

# UC Riverside

## UC Riverside Electronic Theses and Dissertations

### Title

Harnessing the Benefits of Indoor PLC Networks

### Permalink

<https://escholarship.org/uc/item/032466wr>

### Author

Alhulayyil, Hisham A

### Publication Date

2022

Peer reviewed|Thesis/dissertation

UNIVERSITY OF CALIFORNIA  
RIVERSIDE

Harnessing the Benefits of Indoor PLC Networks

A Dissertation submitted in partial satisfaction  
of the requirements for the degree of

Doctor of Philosophy

in

Computer Science

by

Hisham Abdulrahman O Alhulayyil

March 2022

Dissertation Committee:

Dr. Srikanth V. Krishnamurthy, Chairperson  
Dr. Nael Abu-Ghazaleh  
Dr. Michail Faloutsos  
Dr. Jiasi Chen

Copyright by  
Hisham Abdulrahman O Alhulayyil  
2022

The Dissertation of Hisham Abdulrahman O Alhulayyil is approved:

---

---

---

---

Committee Chairperson

University of California, Riverside

## Acknowledgments

I would like to thank the University of California, Riverside, and the Department of Computer Science and Engineering for the facilities and resources that are provided to me throughout my PhD study.

In addition, I am incredibly grateful to my advisor, Professor Srikanth Krishnamurthy for his invaluable advice, continuous support, and patience during my PhD study. His immense knowledge has encouraged me in all the time of my academic research and daily life.

I would also like to thank Professor Jiasi Chen and Professor Karthikeyan Sundaresan for tirelessly training and preaching me to be a better scientist everyday.

My sincere appreciation to the government of the Kingdom of Saudi Arabia for providing the financial and social support throughout my undergraduate and graduate studies. I would like to thank Al Imam Mohammad Ibn Saud Islamic University for their financial support to study in the US during both my master's and PhD degrees. Also, my thanks goes to the Saudi Arabian culture mission for managing my scholarship.

My appreciation also goes out to my friends—Kittipat Apicharttrisorn, Abdulrahman Bin Rabiah and Ahmad Bin Rabiah for a cherished time spent together.

Special thanks to my wife, Eman, who has stood by me through all my travails, my absences, my fits of pique and impatience. She gave me support and help, discussed ideas and prevented several wrong turns. She also supported the family during all my graduate studies. *You have made me stronger, better and more fulfilled than I could have ever imagined. I love you to the moon and back.*

I am forever indebted to my father, Professor Abdulrahman, and my mother, Munirah, for their unconditional love and for giving me the opportunities and experiences that have made me who I am.

To my children, Sadan and Abdulrahman. You are the most memorable and valuable thing that happened to me during my graduate studies.

To my family in the Kingdom of Saudi Arabia for their continuous and unparalleled love, help and support. I am grateful to my sisters- Maymonah, Omaymah, Asma, Norah, Ghadah and Shaden and my two brothers- Mohammed and Hani for always being there for me.

In memory of my beloved uncle Abdulkarim, who passed away during my graduate studies and I did not have the chance to say good bye to him.

# ABSTRACT OF THE DISSERTATION

Harnessing the Benefits of Indoor PLC Networks

by

Hisham Abdulrahman O Alhulayyil

Doctor of Philosophy, Graduate Program in Computer Science

University of California, Riverside, March 2022

Dr. Srikanth V. Krishnamurthy, Chairperson

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 the end-to-end throughput for the concatenated PLC-WiFi link is determined by the bottleneck segment of the aggregated WiFi-PLC link.

The main objective of this dissertation is to maximize the aggregate of WiFi-PLC networks. The main challenge we address is that PLC capacity could be of a lower capacity than the WiFi link. Thus, users that are connected via a good WiFi link could suffer from the poor PLC link associated with the PLC extender they are attached to. Therefore, we develop WOLT, a system that connects each user to a signal extender for the sake of maximizing the aggregate throughput. It does so while accounting for the PLC link capacities and the users' WiFi link qualities. The results obtained from real testbed experiments and high-fidelity simulations and they show that WOLT is capable of increasing the aggregate throughput by more than 2.5x compared to a greedy user association baseline.

Moreover, users that are occluded from the main router could suffer from well-known WiFi impairments such as deep fading and shadowing. To overcome this issue, Distributed Antenna Systems (DAS) have been shown to improve the robustness and stability of the WiFi link. We approach this problem by two solutions. First, we develop PLC-DAS, a system that finds the correct DAS set of extenders for each user in home settings. We compare PLC-DAS against three baselines and across three fairness models. The results from our simulations show an improvement up to 4.5x compared to blindly using all WiFi-PLC extenders to form a DAS transmitter, while maintaining a fairness Jain's index value of at least 0.97 with proportional and max-min fairness models.

Secondly, we develop Priza, a scalable system that clusters WiFi-PLC extenders into cells and assigns frequencies to each cell. Priza is evaluated via real testbed experiments and high-fidelity simulations and it can increase the aggregate throughput by more than 3x compared to the baseline that creates as large DAS cells as possible.



# Contents

<b>List of Figures</b>	<b>x</b>
<b>List of Tables</b>	<b>xi</b>
<b>1 Introduction</b>	<b>1</b>
<b>2 WOLT: Auto-Configuration of Integrated Enterprise PLC-WiFi Networks</b>	<b>7</b>
2.1 Introduction . . . . .	8
2.2 PLC Background in Brief . . . . .	12
2.3 Measurement Study . . . . .	13
2.3.1 Medium sharing in the PLC and WiFi domains . . . . .	13
2.3.2 Showcasing the need for informed user association . . . . .	18
2.4 Problem Statement & Solutions . . . . .	19
2.4.1 Problem Statement and Hardness . . . . .	20
2.4.2 Solutions for PLC-WiFi User Assignment . . . . .	24
2.5 Implementation & Evaluation . . . . .	30
2.5.1 Implementation Details . . . . .	31
2.5.2 Greedy baseline for comparison (called Greedy) . . . . .	32
2.5.3 RSSI baseline for comparison (called RSSI) . . . . .	33
2.5.4 Experimental evaluations . . . . .	34
2.5.5 Simulation results . . . . .	37
2.6 Related Work . . . . .	39
2.7 Conclusions . . . . .	40
<b>3 Boosting Home WiFi Throughputs via Adaptive DAS Clustering of PLC Extenders</b>	<b>42</b>
3.1 Introduction . . . . .	43
3.2 PLC Background in Brief . . . . .	48
3.3 Measurements . . . . .	49
3.3.1 Feasibility of PLC as a backhaul for DAS . . . . .	49
3.3.2 DAS side issues . . . . .	51
3.4 Problem Statement & Solutions . . . . .	55

3.4.1	Problem Statement . . . . .	55
3.4.2	Algorithms for DAS Cell Formation and WiFi Time Allocation . . . . .	63
3.5	Evaluations . . . . .	65
3.5.1	Simulation Details . . . . .	65
3.5.2	Results . . . . .	67
3.6	Related Work . . . . .	73
3.7	Conclusions . . . . .	74
<b>4</b>	<b>Priza: Throughput-efficient DAS Clustering of WiFi-PLC Extenders in Enterprises</b>	<b>76</b>
4.1	Introduction . . . . .	77
4.2	PLC Background in brief . . . . .	82
4.3	Factors influencing PLC based DAS clusters . . . . .	83
4.4	Measurement Study . . . . .	85
4.4.1	Distribution of PLC capacities . . . . .	88
4.4.2	The need for informed WiFi-PLC DAS clustering . . . . .	90
4.5	Formulating the problem . . . . .	97
4.6	Priza and its components . . . . .	105
4.7	IMPLEMENTATION & EVALUATION . . . . .	108
4.7.1	Implementation Details . . . . .	108
4.7.2	Simulations . . . . .	109
4.7.3	Baselines & Performance Metrics . . . . .	110
4.7.4	Experimental Evaluations . . . . .	110
4.7.5	Simulation Results . . . . .	111
4.8	Related Work . . . . .	114
4.9	Conclusions . . . . .	115
<b>5</b>	<b>Conclusions</b>	<b>116</b>
	<b>Bibliography</b>	<b>117</b>

# List of Figures

2.1	Powerline communications extend WiFi coverage through existing interior power lines in a home/enterprise. . . . .	9
2.2	Medium sharing in the PLC and WiFi domains. . . . .	14
2.3	Experiments in a scenario observed on our testbed with user association policies. . . . .	17
2.4	Experimental results . . . . .	34
2.5	WOLT's Effects on Users' Throughputs . . . . .	34
2.6	Simulation results. . . . .	37
3.1	Powerline communications extend WiFi coverage through existing interior power lines in a home. . . . .	44
3.2	CDF of latencies experienced on different PLC links. . . . .	50
3.3	DAS experiments showing feasibility of PLC as a backhaul for DAS, DAS performance gains, and DAS synchronization issues. . . . .	53
3.4	Toy example of different possible DAS cell formation solutions. . . . .	59
3.5	CDF of aggregate throughput with different fairness models. . . . .	67
3.6	Average aggregate throughput. The results show how PLC-DAS outperforms other baselines. . . . .	68
3.7	PLC backhaul impact. . . . .	70
3.8	Home area impact. . . . .	71
4.1	PLC capacities in different labs in our university setting. . . . .	89
4.2	Snapshots of the area surrounded by solid line in lab 3 from Fig.1 with three different network setups. . . . .	91
4.3	Algorithm for frequency assignments and DAS Cell Formation . . . . .	104
4.4	Experimental results. . . . .	111
4.5	Simulation Results. . . . .	112
4.6	Average aggregate throughputs. The results show how the size of the topologies affects the Reuse and B-DAS baselines. . . . .	112

# List of Tables

2.1	Table of Notations . . . . .	20
3.1	Table of Notations . . . . .	56
3.2	Jain’s fairness index . . . . .	69
3.3	Aggregate throughput comparison between real house layouts and randomly generated layouts . . . . .	73
3.4	Fairness comparison between real house layouts and randomly generated layouts	73
4.1	Table of Notations . . . . .	99
4.2	Simulations Jain’s fairness index . . . . .	114

# Chapter 1

## Introduction

Powerline communication (PLC) is a technology that enables sending and receiving data over existing electricity wires. Its importance stems from the fact that it does not require building any networking infrastructure. Unlike the traditional way of establishing a dedicated physical medium for network communications, PLC leverages the pre-existing power wires in the walls as their medium of communication. Thus, in buildings that lack a proper networking infrastructure, PLC provides an easy way to deploy network devices such as access points, printers and range extenders. The PLC is controlled by the IEEE 1901 protocol which, in essence, is similar to the IEEE 802.11 protocol but with some differences in terms of fairness, complexity and performance. Modern PLC extenders are empowered with WiFi interfaces which allow users to connect to PLC extenders wirelessly. Such extenders are capable of expanding the coverage of the networks to the areas where the signal from the main router is weak or unstable. Large vendors such as Cisco, Netgear and TP-Link supply the market with PLC extenders that support up to 2024 Mbit/sec on the PHY rate

which makes these extenders attractive alternatives for conventional range extenders that use Ethernet as a backhaul.

Despite the fact that PLC extenders are capable of widening the coverage of the network, they could, in fact, reduce the aggregate throughput if they deployed naively. Specifically, WiFi-PLC extenders have two interfaces, a PLC interface and a WiFi interface. These two interfaces are jointly used to deliver the data to the users attached to the WiFi-PLC extenders. For the case of the downlink traffic, the packets are sent from the gateway to the WiFi-PLC extenders over the PLC backhaul. The WiFi-PLC extenders then transmit the packets to the intended users over the WiFi backhaul. Therefore, when a user is connected to the WiFi-PLC extenders, the link between the user and the gateway in the network is composed of two different segments: a PLC segment and a WiFi segment. The achievable throughputs on each of these segment could significantly vary. As a consequence, the end-to-end throughput of each user is throttled by the link segment with the minimum throughput. Having said that, one must be careful when using the WiFi-PLC extenders and not only consider the quality of the WiFi part of communication (which could be easily assessed by the user). A comprehensive understanding of the WiFi links' qualities and their underlying PLC capacities is crucial to ensure that the aggregate throughput of the whole network is improved.

In this dissertation we address some of the issues that could arise in WiFi-PLC networks and propose solutions as follows:

- First, we approach the problem of assigning users to WiFi-PLC extenders. In this problem, we consider an enterprise WiFi-PLC network with plurality of extenders.

Users in this network seek to attach themselves to the extenders to maximize the aggregate throughput. Traditionally, each users will associate with the WiFi-PLC extender that yields the highest Received Signal Strength (RSS). Even though the RSS is an indicator for the WiFi signal quality, the PLC link segment could manifest as a bottleneck of the concatenated WiFi-PLC link. When the PLC link segment is of a lower capacity than what the WiFi link segment can deliver, then the end-to-end throughput for the user is determined by the capacity the PLC link segment can accommodate. As consequence, when users are blindly affiliated with WiFi-PLC extenders or merely based on their RSS, the users' end-to-end throughputs could suffer and the total network throughput degrades. Having said that, we propose a solution to the problem of associating users to WiFi-PLC extenders in enterprise settings. First, we proposed a formal mathematical formulation to the problem. We show that the problem of assigning users to WiFi-PLC extenders is NP-hard by deriving a reduction from the set partitioning problem. Subsequently, we develop a framework, WOLT, that cleverly assigns users to WiFi-PLC extenders with the goal of maximizing the aggregate (network-wide) throughput. It does so in two phases: first, it assigns a single user to each extender in the network such that the PLC-side throughput is maximized. Then, it assigns the rest of the users such that the degradation of the aggregate throughput (that has been achieved in the first phase) is minimized. We evaluate WOLT's performance via both real testbed experiments and high-fidelity simulations and our results show that WOLT can improve the aggregate throughput by up to 2.5x compared to a greedy user association baseline. Our solution as well as

the results are detailed in Chapter 2 of this dissertation.

- Secondly, we consider the problem of weak and unstable WiFi links of WiFi-PLC networks in home settings. Despite the fact that WiFi-PLC extenders can expand the network coverage, users located in areas occluded from the main router’s signal could suffer from well-known WiFi impairments such as fading, shadowing and path loss. Such effects have been shown to be common in indoor environments and they could severely degrade the end-to-end throughput of the users. One of the well-studied solutions to counter for unstable WiFi links is via the usage of Distributed Antenna Systems (DAS). The idea is to simultaneously transmit the data from multiple (distributed) antennas, the signals from which traverse different paths and thus, when combined at the receiver, the likelihood of packet losses is reduced. Therefore, DAS systems are capable of improving both the quality and the stability of WiFi links. Having said that, we seek to exploit the existence of multiple WiFi-PLC extenders in home settings by combining their WiFi transmissions into a DAS system with the goal of enhancing the users’ WiFi links, which in turn, increases the aggregate throughput of the whole network. Toward achieving our goal, we first conduct an extensive measurement study to assess the feasibility of employing DAS on top of PLC network and to quantify the level of synchronization needed by the group of extenders in each DAS cluster. In light of our measurements, we formulate the problem of DAS cell clustering as well as the problems of two fairness models, proportional fairness and max-min fairness. Guided by our insights from the measurement study, we develop PLC-DAS; a framework that groups WiFi-PLC extenders into DAS clusters to max-



imize the aggregate throughput of WiFi-PLC networks. PLC-DAS is evaluated by simulations against three baselines and the results show that the algorithms within PLC-DAS not only improve the aggregate throughput by up to 4.5x compared to blindly using all WiFi-PLC extenders to form a DAS transmitter, but they also result in a fairness level better than or at least comparable to other baselines. The details of this problem formulation and the underpinning algorithms for PLC-DAS is described in details in Chapter 3.

- Lastly, we expand our work in Chapter 3 by considering the same problem but in enterprise settings. When such larger settings are considered, the WiFi-PLC extenders (which act like independent access points) contend for the WiFi frequencies. In such case, a dense deployment of these extenders is expected. However, the number of available WiFi frequencies (channels) are limited and, thus; multiple extenders are expected to share the same frequency. This could lead to a smaller airtime share for each extender contending on the same frequency which could reflect negatively on the achievable WiFi throughput for the users attached to these contending extenders. Moreover, having multiple extenders operating solely (not jointly in a DAS cluster) could increase the number of contending extenders on the PLC backhaul wherein the time is equally shared. As a consequence, each extender will have a time share that commensurate with its PLC capacity as well as the number of extenders sharing the PLC backhaul. DAS, in this case, can help alleviate frequency contention and improve the PLC backhaul sharing (as shown in details in Chapter 4). The idea here is to group multiple extenders into DAS clusters. By doing so, the number of entities

sharing the WiFi domain and the PLC backhaul is reduced to a few number of clusters instead of a large number of independent extenders. Therefore, the limited number of WiFi frequencies as well as the PLC backhaul time is now distributed among a few number of clusters (as opposed to many individual extenders). This results in a better utilization of the network time resource (on WiFi and PLC backhauls) and boosts the throughputs of both WiFi and PLC links. The PLC throughput gains from DAS is first identified in this dissertation. We approach this problem by first formulating the problem and proving that it is NP-hard. As a result, we proposed a measurement-driven algorithm that runs in two stages: (a) it assigns frequencies to a subset of the extenders in the network as to maximize the spatial spectrum efficiency. Then (b) it clusters the remaining extenders into DAS cells with the extenders that had been assigned frequencies in first stage as to maintain the spectral efficiency obtained from the first stage and improve the PLC backhaul sharing. Our algorithm is incorporated in a framework we call Priza. Priza has a polynomial runtime complexity which makes it a feasible solution for WiFi-PLC networks with a large number of extenders. Priza is evaluated against three baselines via both real testbed experiments and high-fidelity simulations. The results from our evaluations show that Priza can increase the aggregate throughput by up to 331.3% over a greedy DAS baseline that creates as large DAS cells as possible and maintain a desirable fairness level among the users even though it does not consider fairness explicitly. The detailed Priza's implementation and evaluations can be found in details in Chapter 4.

## Chapter 2

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

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.

## 2.1 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. 2.1. 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 improve-



Figure 2.1: Powerline communications extend WiFi coverage through existing interior power lines in a home/enterprise.

ment 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.*, [30, 57, 83]) 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 [21, 26]. 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.

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 [89] 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 associating 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.

## 2.2 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 [43]. In brief, CSMA/CA mode arbitrates the communication channel among transmitting nodes within the same contention domain. As discussed in [89], the 1901 CSMA/CA is similar to what is used in 802.11, with some differences in terms of performance, complexity and fairness [91]. 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 [24]. 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 [98]. This makes PLC suitable even for bandwidth intensive applications such as video streaming [56].

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 [8, 9, 23, 47]). 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.




## 2.3 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.

### 2.3.1 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 [47] 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` [7] 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. 2.2a. Initially, we position the two clients at the same location i.e., at equal distance



(a) WiFi only: Throughput-fair medium sharing behavior of WiFi clients. (b) PLC only: Different PLC extenders yield in different medium sharing between active PLC extenders. (c) PLC only: Time-fair medium sharing between active PLC extenders.

Figure 2.2: Medium sharing in the PLC and WiFi domains.

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 further client see a throughput degradation, but the stationary client’s throughput decreases as well, in accordance with the reported performance anomaly [47]. 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 [47].

**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. 2.2b. To see the impact of medium sharing on PLC backhaul throughput, we activated multiple extenders simultaneously, and plot the results in Fig. 2.2c. 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 [75]. 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:

$$\mathcal{T}_j^{\text{WiFi}} = \sum_{i \in N_j} \frac{1}{\sum_{i' \in N_j} \frac{1}{r_{i'j}}} \quad (2.1)$$


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 [23, 57]). 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:

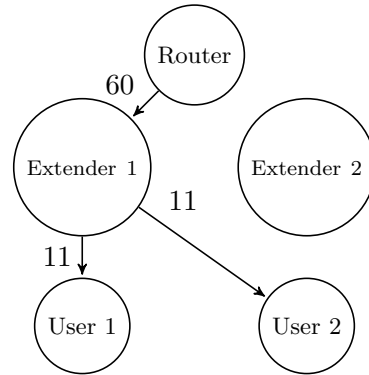
$$\mathcal{T}_j^{\text{PLC}} = \frac{c_j}{A} \quad (2.2)$$

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




(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.

Figure 2.3: Experiments in a scenario observed on our testbed with user association policies.

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 § 2.3.2 (when discussing greedy association).

### 2.3.2 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 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 2.3. Fig. 2.3a 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. 2.3b, we depict this method of user association with an AP; this method has been commonly considered [57] 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. 2.3c. 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. 2.3d, 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.

## 2.4 Problem Statement & Solutions

Our overarching problem is to maximize the total end-to-end-throughput of all users. To do this, in § 2.4.1 we develop a system model based on the take aways from § 2.3, and formulate the problem of maximizing the total throughput. We show that this problem is NP-hard; then, we propose solutions in § 2.4.2 based on certain intuitive properties of the model.

### 2.4.1 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 § 2.3, 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.*, [66, 68, 95]). 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 [24].

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 2.1: Table of Notations



We formulate the throughput maximization problem in Problem 1 below (notations in Table 2.1):

**Problem 1** *PLC-WiFi User Assignment*

$$\max_{x_{ij}} \quad \sum_{j=1}^{|A|} \min \left( \mathcal{T}_j^{WiFi}, \mathcal{T}_j^{PLC} \right) \quad (2.3)$$

$$s.t. \quad \mathcal{T}_j^{PLC} = \frac{c_j}{A}, \quad \forall j \in A \quad (2.4)$$

$$\mathcal{T}_j^{WiFi} = \sum_{i=1}^{|U|} t_{ij}, \quad \forall j \in A \quad (2.5)$$

$$t_{ij} = \left( \frac{1}{\sum_{i' \in N_j} \frac{1}{r_{i'j}}} \right) x_{ij}, \quad \forall j \in A, \forall i \in U \quad (2.6)$$

$$\sum_{j=1}^{|A|} x_{ij} = 1, \quad \forall i \in U \quad (2.7)$$

$$\sum_{i=1}^{|U|} x_{ij} \leq B_j, \quad \forall j \in A \quad (2.8)$$

$$N_j = \{i : x_{ij} > 0\}, \quad j \in A \quad (2.9)$$

$$x_{ij} \in \{0, 1\}, \quad \forall i \in U, \forall j \in A \quad (2.10)$$

The objective (2.3) 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 (2.4) 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 (2.5) specifies the WiFi throughput at extender  $j$ , summed across all users. Constraint (2.6) specifies the throughput of user  $i$  connected to extender  $j$ , based on throughput-fair sharing. Constraint (2.7) postulates that each user must be connected to one extender.

Constraint (2.8) postulates that each extender  $j$  should have no more than  $B_j$  connected users (this constraint will be relaxed later). Constraint (2.9) 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 (2.10) 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 [63], 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 (2.7), 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 is NP-hard, as proved in theorem 1 below. The key idea in the proof is to show a reduction from the partition problem [37] 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  $S = \{w_1, w_2, \dots, w_M\}$  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  $r_{ij} = -\frac{1}{w_i} \forall i = 1, 2, \dots, M$ , the WiFi rates of the dummy users are set as  $r_{ij} = -\infty \forall i = M + 1, M + 2, \dots, M + k$ . 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  $\frac{M+k}{2}$ .

■

## 2.4.2 Solutions for PLC-WiFi User Assignment

---

**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  $i \leftarrow 1$  to  $|U|$  do
2:   for  $j \leftarrow 1$  to  $|A|$  do
3:      $u_{ij} \leftarrow \min(\frac{c_j}{A}, r_{ij})$ 
4:  $(U_1, \{x_{ij}^*\}_{i \in U_1}) \leftarrow \text{ASSIGNMENT\_SOLVER}(\{u_{ij}\}, U, A)$ 
5:  $\{x_{ij}^*\}_{i \in U_2} \leftarrow \text{PROBLEM\_2}(\{r_{ij}\}, \{c_j\}, \{x_{ij}^*\}_{i \in U_1}, U, A)$ 

```

---

In this section, we describe our proposed solutions towards solving Prob. 1. We first provide intuition for our method 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 (2.7), so that not every user has to be connected to an extender, and (b) modifying constraint (2.8) 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 [38] to solve the transformed assignment problem.

Then, in Phase II of our algorithm, we add back constraint (2.7) 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 (2.7) and (2.8), 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 2, which states that a sufficient condition to increase (or maintain) the aggregate throughput is to disconnect the user with worse reciprocal of the WiFi rate than the average of its peers connected to the same extender.

**Lemma 2** *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 (2.3) 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 (2.3) 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:

$$\frac{|N_j|}{\sum_{i' \in N_j} \frac{1}{r_{i'j}}} \geq \frac{|N_j| + 1}{\sum_{i' \in N_j} \frac{1}{r_{i'j}} + \frac{1}{r_{ij}}} \quad (2.11)$$

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 (2.3) increased or stayed the same, proving the claim. A similar analysis follows for the second claim. ■

**Lemma 3** *There exists an optimal solution for the modified Prob. 1 (with (2.7) relaxed, and (2.8) 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 2, 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 2, 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 utility  $u_{ij}$  of user  $i$  to task  $j$  as:

$$u_{ij} \equiv \min\left(\frac{c_j}{A}, r_{ij}\right) \tag{2.12}$$

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 (2.12) 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 3 and (2.12) simplifies to the assignment problem.

**Theorem 2** *The modified Prob. 1 (with (2.7) relaxed, and (2.8) 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 3 that there is at least one optimal solution  $\{x_{ij}^*\}$  such that:

$$\sum_i x_{ij}^* = 1 \quad \forall j \quad (2.13)$$

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 (2.13) as a constraint to Prob. 1. Then, since (2.13)  $\implies |N_j| = 1$ , we can simplify (2.6) as  $t_{ij} = r_{ij}x_{ij}$ , and (2.5) as  $\mathcal{T}_j^{\text{WiFi}} = \sum_i r_{ij}x_{ij}$ . Then the objective function (2.3) 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 (2.13). Thus, Prob. 1 without constraint (2.7) is transformed into an equivalent problem of maximizing  $\sum_j \sum_i \min(r_{ij}, \frac{c_j}{A})x_{ij}$ , with constraints (2.13) and (2.10), 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. 4 below:

**Problem 4** *WiFi User Assignment Only*

$$\max_{x_{ij}, i \in U_2} \sum_{j=1}^{|A|} \mathcal{T}_j^{\text{WiFi}} \quad (2.14)$$

$$s.t. \quad (2.5), (2.6), (2.7), (2.9) \quad (2.15)$$

$$0 \leq x_{ij} \leq 1, \quad \forall i \in U, j \in A \quad (2.16)$$



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 (2.14) maximizes the WiFi throughput  $\mathcal{T}_j^{\text{WiFi}}$  only, compared to the original objective function (2.3) which included the PLC backhaul throughput  $\mathcal{T}_j^{\text{PLC}}$ . 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. 4 allows fractional solutions of  $x_{ij}$ ; however, we prove next in Theorem 3 that Prob. 4 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. 4 results in these integral solutions.

**Theorem 3** *There exists an integer solution to Prob. 4.*

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

$$\sum_{j: x_{ij} > 0} \mathcal{T}_j^{\text{WiFi}}[\text{after}] - \mathcal{T}_j^{\text{WiFi}}[\text{before}] \tag{2.17}$$

$$= \sum_{j: x_{ij} > 0} \left( \frac{x_{ij}}{\sum_{i' \in N_j} \frac{1}{r_{i'j}} + r_{ij}} - \left( \frac{|N_j|}{\sum_{i' \in N_j} \frac{1}{r_{i'j}}} - \frac{|N_j|}{\sum_{i' \in N_j} \frac{1}{r_{i'j}} + r_{ij}} \right) \right) \tag{2.18}$$

where the equality happens after re-arranging the  $\mathcal{T}_j^{\text{WiFi}}$  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 (2.18) 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* [38, 74]. 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  $e^{-5}$ .

## 2.5 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.

### 2.5.1 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 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 [57], 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` [7]. 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 [6] 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 [2]. A  $100\text{ m} \times 100\text{ m}$  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.

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

### **2.5.3 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.

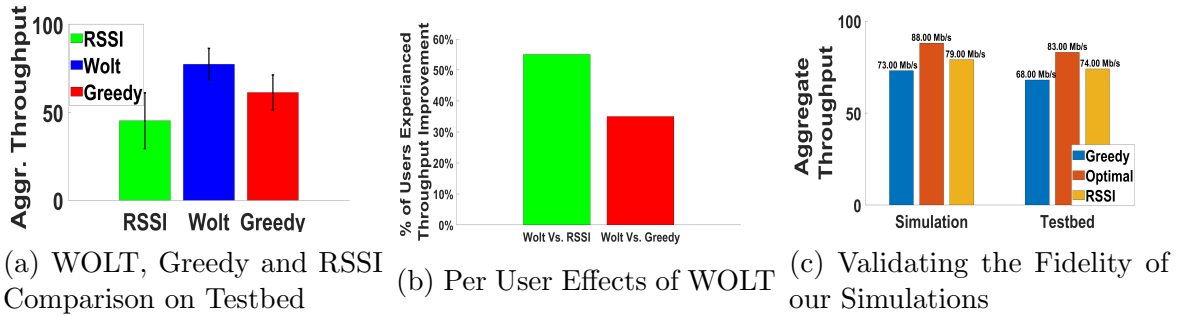


Figure 2.4: Experimental results

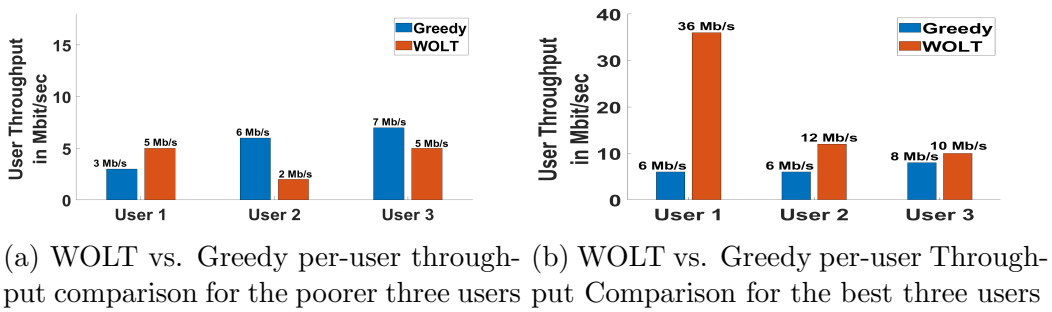


Figure 2.5: WOLT's Effects on Users' Throughputs

## 2.5.4 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 2.4a and 2.4b. 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 association, 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 un-associated (constraint (2.7) in Prob. 9).

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<sup>1</sup>, 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 2.5a and 2.5b, we depict the individual *Greedy* and WOLT throughput

---

<sup>1</sup>Since the set up is small here, we do not consider a fairness metric such as the Jain's fairness index [48] here; we do so later when discussing our simulation results.

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 throughput 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. 2.5a is only (in total) about 6 Mbps. On the other hand, the best three users (depicted in Fig. 2.5b) improve their throughputs 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 2.4c. 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



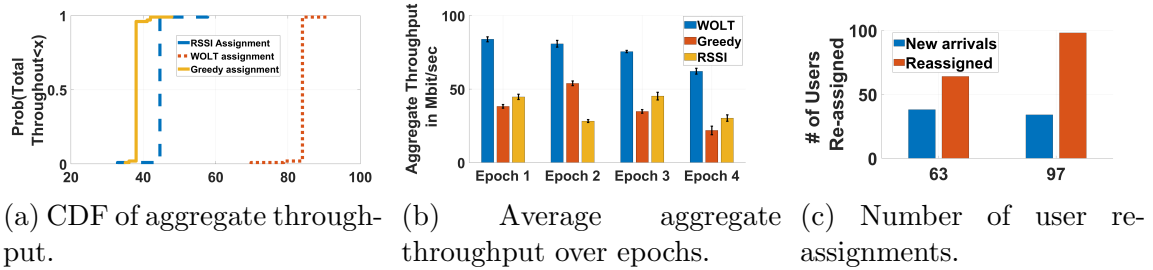


Figure 2.6: Simulation results.

present some larger scale scenarios that we simulate to demonstrate that WOLT performs well even in such cases.

### 2.5.5 Simulation results

**Total throughputs:** We simulate the performance of the WOLT and the greedy algorithm with the simulation settings discussed in §2.5.1. 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. 2.6a. 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. 2.4b, 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 §2.5.1, users arrive and depart from the system according to the Poisson

distribution, with a net average increase of 33 users per epoch. In Fig. 2.6b, 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. 2.6c, 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.

## 2.6 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 [19, 49, 55] 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., [26] propose BOLT a learning-based algorithm to orchestrate flows in a PLC network. Vlachou et al. [88] 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 [91], [89], and [88], 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 [97] 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 [81], 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., [21] perform a measurement study of HomePlug AV2-compliant WiFi extenders. These studies however, do not consider user association problems.

EMPoWER [45] 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 network throughput via intelligent user association. Electri-Fi [90] studies the characteristics of PLC and WiFi networks in terms of their 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 [61,67], extending cellular network coverage with wireless adhoc connectivity is considered. While such networks contain concatenated links, cellular access 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.

## 2.7 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.

## Chapter 3

# Boosting Home WiFi Throughputs via Adaptive DAS Clustering of PLC Extenders

WiFi-capable PLC (Power Line Communications) plug-and-play extenders are becoming popular to improve WiFi range and coverage in homes and enterprises. As shown in prior work, unlike an Ethernet backhaul, the PLC backhaul may not support high data rates. In addition, clients (users) that are either far or partially occluded from the WiFi-PLC extender they associate with can experience fading and shadowing, which degrades the throughput on the wireless link. Thus, both the PLC and WiFi backhauls will influence a user's end-to-end throughput. In this paper, we seek to exploit the presence of multiple PLC extenders that may be plugged in, by combining their transmissions in a distributed antenna system (DAS), to boost client throughputs in a home setting. Specifically, we de-

sign PLC-DAS to determine which PLC extenders are the best candidates for forming a joint DAS transmitter cluster to each client. PLC-DAS is designed based on a real measurement study and not only accounts for the WiFi link qualities from the extenders to the users, but also the PLC link qualities from each extender to a master router which is typically deployed in homes. PLC-DAS is flexible and can maximize the throughput under different fairness objectives. We evaluate PLC-DAS via extensive simulations and show that it can increase the aggregate throughput by up to 4.5x compared to blindly using all WiFi-PLC extenders to form a DAS transmitter, while maintaining a fairness Jain's index value of at least 0.97 with proportional and max-min fairness models.

### 3.1 Introduction

PLC based WiFi extenders that can be plugged into standard power outlets and do not need an Ethernet backhaul are gaining popularity in the market [41] [13] [14]. Typically deployed in homes and enterprises, a master PLC unit connects to the main router and acts as a bridge that connects clients associated with WiFi-PLC extenders to this main router, thus enabling access to the Internet, as shown in Fig. 3.1. It has been shown that WiFi-PLC extenders improve the WiFi coverage in the area of deployment [1]. However, while WiFi-PLC can potentially improve the WiFi coverage, users can still suffer from bad or unstable WiFi links. For example, in indoor deployments, users occluded (no direct link) from WiFi-PLC extenders can see throughput degradation due to deep fading and shadowing. Such effects have been shown to be common in indoor environments [28]. Specifically, a home WiFi-PLC user can experience varying WiFi link qualities across the

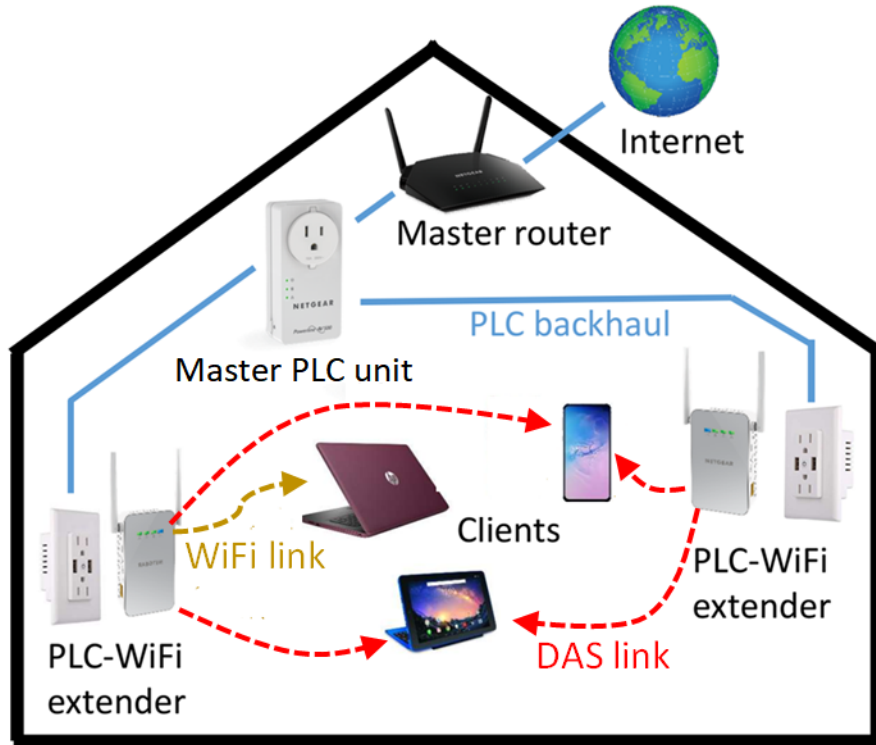


Figure 3.1: Powerline communications extend WiFi coverage through existing interior power lines in a home.

areas in the home (*e.g.*, poor coverage in the garage when all the WiFi-PLC extenders are in various bedrooms inside the home).

A well studied approach to mitigate WiFi link degradation is via the usage of Distributed Antenna Systems (DAS) [100]. The idea is to simultaneously transmit the data from multiple (distributed) antennas, the signals from which traverse different paths and thus, when combined at the receiver reduce the likelihood of packet losses. DAS systems have been shown to improve both the quality and the stability of WiFi links [79].

The key question that we ask in this paper is whether the different PLC extenders can be clustered together to form a DAS transmitter to improve the robustness to indoor fading, and thereby improve the throughputs achievable by clients, regardless of their loca-



tions within a home. However, as discussed in what follows, there are several issues that make the task of forming such a DAS cluster non-trivial because of the composition of the PLC and WiFi parts of the network. Specifically, as shown later, if PLC extenders are blindly clustered to form a DAS transmitter, the client may even suffer a throughput degradation compared to when it uses a single extender. Below, we discuss why this could be the case.

**DAS clustering of extenders could reduce PLC side throughput.** Realizing a DAS cluster [64] [77] requires that a plurality of antennas are synchronized when they transmit the data. While this is inherently satisfied when the antennas are connected to a data source via high-bandwidth backhauls (*e.g.*, fiber), in the scenario of interest, the different PLC extenders are likely to have different backhaul capacities on the links to the master router (from where the packets are delivered to the extenders). Thus, in order to synchronize packet transmissions, the extenders with higher PLC capacities (faster) will need to await the delivery to the slowest PLC extender (the one with the lowest PLC capacity to the master router). In other words, the throughput achieved on the PLC part of the link might in fact decrease compared to what might be achieved on this part, if the client were to connect to a single extender. One must ensure that the gains achieved due to DAS clustering more than offset this decrease; a blind approach to clustering could thus degrade the throughput compared to simply having the client connect to a single extender.

**Large propagation delays from the different extenders can be detrimental to DAS.** Even when the transmitters are synchronized with regards to when they perform their transmissions, there will be variations in propagation delays between the dif-

ferent antennas (WiFi-PLC extenders) and the client. If there is a large difference in these delays, we observe that combining the signals fails, which can adversely affect the throughput. Thus, it is critical that the extenders have similar propagation delays to the client in order to ensure that the gains expected from DAS are derived; otherwise, the throughput might (again) degrade instead of being enhanced.

**Contributions.** In this paper, our key contribution is the design of a framework that we call PLC-DAS to adaptively choose the right set of extenders to transmit to a client, based on its location. This choice is made based on the PLC capacities to the different extenders, and their positions relative to each other and the client (reasons discussed later).

The design of PLC-DAS is driven by a set of experiments that we conduct to understand the above issues. Specifically, we perform extensive WiFi experiments using WARP boards [12] to understand the achievable gains from employing DAS as well as the level of synchronization required to combine the signal from the different antennas. We also conduct experiments on the PLC backhaul using commodity PLC TP Link TL-WPA8630 PLC extenders. The study sheds light on the interaction between the PLC extenders and how the PLC backhaul is shared, which influences the delays from the master router to the various extenders.

As our main contribution, we design a framework, PLC-DAS, that incorporates a measurement-driven online algorithm to determine the right set of antennas to transmit to a client which is at a given location. Specifically, the algorithm results in the client choosing a primary extender to associate with, based on the PLC capacities from the master router to the various potential extenders, as well as the wireless link qualities to the extenders.

Next, other extenders are chosen for joint transmissions along with this primary extender, based on their PLC link capacities and their propagation delays to the client (relative to the primary extender), to form a DAS cell.

We perform extensive evaluations of PLC-DAS using realistic PLC and wireless channel models derived from experiments with realistic home configurations. We consider various fairness models and show that PLC-DAS not only boosts individual client throughputs, and therefore the overall network throughput that is shared across clients in a home, but can also provide max-min or proportional fairness. This is achieved by reducing discrepancies across client throughputs due to better robustness to fading effects, compared to baselines that do not use DAS or apply DAS blindly.

A summary of our contributions in this paper are:

- We perform real experiments to gain an understanding of the feasibility of employing DAS on top of a PLC network.
- We leverage the insights obtained from the measurement experiments to design PLC-DAS. We show that the algorithms within PLC-DAS, which drive the choice of the appropriate antennas to form the DAS cluster to maximize the throughput (or fair throughput) for the client, have polynomial time complexity and can be practically deployed.
- We perform extensive simulations based on realistic channel models and real home layouts to show that PLC-DAS outperforms other baselines approaches that dictate how clients associate with extenders (without DAS or via a blind application of DAS). PLC-DAS achieves up to a 62.7% increase in aggregate throughput compared to a non-

DAS baseline in which each user associates with the PLC extender that offers the best end-to-end throughput; this is the best baseline in terms of the achieved aggregate throughput. The results also show that PLC-DAS provides better fairness across users that share the in-home WiFi capacity, with both max-min and proportional fairness models.

### 3.2 PLC Background in Brief

The MAC 1901 protocol governs backhaul access in PLC networks. It is similar to 802.11, with some differences in terms of the complexity, fairness and performance [89]. It can be configured to operate using a CSMA (throughput-fair) or a TDMA (time-fair) mode. It supports different QoS classes by granting the flows with higher priority a larger number of time slots in the TDMA mode. Most large vendors such as Cisco, Netgear and TP-Link, support both medium access modes with a PHY rate up to 2024 Mbit/sec [98], which makes PLC extenders attractive for expanding the network without needing pre-existing infrastructure.

Most current PLC extenders are empowered by a WiFi interface that increases the network range. This is especially attractive in areas where the main router's signal is low or poor, causing lowered data rates. The WiFi link between the PLC extender and the end user (also referred to as a client) is controlled by the 802.11 protocol. Since 802.11 shares the medium in throughput-fair manner, users connected to the same WiFi-PLC extender will have similar throughputs, and extenders operating on the same WiFi channel will have to share the frequency associated with that channel [16].

### 3.3 Measurements

In this section, we describe our experiments on real testbeds, to get an understanding of the issues relating to realizing DAS in the PLC-WiFi home setting. Specifically, our measurements relate to three aspects: (a) first, we seek to understand the variations in the rate across various PLC links; (b) second, we seek to quantify the gains that might be expected with DAS without using any precoding [59] [78] (note that precoding cannot be applied since we are using plug and play WiFi extenders as our antennas); and (c) finally, we seek to understand the extent to which time synchronization needed across the multiplicity of signals for them to be combined successfully at the receiver.

#### 3.3.1 Feasibility of PLC as a backhaul for DAS

While in general, the PLC backhaul is time-shared across the various extenders, when a DAS transmission is to be enabled, the same data is multicast to the extenders that belong to the DAS cluster, *i.e.*, there is no time-sharing of the backhaul for that transmission. The feasibility of multicasting on the PLC backhaul has been previously demonstrated [35,69]. By creating multicast groups – one group for each DAS cell – we can send the data simultaneously to multiple extenders. However, the latency incurred by each packet transmission to the plurality of PLC extenders in the multicast group varies. In Fig. 3.2 we show, from our real measurement study, the distribution of these latencies across twenty extenders in a two-bedroom apartment in the USA. The setup is as follows: for each power outlet, we include a PLC extender and laptop, which connects to a master router connected to a server, as shown in Fig. 3.1. Specifically, the laptop client is connected by

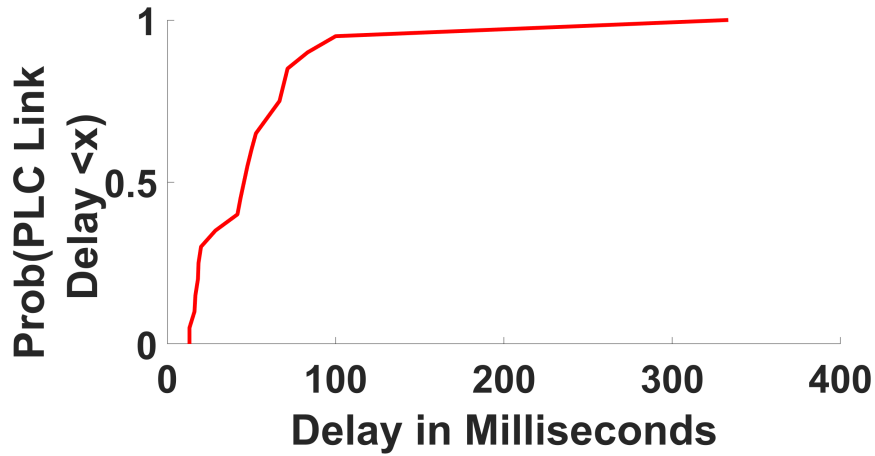


Figure 3.2: CDF of latencies experienced on different PLC links.

a Gigabit Ethernet cable to the PLC extender, for this experiment only (the high capacity of the cable ensures that any degradation in user throughput is because of limited PLC link capacities, not because of the Ethernet connection). The extender is connected over PLC backhaul to the main PLC unit, which in turn is connected to the master router and server over Ethernet. We transmit a saturated downlink TCP flow from the server to laptop clients. We repeat the experiment with twenty different outlets (different PLC links) and we use *iperf3* [7] to measure the throughput/delay on each PLC link.

The measurement results suggest that transmissions on different PLC links experience different delays. Because of this, the delay (and consequently throughput on the PLC network) incurred by a DAS transmission is governed by the delay of the slowest extender in its multicast group. In fact, the difference in the delays as seen by the experiments could be as much as  $3\times$  between pairs of extenders. We reiterate here that the transmissions from the extenders need to have fine-grained synchronization in order to enable a successful DAS transmission; thus, those extenders which receive the packet earlier, will need to await the laggards prior to performing the transmission (no such wait is necessary when there is a

single extender performing the WiFi transmission). In other words, the PLC capacity for a DAS transmitter is dictated by the capacity of the lowest PLC link in its multicast group and is given by:

$$p_i = \min_{\{j \in A: x_{ij} > 0\}} c_j \quad (3.1)$$

where  $p_i$  is the DAS transmitter PLC link capacity for user  $i$ ,  $j$  is the index of extender in the multicast group (DAS cell),  $c_j$  is the capacity of the extender,  $A$  is the set of extenders in that DAS cell, and  $x_{ij}$  is a binary variable indicating whether user  $i$  is connected to extender  $j$  in the DAS cell. This showcases the importance of carefully choosing PLC extenders when forming DAS transmitter; blindly grouping or choosing all extenders could cause the end-to-end throughput of the users to degrade due to this artifact on the PLC side. Later, in our simulation experiments to evaluate PLC-DAS, we emulate the depicted distribution of PLC link delays from Fig. 3.2.

### 3.3.2 DAS side issues

Next, we implement DAS and conduct experiments on the WARP [12] platform to quantify DAS' gains in terms of SNR improvement. In a nutshell, we find that this gain is logarithmic with the number of antennas (extenders) as we show later in this section. We also examine how DAS performs with different transmission powers; specifically, when transmission powers change (increase or decrease), the receiving node experiences different received powers, and we investigate whether this has an impact on DAS gains or not. This


emulates different proximities of a client to the DAS transmitters.

**Experimental setup:** We first describe the set up for our experiments.

*DAS with two antennas.* We use two WARP V3 nodes, one of which acts as the transmitting node (Tx node) and one as the receiving node (Rx node). The Tx node has two SMA output ports. Each SMA port has one antenna attached to it (total two antennas). The Rx node has only one antenna. No precoding [59] [78] is applied as discussed earlier to reflect scenarios with off-the-shelf PLC extenders. We run our experiments with twenty different topologies in which we change both the Rx (client) and Tx antennas' locations. We examined the benefits of constructive signal combining of DAS at the Rx node, with varying Tx powers (from 10 dB to 15 dB). With each Tx power level, we send one hundred transmissions. In order to ensure that the reported average SNR covers a wide range of values, we send two thousand OFDM symbols with each transmission, which is the maximum number of symbols the WARP node can buffer [72], encoded with BPSK modulation. The Rx node captures the superposed transmitted signal from the two transmit antennas and attempts to decode the received combined signal. When the decoding process is successful, the payload is retrieved. After that, the average SNR is determined by computing the Error Vector Magnitude (a.k.a Relative Constellation Error or RCE).

*DAS with more antennas.* To construct DAS clusters with more antennas, we use "Y" shaped splitters to increase the number of Tx antennas. Each splitter has two ends. The first end is attached to the Tx node and the other end is used to connect two antennas. We connect one splitter to each SMA output port (there are two of them) and, consequently, increase the number of Tx antennas to up to four. The Rx node has only one antenna. As





(a) DAS gain as a function of Tx power when 1-2 Tx antennas are used. (b) DAS gain when 2, 3, or 4 antennas are used. (c) The effect of non-synchronized transmissions for two Tx antennas.

Figure 3.3: DAS experiments showing feasibility of PLC as a backhaul for DAS, DAS performance gains, and DAS synchronization issues.

with the two antenna case, we run our experiment with twenty different topologies in which we change the locations of both the Tx and Rx (client) antennas. With each change in the locations of the Tx and Rx antennas, we perform one hundred transmissions. The average SNR value across all the runs and the different topologies is then computed.

*Results on gains with DAS:* Our experiments show that DAS with two antennas, on average, provides a 3 dB increase in the signal-to-noise ratio (SNR) when two antennas are used. This is found to be true across a range of transmission powers as shown in Fig. 3.3a.

The result in Fig. 3.3b shows a logarithmic increase in SNR value at the Rx node as the number of the Tx antennas increases. A 3 dB increase is observed with two antennas. With three antennas, the total DAS gain is about 4.75 dB. The total DAS gain with four antennas is 6 db. Specifically, the results show that the the average resulting SNR due to DAS can be modeled by:

$$w_i = 10 \log\left(\sum_j^{|A|} 10^{snr_j^i/10}\right) \quad (3.2)$$

where  $w_i$  is the resulting average SNR from DAS,  $i$  refers to the index of a specific user, and  $j$  is the index of the antenna.  $A$  is the set of Tx antennas and  $snr_j^i$  is the SNR value user  $i$  experiences from antenna  $j$  alone. We find that [29] reports the same observations as we do here.

**Synchronization:** Our final experiment seeks to quantify the level of synchronization needed across a set of transmitter antennas in DAS, in order to guarantee a constructive signal combining at the receiver. Here, we use one Tx node with two antennas. Then, we induce delays at one of Tx antennas (prior to transmission) at the granularity of nanoseconds to see how this impacts the received SNR at a Rx node. Specifically, the signal is modulated using BPSK and stored in a buffer corresponding to the Tx antennas. Then, we stagger the transmissions of one of the Tx antennas to induce differences in times when the signals are received by the Rx node. The two transmitted signals mix and superpose in the air before arriving to the Rx node. The Rx node receives the mixed signal and starts decoding it. Once the decoding process succeeds, the average SNR is computed.

The result of our experiments, shown in Fig. 3.3c, suggest that if the difference in transmission times between two signals is equal to or larger than 600 ns (nanoseconds), the SNR starts to sharply decline. This happens because the cyclic prefix serves as a guard interval against inter-symbol interference (ISI). In our experiment, the cyclic prefix of each OFDM symbol is equal to 600 ns. Once the time difference between the two signals exceeds the length of the cyclic prefix, ISI is more likely to be severe. This result suggests that extenders that are chosen to serve in one DAS cell must tightly synchronize their transmission times to less than 600 ns.

## 3.4 Problem Statement & Solutions

Our goal is to maximize the aggregate throughput (with different types of fairness objectives) of WiFi-PLC users in a home setting. In order to do this, in §3.4.1 we propose a system model based on the insights from §3.3. Then we formulate the problem of maximizing the total utility (discussed later) of the WiFi-PLC network. We decompose this problem into two subproblems, DAS cell formation and WiFi time assignment, and propose an algorithm in § 3.4.2 to solve these. Our solution can optimize throughput with respect to different fairness functions (specifically, max-min fairness and proportional fairness).

### 3.4.1 Problem Statement

The network consists of a PLC backhaul with a WiFi air interface. Each user connects to the master router over a concatenated PLC-WiFi link. A group of PLC links can deliver data to more than one PLC extender on the PLC backhaul. We refer to such a grouping of PLC extenders as a DAS cell or a DAS cluster. Since we consider a home network in this work, we assume a single WiFi contention domain (multiple interfering contention domains such as in enterprises is left for future work). Therefore, there is minimal inter-domain interference, and each DAS cell serves a single user at a time, by simultaneously transmitting the same data over the WiFi interface to the end user. Multicast is used to efficiently deliver the data to all PLC extenders in each DAS cell, rather than inefficiently sending the same data via unicast to each extender in that group [35, 69].

Our objective is to maximize the total network utility, where utility is defined as a function of the throughput. We formulate this optimization problem in Problem 5 below.

The notations used in these formulations are summarized in Table 3.1; note that in this paper, *rate* refers to the PHY bit rate of the WiFi or PLC links, while *throughput* refers to the achieved bit rate a user could enjoy on a PLC or WiFi link given a time allocation.

Variable	Description
$\alpha$	The fairness factor.
$A$	Set of PLC-DAS extenders.
$c_j$	The PLC capacity of extender $j$
$f(\cdot)$	A function that takes the SNR value in dB and returns the corresponding WiFi modulation scheme rate
$\gamma_j$	The delay difference between the primary extender and extender $j$
$\lambda_i$	The WiFi time allocation for user $i$
$p_i$	The capacity of the PLC backhaul for user $i$
$snr_j^i$	The SNR value experienced by user $i$ from extender $j$
$U$	Set of users.
$v_i$	bitrate of user $i$ .
$u(\cdot)$	utility function defined in (3.11)
$w_i$	The cumulative DAS-SNR for user $i$
$x_{ij}$	Binary variable indicating whether extender $j$ serves user $i$ .

Table 3.1: Table of Notations

**Problem 5 Overall Formulation**

$$\max_{x_{ij}, \lambda_i} \sum_{i=1}^{|U|} u(v_i \lambda_i) \quad (3.3)$$

$$s.t. \quad v_i = \min(f(w_i), p_i), \quad \forall i \in |U| \quad (3.4)$$

$$w_i = 10 \log \left( \sum_{j=1}^{|A|} (10^{(snr_j^i/10)}) x_{ij} \gamma_j \right) \quad (3.5)$$

$$p_i = \min_{\{j \in A: x_{ij} > 0\}} c_j, \quad \forall i \in |U| \quad (3.6)$$

$$\sum_{j=1}^{|A|} x_{ij} \geq 1, \quad \forall i \in |U| \quad (3.7)$$

$$\sum_{i=1}^{|U|} \lambda_i = 1 \quad (3.8)$$

$$x_{ij} = \begin{cases} 1 & \text{if extender } j \\ & \text{serves user } i, \forall i \in |U|, \forall j \in |A| \\ 0 & \text{otherwise} \end{cases} \quad (3.9)$$

$$\gamma_j = \begin{cases} 1 & \text{if extender } j \text{ and the} \\ & \text{primary extender are} \\ & \text{out of the sync by } >, \forall j \in |A| \\ 600 \text{ ns} \\ 0 & \text{otherwise} \end{cases} \quad (3.10)$$

The objective (3.3) is to maximize the total utility of all users. The utility function  $u$  is defined as the  $\alpha$ -fair utility function as a function of each user's throughput  $v_i \lambda_i$ :

$$u(v_i \lambda_i) = \begin{cases} \frac{(v_i \lambda_i)^{1-\alpha}}{1-\alpha} & \alpha > 0, \alpha \neq 1 \\ \log(v_i \lambda_i) & \alpha = 1 \end{cases} \quad \forall i \in |U| \quad (3.11)$$


where for  $\alpha = 0$ , the system prioritizes efficiency (*i.e.*, the aggregate throughput), and as  $\alpha$  increases, the system prioritizes fairness (*i.e.*, users have an equal share of the throughput).

Constraint (3.4) says that the achievable end-to-end bitrate of user  $i$  is the minimum of its PLC and WiFi link segments. Constraint (3.5) specifies the aggregate WiFi SNR a user receives from all the extenders in its DAS cell. Constraint (3.6) quantifies the capacity of


the PLC backhaul for user  $i$  as the minimum of the extenders to which it is connected (for DAS), due to multicast. Constraint (3.7) states that each user must connect to at least one extender. Constraint (3.8) states that the total time allocation must sum to 1. Constraint (3.9) says that  $x_{ij}$  is a decision variable which is equal to 1 when user  $i$  is connected to extender  $j$ , and 0 otherwise. Finally, constraint (3.10) is a system parameter that describes whether the extenders are synchronized within less than 600 ns. The variables in this optimization problem are  $\lambda_i$ , the WiFi time allocation for user  $i$ , and  $x_{ij}$ , which specifies whether user  $i$  is connected to extender  $j$ .

**Toy example:** To illustrate our problem, we next describe a toy example. We will show that creating DAS cells naively, such as by associating to the extender that gives the highest RSSI or to the extender that offers best end-to-end throughput, may not result in the the optimal solution, and thus solving our optimization problem is non-trivial. Fig 3.4a shows our example network topology with the possible PLC and WiFi links for user 1 and user 2. The edges between the router and the two extenders represent the bitrate of each PLC link, if each PLC link was used in isolation. The edges between the extenders and the two users represent the WiFi links if only one WiFi link was active.


First, consider the case for user 1 in Fig. 3.4b. When the two extenders form a DAS cell, user 1 will enjoy a WiFi link with bitrate of 48 Mbps, and an overall end-to-end throughput of  $\min(48, 40) = 40$  Mbps. If user 1 naively decides to associate with extender 1 alone, because it gives the best end-to-end throughput (36 Mbps), using DAS still yields a higher throughput. DAS gives a higher throughput because (a) the WiFi signal from the two extenders combined is better than what user 1 can achieve with an extender individually,




(a) Maximum achievable rates of each PLC and WiFi links in isolation.



(b) Maximum achievable throughput (rate) for user 1. The resulting DAS-SNR is 17 dB which maps to 48 Mbps. Total throughput is  $\min(48, 40) = 40$  Mbps.



(c) Maximum achievable throughput (rate) for user 2. The resulting DAS-SNR is 28 dB for user 2. Total throughput is  $\min(60, 60) = 60$  Mbps. Total throughput is  $\min(72, 40) = 40$  Mbps.



(d) Maximum achievable throughput (rate) for user 2. The resulting DAS-SNR is 28 dB for user 2. Total throughput is  $\min(60, 60) = 60$  Mbps. Total throughput is  $\min(72, 40) = 40$  Mbps.

Figure 3.4: Toy example of different possible DAS cell formation solutions.

and (b) this boost in WiFi signal more than compensates for the reduced rate on the PLC backhaul due to multicast.

On the other hand, user 2 will suffer if both extenders are naively used to form a DAS cell, as shown in Fig. 3.4c. The reason is that user 2 will experience a hit in throughput from including extender 2, because the PLC link for extender 2 is poor quality (40 Mbps), so adding it decreases the multicast backhaul rate to 40 Mbps, throttling the end-to-end throughput to 40 Mbps as well. A naive solution for user 2 is to connect to the extender that offers the highest RSSI, which is extender 2. However, this assignment is suboptimal since the PLC link segment (extender 2's PLC link) has a capacity of only 40 mbs, throttling the end-to-end throughput of user 2 regardless of the high quality of its WiFi link. The optimal end-to-end throughput for user 2 is achieved through the configuration shown in Fig. 3.4d, where user 2 connects to extender 1 only.

**Problem Decomposition:** We next describe how to decompose problem 5 into two sub-problems. First we formulate the problem of DAS antenna selection (*i.e.*, DAS cell formation), where we solve Problem 6 for  $x_{ij}$ . We refer to this as Problem 6. We then formulate the problem of WiFi time allocation, where we solve Problem 7 for  $\lambda_i$ . We first describe Problem 6 below, which is defined for each user  $i$ :

**Problem 6** *PLC-DAS Extender Selection*

$$\max_{x_{ij}} \quad \min(f(w_i), p_i) \tag{3.12}$$

$$s.t. \quad w_i = 10 \log \left( \sum_{j=1}^{|A|} (10^{(snr_j^i/10)}) x_{ij} \gamma_j \right), \quad \forall i \in |U| \tag{3.13}$$



$$p_i = \min_{\{j \in A: x_{ij} > 0\}} c_j, \quad \forall i \in |U| \quad (3.14)$$

$$\sum_{j=1}^{|A|} x_{ij} \geq 1, \quad \forall i \in |U| \quad (3.15)$$

$$x_{ij} = \begin{cases} 1 & \text{if extender } j \text{ serves user } i \\ 0 & \text{otherwise} \end{cases} \quad (3.16)$$

$$\gamma_j = \begin{cases} 1 & \text{if extender } j \text{ and the primary} \\ & \text{extender are out of sync by } > \\ & 600 \text{ ns} \\ 0 & \text{otherwise} \end{cases} \quad (3.17)$$

The objective (3.12) in Problem 6 says that we want to maximize the throughput of a given user  $i$  (during a given time duration). The constraints in this problem match those in problem (5) relating to  $x_{ij}$ .

Next we formulate the problem of solving for the time allocations  $\{\lambda_i\}$  in Problem 7 below:

**Problem 7** *WiFi time allocation with  $\alpha$ -fairness*

$$\max_{\lambda_i} \sum_{i=1}^{|U|} u(v_i \lambda_i) \quad (3.18)$$

$$s.t. \quad \sum_{i=1}^{|U|} \lambda_i = 1 \quad (3.19)$$

The utility function  $u(\cdot)$  in (3.18) is defined in (3.11), and the objective is to maximize the

summation of user utility with respect to the time allocations  $\lambda_i$ . Constraint (3.19) states that the total time allocation across users is equal to one. Note here that this formulation is general enough to capture a wide spectrum of fairness definitions, depending on the value of  $\alpha$ .

In the special case of  $\alpha = 1$  in (3.11), we have the proportional fair utility function as the objective, as written below in Problem 8:

**Problem 8** *WiFi time allocation with proportional fairness ( $\alpha = 1$ )*

$$\max_{\lambda_i} \sum_{i=1}^{|U|} \log(v_i \lambda_i) \quad (3.20)$$

$$s.t. \quad \sum_{i=1}^{|U|} \lambda_i = 1 \quad (3.21)$$

The constraints of Problem 8 are the same as Problem 7.

We seek to understand whether the solutions to the decomposed problems also solve the overall problem. Theorem 4 below shows this.

**Theorem 4** *A solution to Prob. 6 and Prob. 8 is also a solution to Prob. 5 when  $\alpha = 1$ .*

**Proof.** Denote a solution to Prob. 5 as  $(\lambda_i^{**}, x_{ij}^{**})$ . Denote a solution to Prob. 6 as  $x_{ij}^*$  and a solution to Prob. 8 as  $\lambda_i^*$ . The constraints of Prob. 5 are equal to the union of the constraints of Probs. 6 and 8, and hence their feasible sets are equivalent. It remains to examine their objective functions. The claim is that  $(\lambda_i^*, x_{ij}^*)$  is also a solution to Prob. 5.

Since  $x_{ij}^*$  maximizes  $v_i(x_{ij})$  according to the definition of Prob. 6, we know that  $v_i(x_{ij}^*) \geq v_i(x_{ij}^{**})$  for all  $i$ , which implies that  $\sum_i \log(\lambda_i^* v_i(x_{ij}^*)) \geq \sum_i \log(\lambda_i^{**} v_i(x_{ij}^{**}))$ . The

RHS maximizes (3.3) for  $\alpha = 1$ , and so the LHS must equal the RHS. Next, since  $\lambda_i^*$  maximizes Prob. 8, we have that  $\sum_i \log(\lambda_i^* v_i(x_{ij}^*)) \geq \sum_i \log(\lambda_i^{**} v_i(x_{ij}^*)) = \sum_i \log(\lambda_i^{**} x_{ij}^{**})$ , where the equality comes from the RHS-LHS argument above. Since  $\sum_i \log(\lambda_i^{**} x_{ij}^{**})$  maximizes (3.3), we can replace the first inequality with equality. Therefore,  $(\lambda_i^*, x_{ij}^*)$  achieves the same optimal value in (3.3) as  $(\lambda_i^{**}, x_{ij}^{**})$ , and hence is also a solution to Prob. 5.

■

A similar proof holds for Prob. 5 and Probs. 7, 6.

### 3.4.2 Algorithms for DAS Cell Formation and WiFi Time Allocation

In this section, we describe our algorithms to solve Problems 6, 7, and 8. This is done in two steps: (1) DAS cell formation via antenna selection, and (2) WiFi time allocation. The first step of DAS cell formation is solved via Algorithm 2. Algorithm 2 assigns each user to the extender (primary extender) that gives the highest end-to-end throughput. Then, it adds additional extenders to create a DAS cell for each user. Algorithms 3 and 4 allocate time to each user to achieve max-min fairness or proportional fairness, respectively.

**Step I:** First, Algorithm 2 iterates over all users  $i \in U$  and extenders  $j \in A$ , and finds the extender that gives the best end-to-end throughput for each user (lines 1 to 6). Then, it checks if there is any other extenders that can be added to create a DAS cell for each user (lines 7 to 15). It does so by first checking if the PLC capacity of the extender to be added ( $x_{ij}$ ) is greater than the bitrate the user currently has ( $v_i$ ) (line 11). Second, it checks if adding that extender will result an improvement in the WiFi link (lines 13 and 14).

---

**Algorithm 2** DAS Cell Formation

---

**Inputs:** Set of users  $U$ , set of extenders  $A$ , PLC capacity  $c_j$ , SNR value  $snr_j^i$ ,

**Output:** User assignments  $x_{ij}$

**Variables:** user index  $i$ , extender index  $j$ , end-to-end bitrate  $b_{ij}$

```
1: for  $i \leftarrow 1$  to  $|U|$  do
2:   for  $j \leftarrow 1$  to  $|A|$  do
3:      $b_{ij} \leftarrow \max(\min(f(snr_j^i), c_j))$ 
4:    $\bar{j} \leftarrow \text{argmax}_j(b_{ij})$ 
5:    $x_{i\bar{j}} \leftarrow 1$ 
6:    $v_i = b_{i\bar{j}}$ 
7: for  $i \leftarrow 1$  to  $|U|$  do
8:   for  $j \leftarrow 1$  to  $|A|$  do
9:     if  $\gamma_j == 1$  then
10:      if  $x_{ij} == 0$  then
11:        if  $c_j > v_i$  then
12:           $x_{ij} \leftarrow 1$ 
13:           $w_i \leftarrow 10 \log \left( \sum_{j'=1}^{|A|} (10^{\frac{snr_{j'}^i}{10}}) x_{ij'} \right)$ 
14:           $v_i \leftarrow \min(f(w_i), p_i)$ 
15:        else  $x_{ij} \leftarrow 0$ 
```

---

**Step II:** After determining the DAS cells, we next determine how to allocate time resources fairly across users. Our methods to achieve max-min and proportional fairness are presented in Algorithms 3 and 4, respectively. In the former, we try to maximize the throughput of the user with the minimum throughput in the system. This can be achieved by granting more airtime to the users with poor throughputs. Therefore, users are allocated time based on their achievable bitrates relative to the maximum achievable bitrate across all users. Such an allocation can easily be shown to result in equal throughputs for all users.

---

**Algorithm 3** Max-Min Fairness

---

**Inputs:** Set of achievable users' throughputs  $\{v_i\}$

**Output:** Max-min time allocation for each user  $\{\lambda_i\}$

**Variables:** user index  $i$

```
1: for  $i \leftarrow 1$  to  $|U|$  do
2:    $\lambda_i = \frac{\max_{j'}(v_{j'})}{v_i}$ 
```

---

Algorithm 4 attempts to solve Problem 8, for proportional fairness. The objective (3.20) aims to maximize the summation of the log utility,  $\log(v_i \lambda_i)$ . Since  $x_{ij}$  is fixed in Problem 8 (it was computed in Step I above), Problem 8 has a closed form solution for  $\lambda_i$ , where all the users have equal time shares (*i.e.*, equal  $\lambda_i$ ) [33]. Consequently, we divide the total time 1 by the total number of users in the system,  $|U|$ .

---

**Algorithm 4** Proportional Fairness

---

**Inputs:** Set of achievable users' bitrates  $\{v_i\}$

**Output:** Proportional fair time allocation for each user  $\{\lambda_i\}$

**Variables:** user index  $i$

- 1: **for**  $i \leftarrow 1$  to  $|U|$  **do**
  - 2:      $\lambda_i = \frac{1}{|U|}$
- 

**Algorithm Complexity:** Step I of our algorithm runs in  $\mathcal{O}(|U||A|^2)$ , where  $|U|$  is the number of users and  $|A|$  is the number of the PLC extenders. The runtime of the Step II of our algorithm is  $\mathcal{O}(|U|)$ . Thus the total runtime for both steps is given by  $\mathcal{O}(|U||A|^2)$ .

## 3.5 Evaluations

To capture a diverse variety of home settings, we perform extensive simulations. We make the simulation set ups realistic by using the experimental measurement results reported in Section 3.3, both on the PLC and the WiFi (DAS) parts of the network.


### 3.5.1 Simulation Details

**PLC-DAS Simulation Framework:** We implement PLC-DAS entirely from scratch in MATLAB [71] since other simulating tools do not have models of realistic PLC backhauls. PLC-DAS runs on a Lenovo T460p machine running a 64-bit Windows 10 oper-

ating system. We develop code to simulate 2D homes with different areas. This helps us understand how PLC-DAS behaves in a variety of settings, spanning a small studio apartment to an average home in the USA [3]. Specifically, we simulate four to eight WiFi-PLC extenders and five to twenty users. Each extender is placed randomly in the home area and assigned a capacity from the distribution we observed in our measurements in Section §3.3. Later we show that our home models yield results that are similar to the results when real home layouts and WiFi-PLC extender locations from Pinterest [5] are considered.

The number of users (5-20) and their locations in the house layout are chosen randomly. The WiFi links between users and extenders are assigned SNR values based on the physical distance between each user relative to each extender, as well as the shadowing and fading impacts for indoor users as reported in [62]. The resulting average SNR values that users could experience when using DAS is then computed based on our findings in Section §3.3. Subsequently, each averaged SNR value is mapped to a modulation scheme based on the SNR-to-modulation translations provided in [39]. Each modulation scheme is capable of encoding a specific number of bits within each OFDM symbol. Therefore, each SNR value is mapped to a specific bitrate that the corresponding modulation scheme can provide [87].

**Baselines:** We evaluate PLC-DAS against three baselines: (a) Best End-To-End (BETE), (b) Received Signal Strength (RSS) and (c) All-Extenders (All-EXT). The BETE baseline assigns users to the single extender that provides the highest end-to-end throughput over the concatenated WiFi-PLC link. The RSS baseline assigns a user to the extender with the highest quality WiFi link. without considering the PLC backhaul capacity. This reflects



(a) Aggregate throughput with proportional fairness. (b) Aggregate throughput with max-min fairness. (c) Aggregate throughput with throughput fairness.


Figure 3.5: CDF of aggregate throughput with different fairness models.

the assignment policy that currently exists on off-the-shelf WiFi-PLC extenders. Lastly, we consider the case when all accessible extenders are used to form a DAS cell (All-EXT). The All-EXT baseline demonstrates the pitfalls when DAS cells created blindly without considering PLC link capacity differences.

**Performance Metrics:** Our metrics of interest are: (a) the aggregate network throughput that PLC-DAS can deliver compared to the other baselines and, (b) the fairness achieved with PLC-DAS versus other baselines, with respect to our three fairness models viz., proportional fairness (PF), max-min fairness (MM) and throughput-fair (TF). As discussed in §4.6, PF allocates time slots equally to each user, and MM seeks to maximize the throughput of the user with the minimum throughput. higher rates.

### 3.5.2 Results

**Throughput gains with PLC-DAS:** First we compare PLC-DAS against the three baselines, in terms of aggregate throughput with the different fairness models. The CDFs in Figs. 3.5a, 3.5b and 3.5c show that PLC-DAS outperforms all the baselines in all trials. In Fig. 3.6a we show the aggregate throughputs when using PF. *We find that PLC-DAS with*



(a) Proportional fairness aggregate throughputs. (b) Max-min fairness aggregate throughputs. (c) Throughput-fair aggregate throughputs.

Figure 3.6: Average aggregate throughput. The results show how PLC-DAS outperforms other baselines.

*PF outperforms all the three baselines and yields average throughput improvements of 58%, 112% and 462%, over BETE, RSS and All-EXT, respectively. Similarly, when MM and TF throughputs are maximized, as seen in Fig. 3.6b and Fig. 3.6c respectively, PLC-DAS outperforms the baselines BETE, RSS and ALL-EXT on average by 62.7%, 103% and 457%, respectively.*

The RSS baseline yields a higher throughput under PF compared to MM (by 22%). This is because PF allocates an equal time shares to all users. On the other hand, MM maximizes the throughput of the user with the worst throughput, *i.e.*, it provides larger time allocations to users with poor rates, compared to users with good rates, and consequently suffers from a lower aggregate throughput.

The All-EXT baseline suffers with all fairness models. When all extenders are considered for a DAS cell, the extender with the poorest PLC capacity (the slowest), becomes the bottleneck that limits the throughput for that cell. The poorest extender will always be the last extender to receive data on the PLC backhaul and other extenders have to wait for it before transmitting the data over WiFi to the end user, thus increasing the delay and decreasing the throughput.

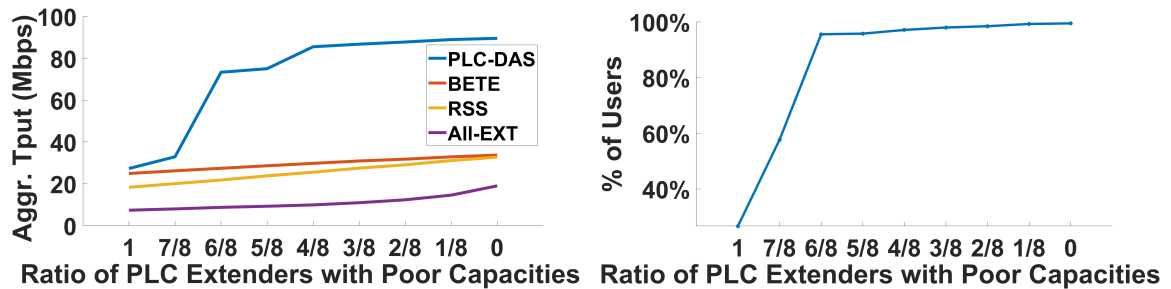


	<b>PF</b>	<b>MM</b>	<b>TF</b>
PLC-DAS	0.97	1	1
BETE	0.93	1	1
RSSI	0.78	1	1
ALL-Exts	1	1	1

Table 3.2: Jain’s fairness index

**Fairness:** We use Jain’s fairness index [48] to evaluate the fairness of PLC-DAS in comparison to the three aforementioned baselines. The closer Jain’s index value is to 1, the more fair. *The Jain’s fairness index results in Table 3.2 show that PLC-DAS provides a better or at least comparable fairness to all other baselines, except for All-EXT. Unlike other baselines, PLC-DAS shows a balance between fairness and maximizing the aggregate throughput.* While PLC-DAS might not be able to benefit some users because either (a) they are not in locations that are amenable to DAS, or (b) because they are already obtaining the best throughput at their locations by connecting to a single primary extender, PLC-DAS can improve the individual throughputs of users with poor throughputs. This alleviates variations across user throughputs thereby enhancing the fairness index.

The reason that All-EXT exhibits the highest fairness index is because all extenders a user can associate with are included for the DAS cluster corresponding to that user. The extender with the poorest PLC link capacity in the home will now dominate the backhaul throughput, and the throughputs of all users will degrade to the capacity determined by that poor extender. Consequently, the individual throughputs for these users become similar and are poor. This boosts the Jain’s index for All-EXT, but at the expense of severely decreased user throughputs. MM and TF always yield the highest index of 1. This is because these models try to achieve equal throughputs for all users. This improves fairness but again, at



(a) The effect of PLC backhaul on the total throughput of PLC-DAS (b) The effect of the PLC backhaul on the percentage of users benefiting from PLC-DAS.

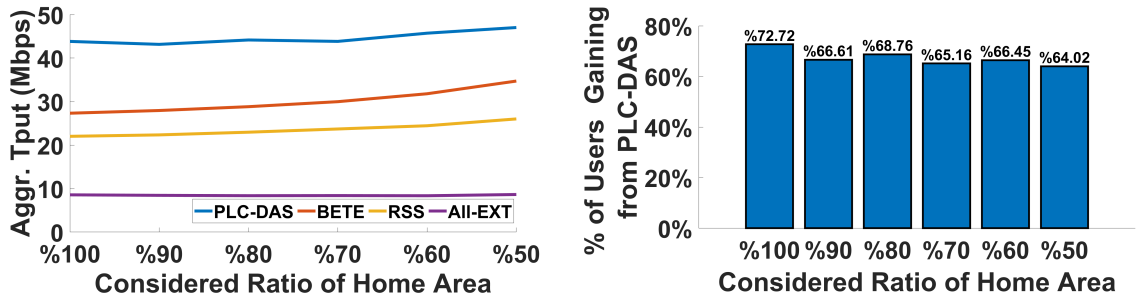
Figure 3.7: PLC backhaul impact.

the cost of higher throughputs. Importantly, PLC-DAS maximizes the aggregate throughput while scoring a fairness index of at least 0.97 across different fairness models.

The remaining results in this section reflect the PF model unless stated otherwise. We omit the results with the other two models (results are similar in spirit with what was reported thus far) due to space constraints.

**The impact of the PLC backhaul on gains with PLC-DAS:** Next, we examine how the different PLC link qualities affect the aggregate throughput gain of PLC-DAS. We simulate a WiFi-PLC network in a home setting with eight PLC extenders, all with good links (a PLC link is classified as good if its capacity  $> 50$  Mbps [21]). Subsequently, we flip one of the good PLC links to a poor link (capacity  $\leq 50$  Mbps). We continue this process, *i.e.*, keep flipping good PLC links, one at a time until all the PLC links are poor. At each switch from a good to a poor link, we simulate the experiment one thousand times and then we take the average of the aggregate throughputs.

The results of this experiment, captured in Fig. 3.7a, show that PLC-DAS still offers an improvement (albeit small) in terms of aggregate throughput even when PLC



(a) Home area effect on the total throughput (b) Home area effect on the percent of users when using PLC-DAS benefiting from PLC-DAS.

Figure 3.8: Home area impact.

backhaul is all poor (ratio=1). PLC-DAS shows an improvement of 1.9 times compared to BETE and up to 7.4 times over All-EXT when 75% of the PLC links are poor (ratio=6/8). This demonstrates that PLC-DAS is very effective in improving throughput even with a mostly poor PLC backhaul. This is because some users can exploit the good PLC extenders to form DAS clusters. Beyond this point however, as the ratio of poor to good quality PLC links increases, the penalties incurred due to PLC backhaul links causes the overall throughput to drop sharply; almost no client benefits from using DAS in such cases.

**Impact of the home area (large house vs. small studio):** In this next set of experiments, we consider homes with sizes similar to an average US house. The goal is to understand the gains with PLC-DAS in homes of different areas. In each case, we consider four to eight extenders with five to twenty users, and randomly generate one thousand different topologies with regards to user locations. The results in Fig. 3.8a show that when the home area is decreased by 50%, PLC-DAS, BETE, and RSS improve their aggregate throughputs compared to the original larger home area. In contrast, the fraction of users benefiting from PLC-DAS when the home area is decreased by 50% declines by

8.7% ,compared to when the whole home area is considered (ratio=100%), as shown in Fig. 3.8b. This is driven by users who are now connected to their primary PLC extender with excellent WiFi link segments, because they are now closer due to the reduced home area. Consequently, *the gains from PLC-DAS are reduced in smaller homes; larger houses, however, yield much higher throughput gains.*

**Realism of our home layouts:** Finally, we show that the randomly generated house layouts and extender locations realistically reflect real-life home settings. We consider ten real house layouts with electrical diagrams obtained from [5]. We study the electrical wiring on these layouts and extract the physical locations of the power outlets, with up to 20 power outlets observed per home based on the electrical diagrams. Each power outlet is a potential location where a WiFi-PLC extender can be plugged in. We run our simulation with ten real-life house layouts. We simulate four to eight extenders and five to twenty users. Extenders are assigned capacities as per the distribution observed in Section §3.3. Then we compute the average aggregate throughputs and the fairness indices. We repeat the experiment, but this time with the house layouts and extender locations generated randomly with our approach. *Tables 3.3 and 3.4 show that our randomly generated topologies yield very similar results to real house layouts and extender locations.* This provides confidence that our simulations reflect realistic home settings.

	Aggregate Throughputs with PF (Mbps)	
	Real House Layouts	Randomly Generated Layouts
PLC-DAS	40	42.1
BETE	25.5	26.7
RSS	21	21.7
All-EXT	9.3	7.5

Table 3.3: Aggregate throughput comparison between real house layouts and randomly generated layouts

	Jain’s Fairness Index with PF	
	Real House Layouts	Randomly Generated Layouts
PLC-DAS	0.97	0.97
BETE	0.95	0.94
RSS	0.85	0.84
All-EXT	1	1

Table 3.4: Fairness comparison between real house layouts and randomly generated layouts

### 3.6 Related Work

In this section, we describe relevant related work.

**WiFi User Association:** The authors of [19,49] develop user association policies in WiFi networks to optimize a performance metric (*e.g.*, throughput) with some fairness models. Our work differs for multiple reasons. First, these efforts assume an Ethernet backhaul of a higher capacity than the WiFi links. The PLC backhaul may not satisfy this property, *i.e.*, the PLC segment could be the bottleneck rather than the wireless link. Second, they do not consider the use of DAS.

**Distributed Antenna Systems:** There are several efforts on creating DAS clusters, either to improve robustness or energy efficiency (*e.g.*, [32, 34, 51, 76]). These efforts

however, do not consider the impact of the backhaul on DAS transmissions as we do here.

**PLC:** There are efforts like [31, 40, 65, 102] on broadband PLC networks. They do not consider concatenated WiFi-PLC links.

**Hybrid WiFi-PLC:** Vidyut [97] considers using electrical wiring to deliver a synchronization signals to wireless APs to improve performance in MIMO deployments. However, they are not concerned with the characteristics of the PLC backhaul. Apichart-trisorn *et al.* [21] measure the performance of PLC extenders equipped with HomePlug AV2, which are WiFi-compliant extenders. In [81], the authors study if a PLC backhaul can serve as a backbone for WiFi in home settings. None of these studies however, consider DAS deployment. The authors of [16] propose a framework to assign users to the appropriate extenders with the objective of maximizing the aggregate throughput in hybrid WiFi-PLC networks. However, they do not consider using DAS as a mechanism for providing better indoor coverage.

### 3.7 Conclusions

In this paper we propose PLC-DAS, a framework to maximize the total network throughput in WiFi-capable PLC networks by using distributed antenna systems or DAS. PLC-DAS intelligently chooses the best set of WiFi-PLC extenders for each user, that gives the highest end-to-end throughput for that use. The challenge we handle is that we ensure that we eliminate PLC backhaul links of inferior qualities when creating the DAS cluster; otherwise, we show that this can degrade the user throughput instead of improving the same. We formulate the problem of choosing the set of extenders that yields the highest

throughput and propose variations that account for different fairness models. Our problem formulation is driven by real experiments on PLC and DAS testbeds. We subsequently design a set of algorithms to build a framework PLC-DAS, which significantly boosts the achievable throughputs of users within homes. We show via simulations that PLC-DAS significantly outperforms non-DAS and naive DAS baselines in terms of aggregate throughput while scoring high in terms of the Jain's fairness index (at least 0.97).

## Chapter 4

# Priza: Throughput-efficient DAS

## Clustering of WiFi-PLC Extenders

### in Enterprises

WiFi-enabled Power Line Communications (PLC) range extenders use electrical wiring as their infrastructure and can extend coverage in homes and enterprises. However, a dense deployment of a large number of PLC extenders in enterprise settings, can cause an inefficient sharing of the PLC capacity, wherein many extenders contend for a share of access to the PLC backhaul, thereby drastically impacting any gains from using these extenders. In this paper, we seek to address this issue by developing a framework, Priza, for clustering the WiFi-PLC extenders to intelligently form DAS (distributed antenna system) to mitigate the inefficiency of sharing on the PLC backhaul. By appropriately managing clustering and reuse, Priza improves the PLC backhaul sharing, while at the same



time harnesses power pooling and diversity gains provided by DAS on the wireless part of the network, to boost user throughputs. We evaluate Priza via real testbed experiments and high-fidelity simulations and demonstrate that it can increase the aggregate throughput by up to 131.5% over the non-DAS reuse baseline, 74% over a DAS baseline that constructs equally-sized DAS cells based on extender proximity, and 331.3% over a greedy DAS baseline that creates as large DAS cells as possible.

## 4.1 Introduction

Powerline communication (PLC) extenders offer a viable technology to expand network coverage in homes and enterprises without the need for buttressing the underlying wired infrastructure [1]. PLC extenders are plug-and-play devices that have recently gained popularity [14, 20]. They can be plugged into power outlets to facilitate networking over electrical wiring. Via standard electrical wall outlets, they communicate with a central controller over existing electrical wiring which, in turn, connects to a master router. The router’s role is to connect the local network to the Internet (e.g., via coaxial cable or fiber). The central controller relays packets between the PLC extenders and the router, and regulates access to the PLC backhaul (from the extenders). Users (clients) can connect to the PLC extenders via Ethernet cables or wirelessly. In other words, an extender plugged into a nearby outlet mimics an AP and can offer a user in an area occluded from the master router, good signal quality, thereby potentially delivering higher throughput than via a direct wireless connection to the master router.

Since there are no guidelines on plugging in extenders into outlets, in an enterprise

setting, one can envision a large number of users plugging in such extenders to improve their wireless throughputs (they are fairly cheap with a price ranging from \$50 to \$120 for a pair of extenders [50]). In cases where there is a dense deployment of such extenders (e.g., in enterprises with closely packed office spaces), the PLC shares of the extenders can shrink, causing the PLC backhaul to become a bottleneck, thus negating the gains from the signal strength benefits on the wireless side due to shorter wireless links. Specifically, the PLC backhaul capacity is time-shared by the extenders as reported in [22] [17] [18]. Thus, when a plurality of WiFi-PLC extenders are active simultaneously, access to the PLC backhaul is equally time shared among these extenders. Thus, as the number of extenders increase, each extender gets a much smaller time share, and thus, experiences a shrinkage in its PLC throughput. This especially affects extenders that have poorer PLC links, for which, the PLC part becomes the bottleneck in terms of the achievable throughput. This in turn, in many cases can completely neutralize the gains from better signal strength due to the closer extender (because the PLC backhaul becomes the bottleneck).

One way to counter the aforementioned problem is to have the extenders “cooperate” by performing joint transmissions in a distributed antenna system (DAS) configuration. Traditionally, DAS systems which entail synchronized transmissions from a cluster of antennas, have been proposed for alleviating wireless channel impairments i.e., fading. With DAS, the transmissions of multiple antennas are constructively combined at a receiver to reduce the likelihood of packet loss and WiFi link instability. DAS for such purposes, has been extensively studied previously (e.g., [100]). Our vision is to group contending PLC extenders into DAS clusters, to improve the sharing of the PLC backhaul capacity.

Specifically, with our approach, the number of competing entities for the PLC capacity, will now commensurate with the much smaller number of DAS clusters rather than the large number of extenders. We note that the additional advantages of WiFi link stability, and robustness to fading that are inherent to DAS will contribute to improving the achievable overall network capacity as well.

However, there are three key factors that make the task of creating DAS clusters on top of WiFi-PLC networks in enterprise settings a non-trivial task. First, one should be careful not to significantly compromise reuse by creating very large DAS clusters; if clusters are formed in an uninformed way, this can cause a degradation in throughput compared to when no DAS is utilized due to poorly utilized frequency bands. Second, because the PLC capacities to the different extenders could themselves vary in capacity [22] [17] (i.e., some extenders could have good PLC links while others could have bad ones), naively clustering PLC extenders without considering their PLC capacities could result in a degradation of users' throughputs compared to the case when the user is attached to an individual extender. Third, if one were to group extenders with very diverse wireless propagation delays to a receiver, DAS combining can be compromised, leading to reception failures.

In this paper, we first perform an extensive PLC/WiFi measurement study using NI USRP-N210 radios [11] and an OctoClock CDA-2990 [4] and TP Link TL-WPA8630 WiFi-PLC extenders. This study not only helps understand the aforementioned factors, but also sheds light on how the PLC capacities are geographically distributed; specifically, we see that PLC capacities can not be inferred simply based on relative extender locations.

Guided by the understanding gained by our measurement study, we design a frame-

work that we call Priza (electric outlet in Greek), to adaptively assign WiFi frequencies to WiFi-PLC extenders and subsequently to group them into DAS cells with the goal of increasing the aggregate throughput. In brief, Priza first assigns exclusive frequencies (interfering extenders are not assigned the same frequency) to WiFi-PLC extenders with an objective of maximizing reuse towards retaining the wireless part of the capacity; note that it is possible not all extenders get an exclusive frequency. Priza then seeks to group the extenders that were unable to obtain an exclusive frequency with those that did, to form DAS clusters; this in turn, not only boosts users' WiFi link robustness as intended by DAS, but reduces the number of entities sharing the PLC backhaul (multiple extenders are grouped into a fewer DAS transmitters), thereby boosting the PLC time share for the PLC backhaul contenders. An informed DAS cluster construction can drastically reduce inefficient sharing of this PLC backhaul capacity<sup>1</sup>. The key property of Priza is that it strikes a balance between exploiting the available frequencies and the usage of DAS to effectively mitigate PLC inefficiencies.

Importantly, the frequency assignments and clustering decisions are made based on the associated PLC capacities to the extenders. In the second step above, Priza checks if clustering a pair (or group) of extenders to form a DAS transmitter boosts or hurts throughput compared to those extenders sharing the WiFi channel. As discussed earlier, the latter case is possible if the PLC capacities to the two extenders under considerations vary significantly (i.e., one has a high PLC backhaul capacity while the other does not). With Priza, the DAS cells constructed in a way such that (a) WiFi-PLC extenders with high PLC capacity discrepancies are not grouped into the same cluster and (b) distant WiFi-PLC

---

<sup>1</sup>Later in §4.7.4 we show that creating DAS cells first and then assigning frequencies may lead to construction of DAS cells that compromise reuse, and thus lead to a reduced system capacity.

extenders with significantly different propagation delays to the user do not perform a joint DAS transmission. The measurement-driven algorithm within Priza needs to iterate over the WiFi-PLC extenders once before it converges, but runs in polynomial time. We clarify here that both the frequency assignment and DAS cell construction processes are performed while maintaining the initial user associations to extenders, i.e., no user re-assignment is needed.

A summary of our contributions in this paper are:

- We perform extensive measurements to gain an understanding of how the PLC backhaul operates and to quantify the gains that can be expected from DAS. Our measurement study sheds light on the factors that influence whether DAS clustering can indeed provide throughput gains in dense enterprise settings.
- Guided by the understanding from the measurement experiments, as our primary contribution, we design Priza. We show that the algorithm within Priza, which drives the frequency assignments and DAS clustering decisions, has an associated polynomial time complexity. We fully implement the DAS part of Priza, and emulate the PLC part based on our measurements with real extenders to conduct a realistic deployment study.
- We evaluate Priza on a real testbed via comparisons with three baselines; we demonstrate that it is capable of achieving a 33.7% higher throughput compared to the state of the art reuse baseline that does not employ DAS, 56.5% over a baseline that creates equally-sized DAS clusters simply based on extender proximity, and 144.6% over the baseline that blindly creates as large DAS clusters as possible.

- We also evaluate Priza at scale, by using high-fidelity simulations (which we show conform with our experimental results in small settings). The results from our simulations show that Priza can outperform other baselines to even larger extents (upto 331 % over the worst baseline and 131 % over the reuse baseline) than our experiments, due to its inherent ability to cope with the increasingly diverse PLC capacities that arise with scale.

## 4.2 PLC Background in brief

WiFi-PLC extenders use the 1901 MAC protocol (called MAC 1901 hereon) to regulate access to the PLC backhaul. This protocol is similar to the WiFi 802.11 protocol, but there are differences between the two in terms of the fairness that they offer, their complexity and performance [89]. MAC 1901 is configurable to operate either in a CSMA or a TDMA mode, and can manage flows with different priorities i.e., it thus supports different QoS classes. As reported in [98], modern commercial PLC extenders manufactured by large vendors such as TP-Link, Cisco and Netgear can support up to 2024 Mbit/sec in terms of the PHY layer rate, which makes PLC extenders attractive alternatives for conventional range extenders that use a pre-existing Ethernet backhaul. The ability to support high PHY rates also makes these extenders a suitable solution for online gaming and video streaming.

Most current PLC extenders are WiFi-capable, which expands the network converge even more, especially in areas where an AP or wireless router's signal is weak or unstable (which causes lowered data rates). The WiFi interface is governed by the 802.11 protocol, and each WiFi-PLC extender acts as an independent access point. Therefore,

users that are connected to a WiFi-PLC extender share the medium in a throughput-fair manner. On the other hand, extenders sharing the same frequency will have to effectively share access to that frequency [17]. Thus, if there is a dense deployment of extenders in a region, the contention for each available frequency would be high, resulting in a much smaller airtime share for the extenders contending on the same congested frequency (shown later).

### 4.3 Factors influencing PLC based DAS clusters

Next, we discuss some of the factors that influence how PLC extenders should be grouped into DAS clusters (we had briefly alluded to these in § 4.1).

**How we choose extenders for DAS clusters impacts the system capacity.**

Dense extender deployments, can cause the extenders to inefficiently share the PLC backhaul, causing significant degradations in the achievable throughput. As discussed earlier, DAS can mitigate this problem (we are the first to identify the benefit of DAS in mitigating PLC backhaul contention) by grouping several extenders to form a combined transceiver. However, if extenders are grouped into clusters blindly, reuse opportunities could be lost, and thus, result in a wastage/inefficient usage of the system capacity. To showcase this with a trivial example, when there are  $N$  antennas (extenders) and an equal number of available WiFi frequencies, if all of the antennas are grouped into a DAS cluster, only one transmission is performed at a time on one of these frequencies and the remaining  $N - 1$  frequencies remain unused. In contrary, if the antennas operate separately without DAS (i.e., in what is called the reuse mode [84]), then the total number of simultaneous transmissions can be

*N.* Even though DAS improves PLC efficiency, and offers gains due to power pooling (from the various antennas) and diversity, these may not compensate for the loss induced by the reduction in the number of simultaneous transmissions and thus, the network throughput will decrease. Therefore, a naive DAS clustering (e.g., creation of one large DAS cell), could be ill conceived. Determining how DAS and reuse should be combined is a critical challenge that we address in our design.

**PLC capacities need to be accounted for, when forming DAS clusters.**

Traditionally, a DAS setup requires that all antennas be connected to a data source over a high speed network (e.g., fiber or Ethernet). However, PLC extenders use electrical wiring which is not only of lower capacity but importantly, different PLC links tend to be very diverse in terms of their capacities [18]. Thus, when a downlink transfer of a data packet from the master router to a wireless client is considered, this packet is typically delivered at different times (via the PLC backhaul) to the extenders that are grouped to form the DAS cluster. Since the extenders (aka antennas) will need to perform a joint transmission, those antennas that are connected to a higher speed PLC link will need to await those that have associated slower PLC links to receive their data, before commencing their synchronized transmissions. Thus, the PLC capacity of the poorest PLC extender in the DAS cluster will dictate the PLC throughput, and including extenders with poor capacities with those with good PLC capacities, could offset the gains expected from the usage of DAS (due to throttling on the PLC part of the link). This necessitates a careful selection of which PLC extenders are to be grouped in each DAS cell.

**A large variance in propagation delays between transmitting PLC ex-**



**tenders is detrimental to DAS.** A successful DAS signal combining depends not only on the synchronized transmission of the antennas (WiFi-PLC extenders), but also on the propagation delays from the multiple transmitters to the receiver. If there are large variations in the signals' travel times, then the combining fails. This has been reported previously by us in [18], where we experimented with simple DAS scenarios using the WARP platform [12] and show that if the delay is greater than 600 ns, the signals cannot be combined. Our platform here is different in that we use USRP radios (as discussed in § 4.1), and we find that with this platform, slightly higher delays can be tolerated; however, to be conservative, we preclude including two extenders into a DAS cluster if the signal propagation delays from those extenders to the client vary by more than 600 ns.

## 4.4 Measurement Study

In this section, we first ask if simply clustering nearby extenders can suffice in forming effective DAS clusters. To this end, we performed measurements which show that power outlets that are in close proximity of each other could in fact have significantly different PLC capacities. These discrepancies in PLC capacities suggest that such a simple strategy will not work. Second, we showcase a set of measurements that help us understand the factors that could influence DAS clustering of WiFi-PLC extenders in enterprise settings.

Before we present the results of our measurement study, we describe our setups, and how we perform our experiments.

**Experimental setup:** Our experimental setup consists of two parts: (a) a PLC setup and (b) a DAS setup.

*PLC setup.* In our PLC setup, we use two TP Link TL-WPA8630 PLC extenders, a Netgear R7000-100NAR Nighthawk router and two laptops viz., an Acer Aspire E15 and a Lenovo Ideapad 300S-14ISK. The first PLC extender acts as the central controller that interfaces with the second extender over the PLC backhaul and with the router via an Ethernet cable. By using Ethernet cables, we connect one of the laptops to the master router and the other laptop to the second PLC extender. The first extender’s role is to relay the traffic between the first laptop (the one connected to the router) and the second laptop over the PLC backhaul. We plug in the second PLC extender (the one that has a laptop connected to it with the Ethernet cable) into various power outlets distributed in four university labs with cubicles, desks and research equipment as shown in Fig. 4.1.

*DAS setup.* For the DAS setup, we use six NI USRP-N210 radios [11] and an OctoClock CDA-2990 [4]. The radios are connected to the clock over SMA cables. The clock’s role is to synchronize the internal clocks of the radios, i.e., when a group of radios is set to form a DAS transmitter, the signal from these antennas are fired within  $50ns$  of each other. This is important to ensure successful signal combining at the receiver [18]. All the radios are equipped with a CBX-40 USRP daughterboard. An antenna is attached to each radio via a SMA port. Three of the six radios are designated as transmitting antennas ( $Tx = \{Tx_1, Tx_2, Tx_3\}$ ) and the remaining three are the users or receivers ( $U = \{U_1, U_2, U_3\}$ ). Both the Tx antennas and the user antennas are connected to a switch over Ethernet cables. We place  $Tx_1$ ,  $Tx_2$  and  $Tx_3$  next to the power outlets in the area surrounded by a solid line in Lab 3 as to reflect the case of real WiFi-PLC extenders’ locations. User antennas (U) are placed as shown in the same area. A Lenovo T460p running Windows 10 64 bit is used

as the central server. The central server is connected to the same switch over an Ethernet cable and its role is to synthesize the signal and pass it to the Tx radios. The Tx radios transmit the signal which is picked up by the users' antennas. The received signal is then sent back to the server for decoding. Upon a successful decoding of the signal, the server computes the achievable WiFi throughput for each user.

*Synergizing PLC and DAS connections.* Ideally, we would connect the USRP radios to the PLC backhaul by using TP Link TL-WPA8630 PLC extenders. However, these radios are designed to communicate with Labview (a software developed by National Instruments to control various devices including their N210 USRP radios) over Ethernet cables. In fact, it is explicitly stated by National Instruments in [10], that the USRP radios will not work unless they are connected to a Gigabit Ethernet interface. In other words, Labview assumes the Ethernet links to the radios to have a certain capacity. When the PLC links are introduced between the radios and the server which is running Labview, the communication between the radios and Labview is either lost or significantly corrupted. This is because Labview streams the signal to the radios at a certain speed (assuming Ethernet capacity). Since the signal will traverse the PLC backhaul which is often of much lower capacity than Ethernet, the PLC link experiences an overflow and packets at the sender buffer are dropped. In spite of several attempts, we found this to be a problem with many of our PLC links, especially those of poor capacity. In order to bypass this issue, we estimate the PLC link capacities offline and then use these estimates to emulate delays experienced on the PLC backhaul in Labview itself. Each of the emulated PLC capacities used in our measurement study are derived from the corresponding real PLC capacities,

estimated from the corresponding power outlets which are all from our four university labs as shown in Fig. 4.1.

#### 4.4.1 Distribution of PLC capacities

The PLC backhaul consists of multiple electrical wires that connect WiFi-PLC extenders to the central controller. The PLC link capacities on these wires differ from outlet to outlet. The differences in PLC capacities can be caused by the length of the wire, amount of noise generated by appliances operating on the same wire, electrical impedance or the number of branches stemming from each wire [103] [22]. Unfortunately, these characteristics, and their impact on the end-to-end capacity of each PLC link is often not easily available and can be considered opaque. However, one might ask if these factors influence nearby PLC-extendors similarly. The objective of this experiment is to understand if this is the case, and if nearby WiFi-PLC extendors can be grouped into a common DAS cluster, trivially.

When a DAS cluster is to be formed, all the extendors are treated as one multicast group [18], and the data source sends the data on the PLC backhaul at once, to all extendors in the associated DAS cluster. However, the data is delivered to the different extendors in a cluster at different times due to the discrepancies in the PLC capacities to those extendors. Thus, the extendors with the good PLC links will have to await laggard extendors with poor PLC links to receive the data on the PLC backhaul before performing their synchronized DAS transmissions. As a consequence, if extendors with poor PLC capacities are clustered with extendors that have good PLC links, the backhaul throughput of the whole cluster degrades to the throughput of the poorest PLC extender in that cell.

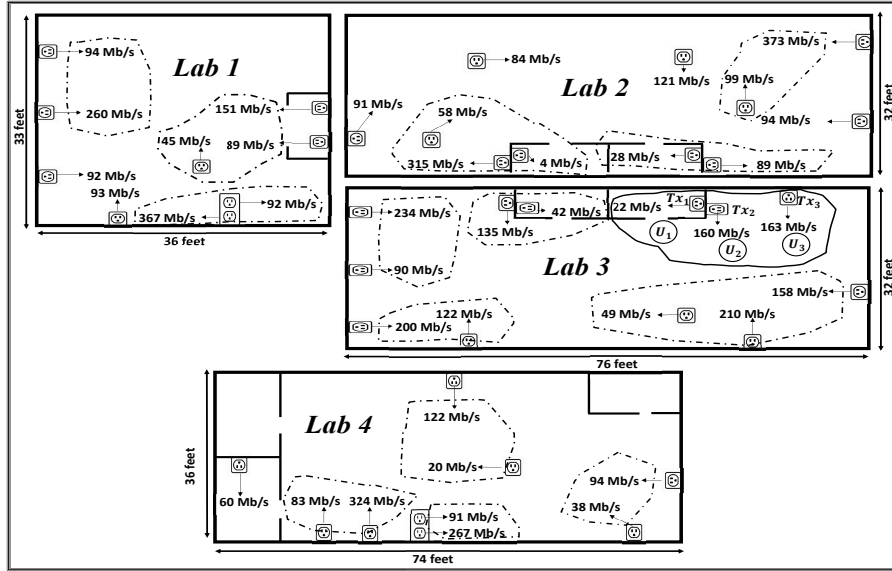


Figure 4.1: PLC capacities in different labs in our university setting.

We use our PLC setup (discussed earlier in this section) to understand the feasibility of naively clustering WiFi-PLC extenders based on their geographical proximity. Specifically, we initiate saturated TCP traffic between the two laptops using *iperf3* [7]. *iperf3* is configured to send the TCP traffic for 30 seconds from the first laptop (the one connected to the router) to the second laptop (the one connected to the target PLC extender). Thus, the traffic traverses the Ethernet cable between the first laptop and the router and then traverses the PLC backhaul to the second laptop. Since the PLC backhaul is of a lower capacity than the Ethernet cable [17] [18] [22], the achievable throughput reported by *iperf3* is determined by the PLC capacity (it is the bottleneck of the concatenated Ethernet and PLC link). We perform this experiment on various power outlets in four university labs.


Our results show that PLC capacities are not consistent across outlets in close proximity i.e., PLC links with poor capacities can exist in close proximity of PLC links

with good capacities as shown in Fig. 4.1 (e.g., one can see the outlet with a PLC capacity of 315 Mbps right next to one with 4 Mbps in Lab 2). This leads us to conclude that one must be careful not to cluster PLC extenders based only on their locations. Grouping extenders naively based on proximity could cause both extenders with good and poor capacities to belong to the same cluster, which as discussed earlier, would cause a degradation of throughput for the entire cluster to that of the extenders with the poor PLC links.


#### 4.4.2 The need for informed WiFi-PLC DAS clustering

As shown earlier in this section, the PLC capacities cannot be inferred merely based on the locations of the power outlets. This observations emphasizes the importance of making informed PLC backhaul decisions when clustering WiFi-PLC extenders together. While as shown previously [84], DAS can indeed improve the system capacity and WiFi link robustness, blindly grouping WiFi-PLC extenders together can lead to a poor network throughput. Specifically, we show with simple experiments that (a) DAS can enable better sharing of the PLC link capacities in addition to improving the WiFi link qualities and robustness, and thus the overall capacity but, (b) arbitrarily grouping extenders into a DAS cluster may result in a degradation in the network throughput compared to the case when DAS is not used.


We use the DAS setup (discussed earlier in this section) to showcase two connection modes as shown in Fig. 4.2. The numbers next to the Tx antennas represent the PLC throughputs associated with the corresponding antennas and the numbers next to the users represent the WiFi throughput that the users can achieve in that configuration. First, if



(a) The network setup in reuse mode.  $U_1$  and  $U_2$  could achieve WiFi data rate of 24 Mbps in isolation. Since  $Tx_1$  and  $Tx_2$  share the same frequency,  $U_1$  and  $U_2$ 's WiFi throughputs become 12 Mbps each.



(b) The network setup with an informed DAS clustering.  $U_2$  and  $U_2$  could achieve WiFi data rate of 36 Mbps in isolation. Since both users share the same cluster, their WiFi throughputs become 18 Mbps each.



(c) The network setup with a bad DAS clustering.  $U_1$  and  $U_2$  could achieve WiFi data rate of 36 Mbps in isolation. Since both users share the same cluster, their WiFi throughputs become 18 Mbps each.

Figure 4.2: Snapshots of the area surrounded by solid line in lab 3 from Fig.1 with three different network setups.

DAS is not used i.e., all extenders use the reuse mode only, we consider the extenders shown in Fig. 4.2a from Lab 3, to show how dense settings could negatively impact the system capacity and lower the aggregate throughput. Subsequently, we consider the usage of DAS. When DAS is used, we investigate two different potential DAS configurations: (a) when  $Tx_1$  and  $Tx_2$  are grouped into one DAS cluster (Fig. 4.2c), we experience the pitfall of grouping extenders with varying PLC capacity together despite the improvement DAS can deliver on the wireless channel, and (b) when  $Tx_2$  and  $Tx_3$  are grouped together (Fig. 4.2b), we see how DAS can help increase the aggregate throughput if performed in an informed way. Recall here that the users are initially associated with their primary Tx antennas and thus, no user re-assignments are performed. That is, when an antenna (extender) is set to join a cluster, the user associated with that antenna becomes a part of that cluster. The initial user assignments are as follows: user 1 is assigned to  $Tx_1$ , user 2 is assigned to  $Tx_2$  and user 3 is assigned to  $Tx_3$ . We limit the number of available frequencies to only two frequencies for the ease of showcasing our take aways.

**DAS can help overcome inefficient PLC sharing caused by dense deployments of extenders and improve throughput.** In this network setup we demonstrate the impact of having multiple extenders sharing the PLC backhaul, when they all operate independently (reuse mode as shown in case (a) in Fig. 4.2 . Specifically, we assign  $Tx_1$  and  $Tx_2$  the same frequency and,  $Tx_3$  a different frequency (recall that we have only two frequencies). We set all the three antennas to the reuse mode; thus, the network contains three reuse cells as shown in Fig. 4.2a, two of which share the same WiFi frequency. In this network configuration, all users  $U$  achieve a WiFi data rate of 24 Mbps in isolation. Since



$Tx_1$  and  $Tx_2$  operate on the same frequency, they will have to share the air equally. Thus, at best, the users associated with  $Tx_1$  and  $Tx_2$  receive half of the time allocations they would enjoy if  $Tx_1$  and  $Tx_2$  have different frequencies. Therefore, the WiFi throughput of  $U_1$  and  $U_2$  take a hit and both users get a WiFi throughput of at most 12 Mbps. On the other hand,  $U_3$  enjoys a 24 Mbps WiFi throughput from  $Tx_3$ . On the PLC side, since all Tx antennas are set to the reuse mode, they will have to share the PLC backhaul. This negatively impacts the end-to-end throughput for  $U_1$  as discussed next.

The PLC backhaul is time-shared [17] [18] [22] between all the Tx antennas. Therefore, the PLC capacity that is obtained by each antenna is one third of the PLC capacity that it obtains when in isolation. Because of this,  $U_1$ 's end-to-end throughput falls to 7.3 Mbps. This is a direct artifact of too many extenders trying to share the backhaul PLC capacity, and thus for some extenders, the PLC part becoming a bottleneck. In this specific case, the reduction in the end-to-end throughput for  $U_1$  is driven by the fact that the PLC link throughput from the router to  $Tx_1$  becomes less than its associated WiFi link throughput due to the many extenders sharing the PLC capacity. In other words, the concatenated WiFi-PLC link for  $U_1$  becomes throttled by the PLC link segment.

On the other hand, because the PLC backhaul capacity is time shared,  $U_2$  and  $U_3$ 's PLC throughputs (from their respective transmitters) are now 53 Mbps and 54 Mbps, respectively. Thus, we see that for these users, the sharing has no impact on their PLC capacities because these (PLC) capacities well exceeded the capacity of their the wireless parts. Because they do not share the same WiFi frequency, they achieve end-to-end throughputs of at most 12 Mbps and 24 Mbps, respectively. The aggregate network throughput for this

setup is given by the summation of the achievable individual users' throughputs which is equal to 43.3 Mbps.

Now let us consider case (b), where we use DAS to combine  $Tx_2$  and  $Tx_3$ . The PLC capacity is now shared between this cluster and  $Tx_1$  i.e., each gets half of the PLC capacity. This has a significant positive impact on  $U_1$ 's PLC throughput. Specifically, since the PLC time share of  $Tx_1$  increases to half, its throughput increases to 11 Mbps (given that its WiFi throughput is 12 Mbps). With respect to  $U_2$ , and  $U_3$ , it may seem that since they now time share the single frequency allocated to their DAS cluster, their WiFi throughputs now fall to 12 Mbps and the total throughput to 35 Mbps. However, this is not the case since the benefits of DAS kick in. Because of power pooling and diversity, a higher MCS (modulation coding scheme) can now be used to these users (specifically 8-QAM instead of 4-QAM) and this improves their WiFi throughputs significantly from 12 Mbps to 18 Mbps. This in turn results in an overall throughput of 47 Mbps (a 12.2 % increase compared to the vanilla reuse case even in this very simple topology).

**Blind DAS clustering can hurt capacity.** In the previous experiment, we showed how DAS clustering can provide significant throughput gains. However, those gains stemmed from a *properly selected* DAS clustering of transmitters. In what follows, we show that an improperly clustering nearby extenders together, without accounting for their PLC capacities can in fact *decrease* the achievable throughput. Specifically, we show how improper clusterings offset any DAS gains and reduce the total network throughput.

We consider the same set of antennas as before, but now, we cluster  $Tx_1$  and  $Tx_2$  together to form one DAS transmitter as shown in Fig. 4.2c (case (c)). Thus, the cluster

consisting of  $Tx_1$  and  $Tx_2$  will jointly use the same frequency and the other frequency is used by  $Tx_3$ .  $U_1$  and  $U_2$  (associated with  $Tx_1$  and  $Tx_2$ , respectively) will receive transmissions from this DAS cluster and improve their WiFi parts of the throughput due to increased SNRs due to power pooling (the powers from the transmitters add up) and spatial diversity. The new improved SNRs for  $U_1$  and  $U_2$  allow them to utilize a modulation scheme with higher bit rates, potentially delivering a WiFi data rate of 36 Mbps each, when in isolation. However, because  $U_1$  and  $U_2$  have to alternate in using the air to receive their data over their WiFi links from the DAS cluster, the achievable WiFi throughputs for both users is 18 Mbps each. Note that because of throughput gains due to DAS, the total throughput for both users is equal to 36 Mbps (as opposed to only 24 when in reuse mode for both users).  $Tx_3$  is assigned a separate frequency and its associated user ( $U_3$ ) enjoys a WiFi throughput of 24 Mbps. Thus, the wireless parts of all users either improve or remain the same compared to the reuse case with this configuration. Unfortunately however, this becomes irrelevant because the configuration causes a throttling in the PLC parts of the concatenated links as shown next.

When  $T_1$  and  $Tx_2$  are grouped together, transmissions to those antennas are performed by treating them as a single multicast group. We note that the PLC links to  $Tx_1$  and  $Tx_2$  are of different capacities, and the router sends the data at once to all the multicast group participants. The data however, is received at different times by each participant antenna ( $Tx_1$  and  $Tx_2$ ) due to the delay differences on the PLC links. Specifically,  $Tx_2$  has to wait for  $Tx_1$  to receive the data, before performing a joint DAS transmission with that transmitter. *This makes the PLC link with the minimum capacity the determinant of the*

*PLC capacity for the whole cluster.* Since  $Tx_1$  has the lowest PLC capacity in the cluster (22 Mbps), the PLC capacity for the whole cluster falls to only 22 Mbps. Beyond this, the PLC backhaul is time-shared. Therefore, the access time for the PLC backhaul is apportioned between the DAS cluster ( $Tx_1$  and  $Tx_2$ ) and  $Tx_3$ . Therefore, the PLC throughput for the DAS cluster becomes only 11 Mbps (and 81 Mbps for  $Tx_3$ ). The end-to-end throughput for  $U_1$  and  $U_2$  will be limited to what the PLC backhaul can deliver to the cluster and it is equal to 11 Mbps (hence, the PLC throughput is the bottleneck and WiFi gains from DAS are entirely undermined). Since  $Tx_3$ 's PLC throughput is higher than the WiFi throughput of its associated user ( $U_3$ ), then the end-to-end throughput for  $U_3$  is equal to its WiFi throughput (24 Mbps). The total network throughput for this network setup is now 35 Mbps, which is less than what we achieve if DAS is not used (case (a)) and obviously if the cluster is constructed in an informed way (case (b)).

**Reuse compromise.** One way to eliminate the PLC sharing entirely is to cluster all three antennas together (configuration not shown). Now the entire set of extenders is in one cluster and becomes a multicast group. In this scenario, only one frequency (out of two) will be used and assigned to the cluster, and the other frequency will be left unutilized. This in essence compromises the benefits of the having multiple simultaneous transmissions enabled by reuse. All users  $U$  have to share the air equally and, thus, they achieve WiFi throughputs of 12 Mbps each. In addition to the wasted WiFi frequency, the PLC capacity of the cluster degrades to only 22 Mbps ( $Tx_1$  has the lowest capacity in the cluster and becomes the determinant). As a result, the end-to-end throughput of the whole cluster is throttled by its PLC link and becomes only 22 Mbps.

*In summary, it is crucial to account for individual PLC capacities when grouping antennas together into a DAS cluster. Specifically, if extenders with varying PLC capacities are grouped together in the same DAS cluster (case (c)), the PLC capacity degrades to the capacity of the poorest extender in the cluster, and as a result, the cluster's throttled PLC link offsets any DAS gains. More importantly, naive, greedy clustering of extenders into DAS clusters, would leave frequencies underutilized and thus, compromise the achievable overall capacity.*

## 4.5 Formulating the problem

In this section, we first formally define the problem that we seek to solve. Specifically, we seek to achieve the best trade-off between reuse and DAS clustering in order to maximize the aggregate throughput of the WiFi-PLC users in enterprise settings. We seek to do so by taking into account the PLC capacities to the different extenders. We formalize our problem guided by our findings in section §4.4. We find unfortunately that the problem is NP-hard (a straight forward reduction from the *graph coloring problem* [54] is possible).

**Problem Statement:** For ease of exposition, we primarily discuss downlink transmissions (which is where DAS is primarily utilized, and which carries an asymmetrically heavy part of the communication [85]) Each DAS cell (cluster) receives data from the central controller via a multicast. When all the extenders have the data, they perform a synchronized transmission to the users over the WiFi air interface.

Our objective is to ensure that we maximize the achievable throughput by eliminating undesirable PLC sharing that might arise in dense deployment of WiFi-PLC extenders.

To ensure that we do not compromise reuse (guided by the last experiment in §4.4), we first assign frequencies to a subset of extenders so as to maximally achieve reuse; then, we cluster the remaining extenders so as to maximize the aggregate throughput. Specifically, we try to grant each transmitter in the network as much airtime as possible so that the throughputs for the users increase, while accounting for individual PLC link capacities and leveraging the gains expected from using DAS. Note here that no user reassignments are performed. Thus, when an extender is included in a DAS cluster, the users that were already associated with that extender, now associate with that cluster. In what follows, “rate” refers to the PHY bit rate of a WiFi or PLC link, while throughput refers to the achieved bit rates for a user on a PLC or WiFi link given a time allocation. We also make a distinction between cell and cluster; a cell refers to one or many antennas that operate jointly to serve a subset of the users. In other words, a cell could be a “reuse” or “DAS” cell, depending on the number of antennas it includes. However, a cluster is always a DAS cluster that has more than one antenna. We formulate our optimization problem in Problem 9 below. The notations used in our formulation, are tabulated in Table I.

<b>Variable</b>	<b>Description</b>
$A$	The set of WiFi-PLC extenders.
$U$	The set of users.
$U_k$	The set of users associated with cell $k$ . $U_k$ form a partition of $U$ .
$c_j$	The PLC capacity of extender $j$ .

$f(\cdot)$	A function that takes the SNR value in dB and returns the corresponding WiFi modulation rate.
$snr_j^i$	The SNR value experienced by user $i$ from extender $j$ .
$v_{ik}$	The achievable end-to-end throughput of user $i$ from cell $k$ .
$p_{ik}$	The throughput of the PLC backhaul for user $i$ from cell $k$ .
$w_{ik}$	The WiFi link throughput for user $i$ from cell $k$ .
$x_{ik}$	A binary variable indicating whether cell $k$ serves user $i$ .
$\bar{x}_{ij}$	A binary variable indicating whether user $i$ was originally associated with extender $j$ .
$y_{jk}$	A binary variable indicating whether extender $j$ is a part of cell $k$ .
$R$	The number of cells with assigned extenders.
$\gamma_k$	A binary variable that indicates whether the signal propagation delays between extenders in cell $k$ is less than 600ns.
$\Omega_{qk}$	A binary variable which indicates whether WiFi channel $q$ is assigned to cell $k$ .
$NC$	The total number of WiFi channels/frequencies.

Table 4.1: Table of Notations

**Problem 9** *WiFi-PLC Clustering*

$$\max_{y_{jk}, \Omega_{qk}} \sum_{i,k} v_{ik} \quad (4.1)$$

$$s.t. \quad v_{ik} = \min(w_{ik}, p_{ik}) x_{ik}, \quad \forall i \in U \quad (4.2)$$

$$p_{ik} = \frac{\min_{\{j: y_{jk}=1\}} c_j}{R \sum_{i \in U_k} x_{ik}} \quad \forall i \in U, \forall k \in A \quad (4.3)$$

$$y_{jk} = \begin{cases} 1 & \text{if extender } j \text{ is} \\ & \text{part of cell } k \quad \forall j, k \in |A| \\ 0 & \text{otherwise} \end{cases} \quad (4.4)$$

$$\sum_{j=1}^{|A|} y_{jk} = 1 \quad \forall k \in A \quad (4.5)$$

$$w_{ik} = \begin{cases} \frac{\Omega_{qk}}{\sum_{i \in U_k} \frac{1}{f(\text{snr}_i^j)}}, 1, \forall i \in U, \forall k \in A, \Omega_{qk} = 1 & \text{if } \sum_{j=1}^{|A|} y_{jk} = 1 \\ \frac{D(i,j,k,y) \Omega_{qk}}{\sum_{i \in U_k} x_{ik}} \gamma_k, 1, \forall i \in U, \forall k \in A, \Omega_{qk} = 1 & \text{if } \sum_{j=1}^{|A|} y_{jk} > 1 \\ 0, \text{ otherwise} & \end{cases} \quad (4.6)$$

$$\gamma_k = \begin{cases} 1 & \text{if the delay differ-} \\ & \text{ence } \forall j \in k \text{ is } < 600 \text{ ns} \\ & \text{if the delay differ-} \\ 0 & \text{ence } \exists j \in k \geq 600 \text{ ns} \end{cases} \quad \forall k \in |A| \quad (4.7)$$

$$x_{ik} = \sum_j \bar{x}_{ij} y_{jk} \quad \forall i, k \quad (4.8)$$



$$R = \sum_k 1_{\sum_j y_{jk} > 1} \quad (4.9)$$

$$z_{k,k'} = \begin{cases} 1 & \text{if } k \text{ and } k' \text{ are neighbors} \\ 0 & \text{otherwise} \end{cases} \quad \forall k, k' \in A \quad (4.10)$$

$$\forall z_{kk'} = 1, \forall k, k' \in A$$

$$\Omega_{qk} + \Omega_{qk'} \leq 1 \quad A \quad (4.11)$$

$$\Omega_{qk} = \begin{cases} 1 & \text{if frequency } q \text{ is assigned to cell } k. \\ 0 & \text{otherwise} \end{cases} \quad \begin{matrix} \forall q \in NC \\ \forall k \in A \\ \forall q \in NC \end{matrix} \quad (4.12)$$

$$\sum_{q=1}^{NC} \Omega_{qk} = 1 \quad \forall k \in A \quad (4.13)$$

The objective (4.1) is to maximize the summation of all users' ( $i$ ) throughputs across all cells ( $k$ ) in the network. Constraint (4.2) defines the end-to-end throughput for each user as the minimum of the throughputs of its (concatenated) PLC and WiFi links. The throughput achieved on the PLC link segment is given by constraint (4.3) which is the minimum PLC link capacity in the cell (the bottleneck PLC capacity dictates the capacity of a DAS cluster as discussed in §4.4) divided by the number of cells in the system. Each DAS cell is a multicast group; thus, there is no sharing between the extenders belonging to the same cell. However, the PLC backhaul capacity in general is time-shared [18] [22]; therefore, the PLC throughput achieved in each cell is proportional to the number of the cells in the system. In

constraint (4.4),  $y_{jk}$  is a decision variable which is equal to one when extender  $j$  is assigned to cell  $k$  and 0 otherwise. Constraint (4.5) ensures that each extender is assigned to one and only one cell. Constraint (4.6) captures the achievable throughput on the WiFi link segment. The first case states that if the cell has a single extender, then the throughput is shared in a throughput-fair manner as reported in [47]. The second case defines the WiFi throughput user  $i$  enjoys if it is connected to DAS cell  $k$ . The function  $D(\cdot)$  in the numerator is given by:

$$D(i, j, k, y) = f\left(10 \log\left(\sum_{j=1}^{|A|} (10^{(snr_j^i/10)} y_{jk})\right)\right) \quad (4.14)$$

where, the  $f(\cdot)$  function in Eqn. (4.14) is one that takes the resulting DAS SNR as the input and outputs the corresponding modulation rate that can be used for that SNR. Constraint (4.7) is a system parameter that indicates whether the propagation delay differences across the extenders in one cell is less than or equal to 600ns. Constraint (4.8) requires each user to associate with the cell to which its primary extender belongs. Constraint (4.9) quantifies the number of cells that time share the PLC backhaul. Specifically, it counts the number of cells that have extenders assigned to them. Constraint (4.10) defines a parameter to indicate whether two cells would interfere in the WiFi domain if they were to be assigned the same frequency. Constraint (4.11) ensures that no two interfering cells (in the WiFi domain) have the same frequency. Finally, constraint (4.12) specifies what frequency is assigned to each cell and (4.13) postulates that each cell must not have more than one assigned frequency.

Since proper throughput sharing (and maximization) is most relevant to heavily loaded scenarios, our model assumes that each user has a saturated flow. Next, we discuss

the hardness of Problem (9).

**Time Complexity Matters.** Consider a setting of a university lab with 30 extenders. Let say we want to form six clusters, each of which has five extenders. The number of possible ways to form these clusters is then given by  $\binom{30}{30-5}$  possibilities. The number of possibilities increases significantly if the clusters are assumed to have variable numbers of extenders. Therefore, if a brute force approach were to be applied, a prohibitively high time complexity is inevitable. Having provided this intuition, we formally prove next that Problem (9) is NP-hard by deriving a reduction from the *graph coloring problem* [54] in Theorem (5) below. The key idea in the proof is to construct a simple instance of Problem (9) to which, we can reduce a graph coloring problem. Thus, since the graph coloring problem is not solvable in polynomial time, the simple instance of Problem (9) is not solvable in polynomial time . Hence, Problem (9) is NP-hard.

**Theorem 5** *Prob. 9 is NP-hard.*

**Proof.** Denote a solution to Problem 9 as  $(y_{jk}^*, \Omega_{qk}^*)$ . Let  $\gamma_k^* = 0 \forall k \in A$  and  $z_{kk'}^* = 1, \forall k, k' \in A$  be the inputs to Problem 9. Then we propose the following polynomial-time transformation from the graph coloring problem [54] to the instance of our problem. If  $\gamma_k^* \forall k \in A$  is equal to 0, then constraint 4.6 reduces to the first case only with the condition  $\sum_{j=1}^{|A|} y_{jk}^* = 1, \forall i \in U, \forall k \in A, \Omega_{qk}^* = 1$ . That is, each cell  $k \in A$  has only and only one  $j$  assigned to it (hence,  $\sum_{j=1}^{|A|} y_{jk}^* = 1$ ). The second case of constraint 4.6 is not feasible since  $\gamma_k^* = 0 \forall k \in A$  which precludes the solution of Problem 9 from creating DAS clusters. Thus,  $y_{jk}^* \neq y_{j'k}^* \neq 1, \forall j, j', k \in A$  as per the first case of constraint 4.6. What remains is the specifications of the values of  $\Omega_{qk}^*$ . The constraints 4.10-4.13 mandate that each

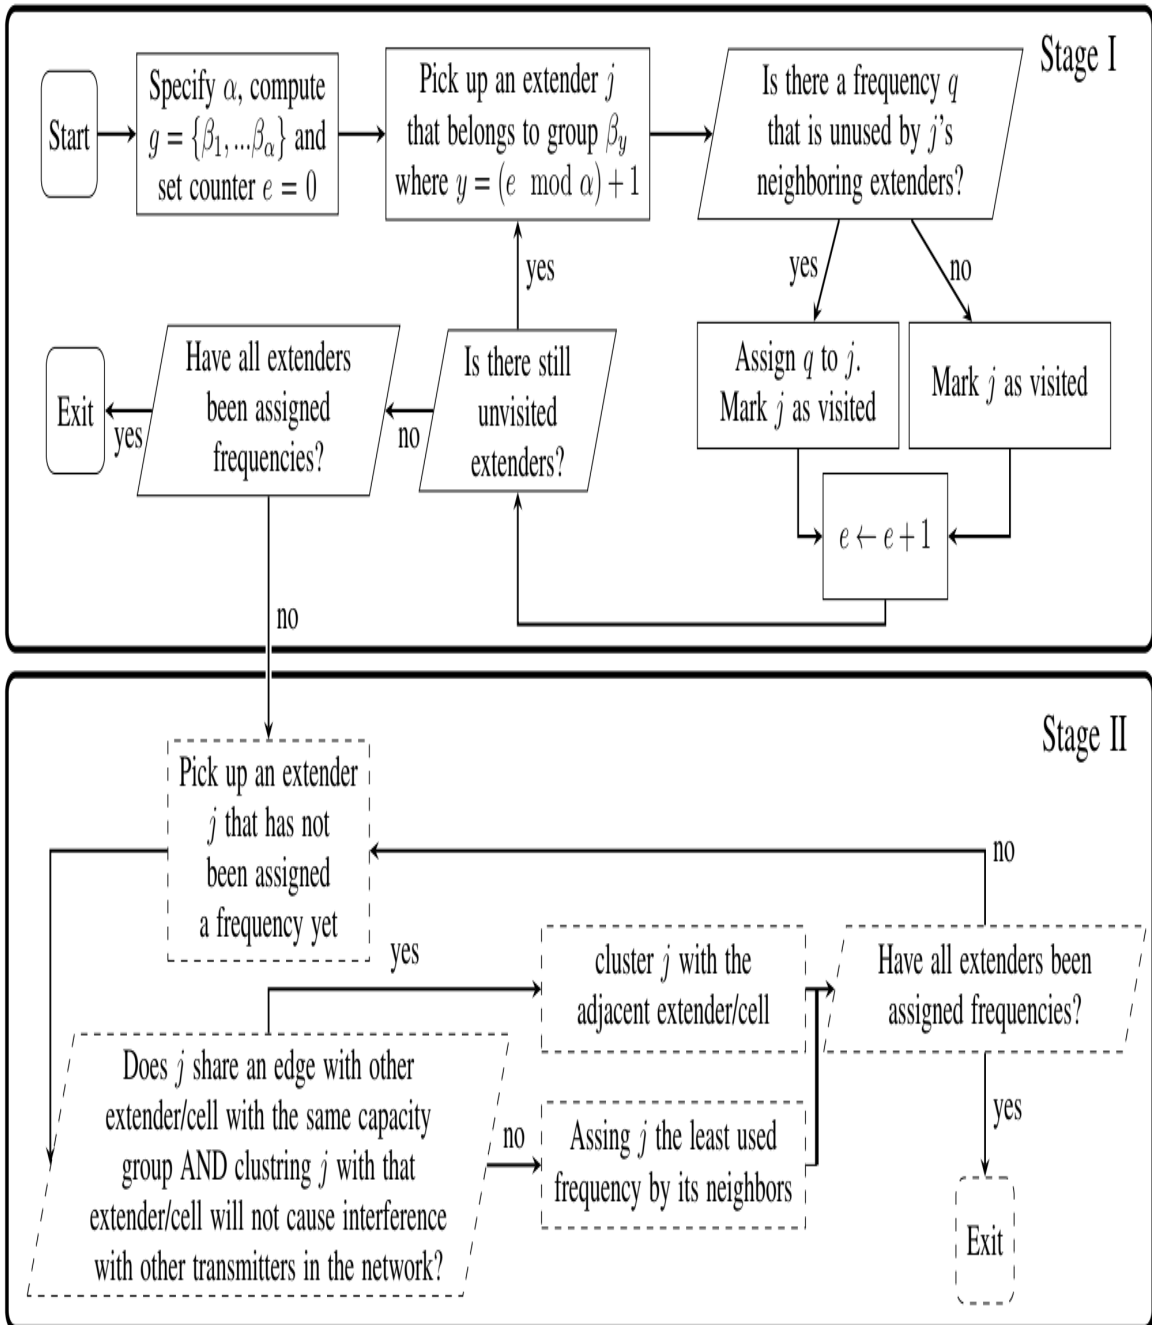


Figure 4.3: Algorithm for frequency assignments and DAS Cell Formation

$k \in A$  must be assigned a frequency that is different from that of its neighbors. Since  $z_{kk'}^* = 1 \forall k, k' \in A$ , then  $\Omega_{qk}^* \neq \Omega_{qk'}^* \neq 1, \forall k, k' \in |A|, \forall q \in NC$ . Now the problem of assigning colors to nodes in a complete graph can be reduced to this by replacing each node in the graph with  $k \in A$  and  $\Omega_{qk}^*$  being the solution that determines whether the color (frequency)  $q$  is assigned to node  $k$  or not. Hence, solving the graph coloring problem on a complete graph is reduced to solving this particular instance of Problem 9. Since the graph coloring problem is NP-complete, solving the general Problem 9 is NP-hard. ■

## 4.6 Priza and its components

In this section we describe how our framework Priza organizes a WiFi-PLC network into DAS cells and how these cells are assigned frequencies. Priza’s workflow is shown in Figure 4.3 and runs in two stages: (i) it assigns frequencies to a subset of the extenders (antennas) such that there is no interference (neighboring extenders are not assigned the same frequency) to maximize reuse, and (ii) it clusters the remaining extenders to form DAS clusters in an informed way, specifically with respect to their PLC backhaul capacities. The reason why Priza assigns frequencies before any clustering is performed is to fully retain any gains that spectrum reuse can provide. If the DAS clusters are constructed first, then we might waste part of the system capacity due to potential overuse of DAS<sup>2</sup>. In section §4.7.4, we discuss two approaches where, DAS clusters are formed first, based on naive policies, and show that they do not perform as well as Priza.

---

<sup>2</sup>One can potentially split the formed DAS clusters to regain losses from inefficient spectral reuse, but that would make the approach inherently complex and convergence may be an issue.

Priza groups PLC extenders based on their PLC capacities. Formally, let  $\alpha$  be the number of available frequencies. Then, we set the number of these PLC groups to  $\alpha$ . If  $c_j$  represents the PLC capacities to the various extenders, then, we compute a set of intervals  $g$ , given by  $g = \{\beta_1, \beta_2, \dots, \beta_\alpha\}$ . Each “capacity group interval,”  $\beta_n$  is of width:

$$\beta_n = \left[ (n-1) \frac{\max(c_j)}{\alpha}, n \frac{\max(c_j)}{\alpha} \right], \forall n \in [1, 2, \dots, \alpha]. \quad (4.15)$$

Each extender is then tagged by the index of the capacity group interval, within which its capacity lies. Specifically, if extender  $i$  has PLC capacity  $c_i \in \beta_l$ , ( $l \in \{1, \alpha\}$ ) then it is tagged as a part of the  $l^{th}$  group.

To assign frequencies, Priza maps the extenders on to an interference graph (as is typically done [80]). Each extender is represented by a node in the graph and the edges are added between each pair of extenders if they interfere with each other in the WiFi domain. Specifically, if an extender is able to receive a signal of more than 4 db from another extender [80], then these two extenders are considered to interfere with each other and an edge is added to the interference graph between them. After tagging the extenders with their corresponding capacity groups and creating the interference graph, Priza is now able to make frequency assignment and clustering decisions. Next, we describe its two stages.

**Stage I:** First, Priza’s algorithm iterates over all WiFi-PLC extenders and picks one extender from each capacity group  $\beta_n \in g$  and assigns a frequency to it. The idea here is to assign a frequency to extenders that differ in terms of their PLC capacities (based on our understanding from our measurements). Each chosen extender is assigned a frequency that is not used by its neighboring extenders. If an extender cannot be assigned frequency

different from its neighboring extenders, then it is tagged as 'visited' and, subsequently, left to be handled in Stage II. The objective of Stage I is to ensure that we maximally exploit the achievable spectral reuse spatially. The reason that Priza iterates over the capacity groups is to ensure that extenders from different capacity groups are separated in frequency. Thus, extenders that are left without frequency assignments in Stage I will have a better chance to be clustered with extenders that are in the same capacity group (as we will show in Stage II), thereby alleviating PLC capacity discrepancies within each group. It is easy to see that, in an extreme case, if we were to assign frequencies to all extenders with capacities in  $\beta_n$ , then we could end up with complete sets of extenders belonging to the other capacity groups  $\beta_{n'}, \forall n' \in \alpha, n' \neq n$  that have to be clustered (due to the lack of available frequencies) with cells that have extenders in  $\beta_n$ . This would cause extenders with diverse PLC capacities to be grouped together, which hurts capacity as discussed in §4.4).

**Stage II:** After determining how frequencies are initially assigned, Priza in Stage II will need to iterate over all the extenders that have not been assigned frequencies by Stage I, and make the decisions on whether each such extender should join a neighboring cell (thus, converting the reuse cell into a DAS cell) or if it should be a separate cell by itself. This is done as follows: if extender  $j$  (that has not been assigned frequency) shares an edge with another cell which is in the same capacity group, and clustering  $j$  with that cell does not create an interference with other cells in the system (due to DAS power pooling [84]), then Priza clusters  $j$  with that cell. Otherwise, extender  $j$  will form a cell by itself and be assigned a frequency that has the least interference from its neighbors. This process ensures that the system reuse achieved in Stage I is maintained to the extent viable by clustering

extenders together unless DAS clustering causes interference that did not exist prior to the clustering. In this case, the system capacity will take a small hit because of creating a new reuse cell.

**Algorithm Complexity:** Stage I of the above algorithm runs in  $\mathcal{O}(|A|)$  where  $|A|$  is the number of extenders in the system. The runtime of Stage II is given by  $\mathcal{O}(|A| - m)$  where  $m$  is the number of extenders with assigned frequencies in Stage I. Therefore, the total runtime for the entirety of Algorithm 1 is  $\mathcal{O}(|A|)$ .

## 4.7 IMPLEMENTATION & EVALUATION

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

### 4.7.1 Implementation Details

**Testbed configuration and equipment:** We refer the reader to §4.4, where our testbed and configurations were described. In brief, we perform experiments with five radios as Tx antennas and three radios as users. The Tx antennas are assigned PLC capacities to emulate a real PLC backhaul (see below). The experiments are performed in a 2664 square feet university lab with chairs, desks, research equipment and two cubicles. All the radios (Tx antennas and users) are randomly moved around to create ten different topologies. We set the number of the available frequencies to less than the number of extenders (Typically, 2 to 3 frequencies are used) to ensure that we have adequate reuse opportunities.

**Software Implementation:** Priza is implemented in Labview as a user space



utility that runs on the server. Five (out of eight) USRP radios are designated as Tx antennas and the remaining three are users. Priza starts with building the interference graph (as described in §4.6). It does so by sending a short beacon packet from each Tx antenna to all other Tx antennas in the testbed. The other Tx antennas compute the SNR upon receipt. If the SNR is found to be more than 4 db between two Tx antennas [80] then they are marked as interfering. The same process is repeated for each Tx antennas and edges are added to the interference graph accordingly.

As discussed earlier in §4.4, the PLC throughputs are estimated and since they remain stable over time to each extender, we use emulations thereof. The concatenation of these emulated PLC links and real WiFi links is performed by Priza. Specifically, we measure the end-to-end throughput for each user as the user’s minimum PLC throughputs and WiFi throughputs as defined in constraint (4.2) of Problem 9.

### 4.7.2 Simulations

To evaluate Priza in larger settings, we implement Priza entirely in Matlab [71]. Specifically, we construct a WiFi-PLC network with 50 to 70 extenders and 70 to 100 users. The PLC capacities are taken from measurements from real outlets in our university building. Both the extenders and the users are randomly and geographically distributed in a 2D-plane with an area of  $108 \times 148$  square feet commensurate with the area of six real university labs (similar to that in Figure 4.1). The WiFi links between each extender and user are assigned SNR values based on the physical distance between them and models capturing the shadowing and fading effects for non-line-of-sight (NLoS) indoor signal propagation as reported in [62]. The resulting DAS SNR is estimated using Eqn. (4.14) (also

reported by [29] and [18]). Each SNR value is mapped to an appropriate modulation scheme as in [39]. Each modulation scheme is designed to encode the appropriate number of bits within each OFDM symbol [15].


### 4.7.3 Baselines & Performance Metrics

**Baselines:** Priza is evaluated against three baselines: (a) Reuse, (b) Balanced-DAS (B-DAS) and (c) Large-DAS (L-DAS). Reuse assigns each extender a frequency that is least used by its neighbors and does not create DAS cells. This is the default mode of spectrum sharing with PLC extenders. B-DAS creates equally-sized DAS cells by grouping the nearest extenders together. L-DAS creates as large DAS cells as possible without violating the 600 ns propagation delay constraint. This baseline represents the naive way of creating DAS clusters without carefully ensuring that reuse gains are not lost.

**Performance Metrics:** Our performance metrics of interest are (a) the aggregate throughput achieved by Priza and (b) the fairness of the users' individual throughputs. We use the commonly used Jain's fairness index [48] to evaluate the latter.

### 4.7.4 Experimental Evaluations

**Gains in Aggregate Throughputs:** We first depict the throughputs with Priza and the baselines in Figures 4.4a and 4.4b. We see that Priza outperforms the other baselines and achieves higher aggregate throughputs. The average aggregate (network-wide) throughput improvements are 33.7%, 56.5% and 144.6% over Reuse, B-DAS and L-DAS, respectively. In the second figure, we observe that Priza outperforms all other baselines in all trials. These are attributable to Priza being able to efficiently share the PLC capacity,



(a) Priza, Reuse, B-DAS and L-DAS comparison on throughputs. (b) CDF of the aggregate throughput. (c) Validating the fidelity of our simulations. testbed.

Figure 4.4: Experimental results.

and the gains from DAS as discussed in § 4.4, unlike the other schemes.

**Fairness:** Even though Priza does not consider fairness explicitly, it achieves comparable fairness to those of the baselines, which inherently are designed to be fair as discussed below. Specifically, Priza yields a Jain’s fairness index of 0.9034 compared to 0.934 with reuse, 0.951 with B-DAS and almost perfect fairness of 1.0 with L-DAS. Reuse by design tries to allocate frequencies in a fair way (least congested), B-DAS has similar sized clusters and thus, each cluster has a similar likelihood of experiencing poor PLC links and all clusters get the same DAS gain, and with L-DAS the poorest PLC link is what dictates the throughput for all users. Priza creates different sized DAS clusters, and groups good extenders in the same group, and not so good extenders into different groups. However, all PLC links see improvements, and so do the WiFi links due to DAS gains; thus, it is able to largely ensure that the throughputs that are achieved are similar.

#### 4.7.5 Simulation Results

**Simulation Fidelity:** To show that our simulations reflect what one might observe in a real testbed, we simulate ten topologies with the same users and antennas locations

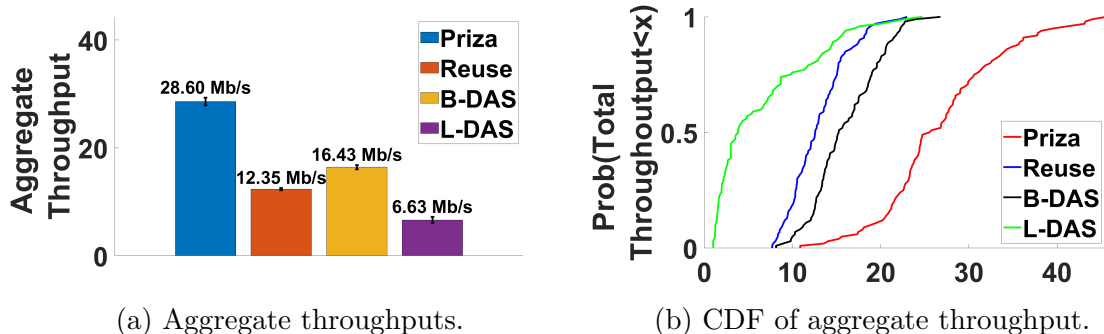


Figure 4.5: Simulation Results.

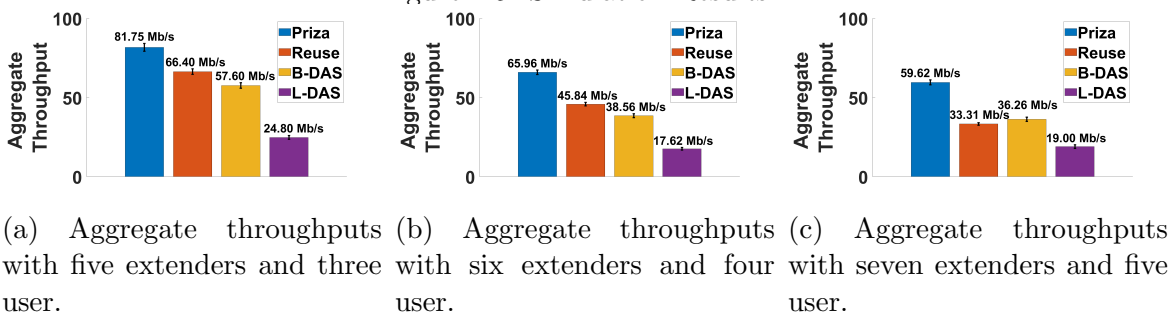


Figure 4.6: Average aggregate throughputs. The results show how the size of the topologies affects the Reuse and B-DAS baselines.

of our testbed experiments. The results are shown in Fig. 4.4c. We see that the simulation yields similar results to those obtained from the real testbed experiments, thus showcasing the high fidelity of the models that we use.

**Aggregate throughputs:** This simulation experiment was as described in § 4.7.2 with the number of available frequencies set to 11 so as to reflect the case of the actual number of WiFi channels in 802.11b [53]. The results of our simulations are provided in Fig. 4.5. Fig. 4.5a shows how Priza outperforms other baselines and yields higher aggregate throughputs. *Aggregate throughput improvements of 131.5%, 74% and 331.3% are observed over Reuse, B-DAS and L-DAS respectively.*

We note however that in contrast with the results from the testbed, B-DAS out-

performs Reuse. We attribute this to the significant increase in extender density and scale. In Fig. 4.6 we show the variations in relative aggregate throughputs of B-DAS and Reuse with scale. When the topology is small (same size as our testbed), Reuse outperforms B-DAS (Fig. 4.6a) but as the network size grows larger, B-DAS starts to improve compared to Reuse (Fig. 4.6b). Fig. 4.6c shows how B-DAS overtakes Reuse when the topology size becomes even larger. This is because as the network size increases, PLC inefficiency increases and B-DAS helps with alleviating the same; this combined with the DAS gains even in a naive way helps improve throughput beyond what Reuse yields (although still significantly lower than that of Priza).

**Fairness:** In Table 4.2, we show the Jain’s fairness index values with respect to both (a) the end-to-end throughput over the concatenated WiFi-PLC link and (b) the WiFi link segment only. Priza’s fairness as well as all other baselines (except L-DAS) take a hit since the number of cells is large in the simulation, and the PLC links can be diverse in terms of their capacities, causing high variations in the users’ end-to-end throughputs. In the lower part of Table 4.2 we show the fairness index values for the same users when the PLC backhaul has a very high capacities (e.g,  $w_{ik} = \min(w_{ik}, p_{ik})x_{ik}, \forall i \in U$ ). In such cases, all methods maintain high fairness index values since, indicating that the methods share the WiFi domain in a fair manner (as defined in constraint (4.6) of Problem 9).

End-to-end throughputs	Priza	0.47702
	Reuse	0.46332
	B-DAS	0.50288
	L-DAS	1
End-to-end throughputs with high PLC capacities	Priza	0.98255
	Reuse	0.98802
	B-DAS	0.97814
	L-DAS	0.99674

Table 4.2: Simulations Jain’s fairness index

## 4.8 Related Work

**Hybrid WiFi-PLC:** There are some efforts on applications that use hybrid WiFi-PLC networks. the authors of [96] develop a synchronization mechanism that uses electrical wiring for MIMO access points. Other papers *e.g.*, [46] consider maximizing the throughput of hybrid WiFi-PLC networks by exploiting multipath routes in the PLC backhaul. These efforts however, do not consider the impact of the PLC backhaul being the bottleneck. In our prior work [17], we present an algorithm to maximize the aggregate throughput via user assignments while accounting for the impact of the PLC backhaul. However, we do not consider using DAS nor the problem of inefficient PLC backhaul sharing that arises in enterprise settings like the work proposed in this paper.

**Distributed Antenna Systems:** There are efforts such as [84] and [94] that use DAS for increasing the throughput via dynamic clustering or transmission power allocations. None of these however, consider a PLC backhaul. In [18], we consider DAS over PLC for home settings, but there was no study of inefficiency across multiple user connections and only a single DAS cell was constructed.

**WiFi throughput maximization:** Papers like [25, 58] (among others) propose

algorithms to maximize throughput in WiFi networks. However, they all assume a high capacity backhaul such as Ethernet or fiber.

## 4.9 Conclusions

In this paper we consider the dense deployment of plug and play PLC based WiFi extenders in enterprises and showcase some of the pitfalls of doing so. Specifically, via a measurement study we find that inefficient sharing of the PLC backhaul can significantly impact the achievable capacity. We then propose Priza, a framework that maximizes the aggregate throughputs in such settings by grouping the PLC extenders into DAS clusters. In particular, Priza’s two prong approach retains the gains from frequency reuse, but clusters PLC extenders so as to mitigate the inefficiencies in sharing the PLC backhaul. It also provides the inherent benefits of power pooling and diversity that DAS provides. We show that if instead of using Priza, a blind and uninformed clustering of extenders is employed, it does not solve the inefficiencies, and could further degrade the achievable throughput. We perform extensive experiments both on a real testbed and via simulations to showcase the superiority of Priza over both the vanilla reuse deployments, and naive DAS baselines.

## Chapter 5

# Conclusions

In this dissertation, we identify some of the issues that could affect the aggregate throughput of WiFi-PLC networks. Specifically, we show that the end-to-end throughput of the concatenated WiFi-PLC link is, in fact, determined by the minimum throughput of its composed link segments. Therefore, we discuss the problems of user-extenders associations and poor WiFi link qualities. For the former we proposed WOLT, a framework that cleverly assigns users to WiFi-PLC extenders. For the latter, we developed two frameworks: (a) PLC-DAS that constructs DAS clusters of WiFi-PLC extenders in home settings to improve users' WiFi signals and (b) Priza that assigns frequencies to WiFi-PLC extenders and then form DAS clusters to improve the sharing of the PLC backhaul and boost WiFi signals. All our solutions share the same common goal of increasing the aggregate throughput of the WiFi-PLC networks.



# Bibliography

- [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] Cisco aironet 1200 series access point. <http://www.vsp.virginia.gov>.
- [3] Highlights of annual 2019 characteristics of new housing. [www.census.gov/](http://www.census.gov/).
- [4] Octoclock cda-2990. <https://www.ettus.com/all-products/octoclock/>.
- [5] Pinterest.
- [6] Qualcomm atheros open powerline toolkit. <https://github.com/qca/open-plc-utils>.
- [7] Tcp/udp bandwidth measurement tool. <https://iperf.fr/>.
- [8] Tp-link 300mbps wi-fi range extender, av600 powerline edition. <https://www.tp-link.com/us/home-networking/powerline/tl-wpa4220kit/>.
- [9] Tp-link powrline ac wifi kit. <https://www.tp-link.com/us/home-networking/powerline/tl-wpa8630-kit/>.
- [10] Usrp hardware driver and usrp manual.
- [11] Usrp n210. <https://www.ettus.com/all-products/un210-kit/>.
- [12] Warp project.
- [13] *Power Line Communication (PLC) Systems - Global Market Trajectory & Analytics*. Global Industry Analysts, Inc, 2020.
- [14] *Power Line Communication Market - Global Industry Analysis, Size, Share, Growth, Trends and Forecast 2016 - 2024*. Transparency Market Research, 2021.
- [15] Abdeldime MS Abdelgader and Wu Lenan. The physical layer of the ieee 802.11 p wave communication standard: the specifications and challenges. In *WCECS'14*.
- [16] Hisham A Alhulayyil, Kittipat Apicharttrisorn, Jiasi Chen, Karthikeyan Sundaresan, Samet Oymak, and Srikanth V Krishnamurthy. Wolt: Auto-configuration of integrated enterprise plc-wifi networks.

- [17] Hisham A Alhulayyil, Kittipat Apicharttrisorn, Jiasi Chen, Karthikeyan Sundaresan, Samet Oymak, and Srikanth V Krishnamurthy. Wolt: Auto-configuration of integrated enterprise plc-wifi networks. In *ICDCS'20*.
- [18] Hisham A Alhulayyil, Jiasi Chen, Karthikeyan Sundaresan, and Srikanth V Krishnamurthy. Boosting home wifi throughputs via adaptive das clustering of plc extenders. In *MASS'21*.
- [19] Mohammed AMER, Anthony Busson, and Isabelle 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*, MSWiM '16.
- [20] Global Industry Analysts. Private tutoring global market trajectory & analytics. *Global Industry Analysts, Inc*, 2020.
- [21] Kittipat Apicharttrisorn, Ahmed Osama Fathy Atya, Jiasi Chen, Karthikeyan Sundaresan, and Srikanth V. Krishnamurthy. Enhancing wifi throughput with plc extenders: A measurement study. In *Passive and Active Measurement*, 2017.
- [22] Kittipat Apicharttrisorn, Ahmed Osama Fathy Atya, Jiasi Chen, Karthikeyan Sundaresan, and Srikanth V Krishnamurthy. Enhancing wifi throughput with plc extenders: A measurement study. In *International Conference on Passive and Active Network Measurement*, pages 257–269. Springer, 2017.
- [23] Ehsan Aryafar, Alireza Keshavarz-Haddad, Michael Wang, and Mung Chiang. Rat selection games in hetnets. In *2013 Proceedings IEEE INFOCOM*, pages 998–1006. IEEE, 2013.
- [24] IEEE Standards Association et al. Ieee standard for broadband over power line networks: Medium access control and physical layer specifications. *IEEE Std*, 2010(2010):1–1586, 1901.
- [25] Ramy Atawia and Haris Gacanin. Self-deployment of future indoor wi-fi networks: An artificial intelligence approach. In *IEEE GLOBECOM'17*.
- [26] Ahmed Osama Fathy Atya, Karthikeyan Sundaresan, Srikanth V. Krishnamurthy, Mohammad A. Khojastepour, and Sampath Rangarajan. Bolt: Realizing high throughput power line communication networks. In *Proceedings of the 11th ACM Conference on Emerging Networking Experiments and Technologies*, CoNEXT '15. ACM.
- [27] Anand Balachandran, Geoffrey M Voelker, Paramvir Bahl, and P Venkat Rangan. Characterizing user behavior and network performance in a public wireless lan. In *ACM SIGMETRICS Performance Evaluation Review*, volume 30, pages 195–205. ACM, 2002.

- [28] Giuseppe Bianchi, Simone Di Domenico, Mauro De Sanctis, Laura Liberati, Valerio Perrotta, and Ernestina Cianca. Unveiling access point signal instability in wifi-based passive sensing. In *ICCCN '17*. IEEE.
- [29] Donald G Brennan. Linear diversity combining techniques. *Proceedings of the IRE*, 1959.
- [30] 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.
- [31] Miguel Elias M Campista, Luís Henrique Maciel Kosmowski Costa, and Otto Carlos Muniz Bandeira Duarte. Improving the data transmission throughput over the home electrical wiring. In *LCN '05*. IEEE, 2005.
- [32] Daniel Castanheira and Atilio Gameiro. Distributed antenna system capacity scaling [coordinated and distributed mimo]. *IEEE Wireless Communications*, 2010.
- [33] Jiasi Chen, Mung Chiang, Jeffrey Ercan, Guangzhi Li, KK Ramakrishnan, and Rakesh K Sinha. Fair and optimal resource allocation for lte multicast (embms): Group partitioning and dynamics. In *2015 IEEE Conference on Computer Communications (INFOCOM)*, pages 1266–1274. IEEE, 2015.
- [34] Xin Chen, Xiaodong Xu, and Xiaofeng Tao. Energy efficient power allocation in generalized distributed antenna system. *IEEE Communications Letters*, 2012.
- [35] Francesco Chiti, Romano Fantacci, Dania Marabissi, and Andrea Tani. Performance evaluation of an efficient and reliable multicast power line communication system. *IEEE Journal on Selected Areas in Communications*, 2016.
- [36] Wan Choi and Jeffrey G Andrews. Downlink performance and capacity of distributed antenna systems in a multicell environment. *IEEE transactions on wireless communications*.
- [37] Sunil Chopra and Mendu R Rao. The partition problem. *Mathematical Programming*, 59(1-3):87–115, 1993.
- [38] Thomas H Cormen, Charles E Leiserson, Ronald L Rivest, and Clifford Stein. *Introduction to algorithms*. MIT press, 2009.
- [39] Manish Dangi and Mahesh Kumar Porwal. Analyses of snr threshold for minimum ber in various modulations schemes and development of an adaptive modulation scheme. *IJISSE-Int. J. Innov. Sci. Eng. Technol*, 2015.
- [40] Virginie Degardin, Martine Lienard, Ahmed Zeddani, Frédéric Gauthier, and P Degauquel. Classification and characterization of impulsive noise on indoor powerline used for data communications. *IEEE Transactions on Consumer Electronics*, 2002.

- [41] Dacfeý Dzung, Inigo Berganza, and Alberto Sendin. Evolution of powerline communications for smart distribution: From ripple control to ofdm. In *2011 IEEE International Symposium on Power Line Communications and Its Applications*. IEEE, 2011.
- [42] Mahmoud Elgenedy, Mostafa Sayed, Naofal Al-Dhahir, and Rakan C Chabaan. Cyclostationary noise mitigation for simo powerline communications. *IEEE Access*'18.
- [43] Shmuel Goldfisher and Shinji Tanabe. Ieee 1901 access system: An overview of its uniqueness and motivation. *IEEE Communications Magazine*, 48(10):150–157, 2010.
- [44] Robert Heath, Steven Peters, Yi Wang, and Jiayin Zhang. A current perspective on distributed antenna systems for the downlink of cellular systems. *IEEE Communications Magazine*, 2013.
- [45] Sébastien Henri, Christina Vlachou, Julien Herzen, and Patrick 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, CoNEXT '16*. ACM.
- [46] Sébastien Henri, Christina Vlachou, Julien Herzen, and Patrick Thiran. Empower hybrid networks: Exploiting multiple paths over wireless and electrical mediums. In *CoNEXT'16*.
- [47] Martin Heusse, Franck Rousseau, Gilles Berger-Sabbatel, and Andrzej 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)*, volume 2, pages 836–843. IEEE, 2003.
- [48] Raj Jain, Arjan Duresi, and Gojko Babic. Throughput fairness index: An explanation. In *ATM Forum contribution*, volume 99, 1999.
- [49] Amit P. Jardosh, Kimaya Mittal, Krishna N. Ramachandran, Elizabeth M. Belding, and Kevin C. Almeroth. Iqu: Practical queue-based user association management for wlans. In *Proceedings of the 12th Annual International Conference on Mobile Computing and Networking, MobiCom '06*. ACM, 2006.
- [50] Simon Jary. Best powerline adapters 2022: Improve your home network speeds. <https://www.techadvisor.com/test-centre/network-wifi/best-powerline-adapters-3490638/>, 2022.
- [51] Jingon Joung, Yeow Khiang Chia, and Sumei Sun. Energy-efficient, large-scale distributed-antenna system (l-das) for multiple users. *IEEE Journal of Selected Topics in Signal Processing*, 2014.
- [52] Tomoya Kageyama, Osamu Muta, and Haris Gacanin. An enhanced selected mapping technique for joint papr reduction and impulsive noise suppression in multi-carrier powerline communications systems. In *IEEE ISPLC'18*.

- [53] Steve Kapp. 802.11: leaving the wire behind. *IEEE Internet Computing*, 6(1):82–85, 2002.
- [54] David Karger, Rajeev Motwani, and Madhu Sudan. Approximate graph coloring by semidefinite programming. *Journal of the ACM JACM'98*.
- [55] 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.
- [56] Srinivas Katar, Brent Mashburn, Kaywan Afkhamie, Haniph Latchman, and Richard Newman. Channel adaptation based on cyclo-stationary noise characteristics in plc systems. In *2006 IEEE International Symposium on Power Line Communications and Its Applications*, pages 16–21. IEEE, 2006.
- [57] Bruno Kauffmann, François Baccelli, Augustin Chaintreau, Vivek Mhatre, Konstantina Papagiannaki, and Christophe Diot. Measurement-based self organization of interfering 802.11 wireless access networks. In *Infocom*, volume 7, pages 1451–1459. Citeseer, 2007.
- [58] Evgeny Khorov, Anton Kiryanov, and Alexander Krotov. Joint power control and time division to improve spectral efficiency in dense wi-fi networks. In *IEEE Black-SeaCom'18*.
- [59] Toshiaki Koike-Akino, Andreas F Molisch, Chunjie Duan, Zhifeng Tao, and Philip Orlik. Capacity, mse and secrecy analysis of linear block precoding for distributed antenna systems in multi-user frequency-selective fading channels. *IEEE Transactions on Communications*, 2011.
- [60] David Law, Dan Dove, John D'Ambrosia, Marek Hajduczenia, Mark Laubach, and Steve Carlson. Evolution of ethernet standards in the ieee 802.3 working group. *IEEE Communications Magazine*, 51(8):88–96, 2013.
- [61] 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.
- [62] Jing Lei, Roy Yates, Larry Greenstein, and Hang Liu. Wireless link snr mapping onto an indoor testbed. In *First International Conference on Testbeds and Research Infrastructures for the Development of Networks and Communities*. IEEE.
- [63] Douglas J Leith and Peter Clifford. Tcp fairness in 802.11 e wlans. In *2005 International Conference on Wireless Networks, Communications and Mobile Computing*, volume 1, pages 649–654. IEEE, 2005.
- [64] Jiaxin Li, Ke Xiong, Jie Cao, Xi Yang, and Tong Liu. Energy efficiency in rf energy harvesting-powered distributed antenna systems for the internet of things. *Sensors*, 2020.

- [65] Jing Lin, Marcel Nassar, and Brian L Evans. Impulsive noise mitigation in powerline communications using sparse bayesian learning. *IEEE Journal on Selected Areas in Communications*, 2013.
- [66] Wang Lingzhi, Hua Cunqing, Zheng Rong, and Ni Rui. Online channel selection and user association in high-density wifi networks. In *2015 IEEE International Conference on Communications (ICC)*, pages 1571–1576. IEEE, 2015.
- [67] 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.
- [68] Rajesh Mahindra, Hari Viswanathan, Karthik Sundaresan, Mustafa Y Arslan, and Sampath Rangarajan. A practical traffic management system for integrated lte-wifi networks. In *Proceedings of the 20th annual international conference on Mobile computing and networking*, pages 189–200. ACM, 2014.
- [69] Ali Maiga, Jean-Yves Baudais, and Jean-François Héland. Increase in multicast ofdm data rate in plc network using adaptive lp-ofdm. In *ICAST '09*. IEEE.
- [70] Kazuki Maruta. Dynamic clustering and coordinated user scheduling for cooperative interference cancellation on ultra-high density distributed antenna systems. *Entropy'18*.
- [71] MATLAB. *version 9.6.0.1072779 (R2019a)*. The MathWorks Inc., Natick, Massachusetts, 2019.
- [72] Munehiro Matsui, Tadao Nakagawa, Riichi Kudo, Koichi Ishihara, and Masato Mizoguchi. Blind frequency-dependent iq imbalance compensation scheme using cma for ofdm system. In *2011 IEEE 22nd International Symposium on Personal, Indoor and Mobile Radio Communications*. IEEE.
- [73] Vivek P Mhatre, Konstantina Papagiannaki, and Francois Baccelli. Interference mitigation through power control in high density 802.11 wlans. In *IEEE INFOCOM'07*.
- [74] G Ayorkor Mills-Tettey, Anthony Stentz, and M Bernardine Dias. The dynamic hungarian algorithm for the assignment problem with changing costs. 2007.
- [75] Rohan Murty, Jitendra Padhye, Ranveer Chandra, A Roy Chowdhury, and Matt Welsh. Characterizing the end-to-end performance of indoor powerline networks. *Harvard University Technical Report*, 2008.
- [76] Tsutomu Niiho, Mariko Nakaso, Koichi Masuda, Hiroyuki Sasai, Kuniaki Utsumi, and Masaru Fuse. Multi-channel wireless lan distributed antenna system based on radio-over-fiber techniques. In *LEOS '04*. IEEE.
- [77] Hassna Ouassal, Torre Rocco, Ming Yan, and Jeffrey A Nanzer. Decentralized frequency synchronization in distributed antenna arrays with quantized frequency states and directed communications. *IEEE Transactions on Antennas and Propagation*, 2020.

- [78] Jaehyun Park and Bruno Clerckx. Multi-user linear precoding in massively distributed polarized antenna systems under imperfect csit. *IEEE Transactions on Vehicular Technology*, 2020.
- [79] Silvia Pato, João Pedro, João Santos, Artur Arsénio, Pedro Inácio, and Paulo Monteiro. On building a distributed antenna system with joint signal processing for next generation wireless access networks: The futon approach. In *7th Conference on Telecommunications, Portugal*, 2008.
- [80] Iyappan Ramachandran and Sumit Roy. Wlc46-2: On the impact of clear channel assessment on mac performance. In *IEEE Globecom 2006*.
- [81] Nirmalya Roy, David Kleinschmidt, Joseph Taylor, and Behrooz 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*, BuildSys'13. ACM.
- [82] Olamide M Shekoni, Ali N Hasan, and Thokozani Shongwe. Applications of artificial intelligence in powerline communications in terms of noise detection and reduction: a review. *Australian Journal of Electrical and Electronics Engineering*, 2018.
- [83] Vivek Shrivastava, Shravan Rayanchu, Suman Banerjee, and Konstantina Papagianaki. Pie in the sky: Online passive interference estimation for enterprise wlans. In *Proceedings of the 8th USENIX Conference on Networked Systems Design and Implementation*, NSDI'11, pages 337–350, Berkeley, CA, USA, 2011. USENIX Association.
- [84] Shailendra Singh, Karthikeyan Sundaresan, Srikanth V Krishnamurthy, Xinyu Zhang, Amir Khojastepour, and Sampath Rangarajan. Trinity: A practical transmitter cooperation framework to handle heterogeneous user profiles in wireless networks. In *ACM MobiHoc'15*.
- [85] Bongyong Song, Rene L Cruz, and Bhaskar D Rao. Downlink optimization of indoor wireless networks using multiple antenna systems. In *IEEE INFOCOM'04*.
- [86] Angus E Taylor. L'hospital's rule. *The American Mathematical Monthly*, 59(1):20–24, 1952.
- [87] David John Urban. Application of 4k-qam, ldpc and ofdm for gbps data rates over hfc plant. *NCTA Technical Papers*, 2013.
- [88] Christina Vlachou, Albert Banchs, Julien Herzen, and Patrick 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*, CoNEXT '14. ACM, 2014.
- [89] Christina Vlachou, Albert Banchs, Julien Herzen, and Patrick Thiran. On the mac for power-line communications: Modeling assumptions and performance tradeoffs. In *Network Protocols (ICNP), 2014 IEEE 22nd International Conference on*, pages 456–467. IEEE, 2014.

- [90] Christina Vlachou, Sébastien Henri, and Patrick Thiran. Electri-fi your data: Measuring and combining power-line communications with wifi. In *Proceedings of the 2015 Internet Measurement Conference, IMC '15*. ACM, 2015.
- [91] Christina Vlachou, Julien Herzen, and Patrick Thiran. Fairness of mac protocols: Ieee 1901 vs. 802.11. In *2013 IEEE 17th International Symposium on Power Line Communications and Its Applications*, pages 58–63. IEEE, 2013.
- [92] Xili Wan, Xinjie Guan, Wentian Zhao, Guangwei Bai, and Baek-Young Choi. Bandwidth cost minimization via user association for enterprise wlans. *Sensors*, 18:4104, 11 2018.
- [93] Wikipedia. Hungarian algorithm — Wikipedia, the free encyclopedia. [https://en.wikipedia.org/wiki/Hungarian\\_algorithm](https://en.wikipedia.org/wiki/Hungarian_algorithm), 2020.
- [94] Weiye Xu, Benben Wen, Min Lin, and Xiangbin Yu. Energy-efficient power allocation scheme for distributed antenna system over spatially correlated rayleigh channels. *IET Communications*'18.
- [95] Qiaoyang Ye, Beiyu Rong, Yudong Chen, Mazin Al-Shalash, Constantine Caramanis, and Jeffrey G Andrews. User association for load balancing in heterogeneous cellular networks. *IEEE Transactions on Wireless Communications*, 12(6):2706–2716, 2013.
- [96] Vivek Yenamandra and Kannan Srinivasan. Vidyut: exploiting power line infrastructure for enterprise wireless networks. *ACM SIGCOMM'14*.
- [97] Vivek Yenamandra and Kannan Srinivasan. Vidyut: Exploiting power line infrastructure for enterprise wireless networks. In *Proceedings of the 2014 ACM Conference on SIGCOMM*, SIGCOMM '14. ACM, 2014.
- [98] Larry Yonge, Jose Abad, Kaywan Afkhamie, Lorenzo Guerrieri, Srinivas Katar, Hidayat Lioe, Pascal Pagani, Raffaele Riva, Daniel M Schneider, and Andreas Schwager. An overview of the homeplug av2 technology. *Journal of Electrical and Computer Engineering*, 2013, 2013.
- [99] Nidal Zarifeh, Mai Alissa, Theo Kreul, and Thomas Kaiser. Antenna selection performance of distributed antenna systems in full-duplex indoor base station. In *IEEE GeMiC'19*.
- [100] Yan Zhang, Honglin Hu, and Jijun Luo. *Distributed antenna systems: open architecture for future wireless communications*. CRC Press, 2007.
- [101] Huiling Zhu. Performance comparison between distributed antenna and microcellular systems. *IEEE Journal on Selected Areas in Communications*, 2011.
- [102] Manfred Zimmermann and Klaus Dostert. Analysis and modeling of impulsive noise in broad-band powerline communications. *IEEE transactions on Electromagnetic compatibility*, 2002.



- [103] Manfred Zimmermann and Klaus Dostert. A multipath model for the powerline channel. *IEEE Transactions on communications*, 2002.