series access point," http://www.vsp.virginia.gov/downloads/STARSContract/Appendix

[29] R. Jain, A. Durresi, and G. Babic, "Throughput fairness index: An explanation," in *ATM Forum contribution*, vol. 99, no. 45, 1999.

[30] M. AMER, A. Busson, and I. Guérin Lassous, "Association optimization in wi-fi networks: Use of an access-based fairness," in *Proceedings of the 19th ACM International Conference on Modeling, Analysis and Simulation of Wireless and Mobile Systems*, ser. MSWiM '16.

[31] O. B. Karimi, J. Liu, and J. Rexford, "Optimal collaborative access point association in wireless networks," in *IEEE INFOCOM 2014 - IEEE Conference on Computer Communications*. IEEE, 2014.

[32] A. P. Jardosh, K. Mittal, K. N. Ramachandran, E. M. Belding, and K. C. Almeroth, "Iqu: Practical queue-based user association management for wlans," in *Proceedings of the 12th Annual International Conference on Mobile Computing and Networking*, ser. MobiCom '06. ACM, 2006.

[33] C. Vlachou, A. Banchs, J. Herzen, and P. Thiran, "Analyzing and boosting the performance of power-line communication networks," in *Proceedings of the 10th ACM International on Conference on Emerging Networking Experiments and Technologies*, ser. CoNEXT '14. ACM, 2014.

[34] V. Yenamandra and K. Srinivasan, "Vidyut: Exploiting power line infrastructure for enterprise wireless networks," in *Proceedings of the 2014 ACM Conference on SIGCOMM*, ser. SIGCOMM '14. ACM, 2014.

[35] N. Roy, D. Kleinschmidt, J. Taylor, and B. Shirazi, "Performance of the latest generation powerline networking for green building applications," in *Proceedings of the 5th ACM Workshop on Embedded Systems For Energy-Efficient Buildings*, ser. BuildSys'13. ACM.

[36] S. Henri, C. Vlachou, J. Herzen, and P. Thiran, "Empower hybrid networks: Exploiting multiple paths over wireless and electrical mediums," in *Proceedings of the 12th International on Conference on Emerging Networking EXperiments and Technologies*, ser. CoNEXT '16. ACM.

[37] C. Vlachou, S. Henri, and P. Thiran, "Electri-fi your data: Measuring and combining power-line communications with wifi," in *Proceedings of the 2015 Internet Measurement Conference*, ser. IMC '15. ACM, 2015.

[38] L. K. Law, S. V. Krishnamurthy, and M. Faloutsos, "Capacity of hybrid cellular-ad hoc data networks," in *IEEE INFOCOM 2008 - The 27th Conference on Computer Communications*. IEEE, 2008.

[39] B. Liu, Z. Liu, and D. Towsley, "On the capacity of hybrid wireless networks," in *IEEE INFOCOM 2003. Twenty-second Annual Joint Conference of the IEEE Computer and Communications Societies (IEEE Cat. No.03CH37428)*. IEEE.