Priza: Throughput-Efficient DAS Clustering of WiFi-PLC Extenders in Enterprises

Hisham A. Alhulayyil, Jiasi Chen, Karthikeyan Sundaresan, *Fellow, IEEE*, and Srikanth V. Krishnamurthy, *Fellow, IEEE*

***Abstract*— WiFi-enabled Power Line Communications (PLC) range extenders 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, where many extenders contend for a share of access to the backhaul network comprising of the electrical wiring (power lines), thereby drastically impacting any gains from using these extenders on the wireless part of the network. In this paper, we address this issue by developing a framework, Priza, for clustering the WiFi-PLC extenders to intelligently form a DAS (distributed antenna system) to mitigate the inefficiency of sharing the PLC backhaul. By appropriately managing clustering and reuse, Priza improves the PLC backhaul sharing, while at the same time, harnessing the power pooling and diversity gains from 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 the best 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.**

***Index Terms*— Powerline communication (PLC), WiFi-PLC extenders, distributed antenna system, DAS clustering, throughput maximization, frequency assignment.**

I. Introduction

POWERLINE communication (PLC) extenders, that have recently gained popularity [1], [2], 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 primary router. The router connects this 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 access point (AP) and can offer good signal quality and throughput to a user in an area occluded from the primary 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. 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 [3], [4], and [5]. Thus, as the number of extenders increases, 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 extender being closer to the client.

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 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., [6]). Our vision is to group contending PLC extenders into DAS clusters, to improve the sharing of the PLC capacity, while ensuring good wireless signal quality.

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 individual extenders (we are the first to identify the benefit of DAS in mitigating PLC backhaul contention). We note that the additional advantages of WiFi link stability, DAS now 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

Manuscript received 5 February 2023; revised 20 November 2023, 8 April 2024, and 3 August 2024; accepted 10 August 2024. Date of publication 26 August 2024; date of current version 13 November 2024. This work was supported in part by the NSF NeTS under Grant 1528095. The associate editor coordinating the review of this article and approving it for publication was Z. Guan. *(Corresponding author: Hisham A. Alhulayyil.)*

Hisham A. Alhulayyil is with the Department of Computer Science and Engineering, University of California at Riverside, Riverside, CA 92521 USA, and also with the College of Computer and Information Sciences, Information Systems Department, Imam Mohammad Ibn Saud Islamic University, Riyadh 11432, Saudi Arabia (e-mail: halhu001@ucr.edu).

Jiasi Chen is with the Department of Electrical and Computer Engineering (ECE), University of Michigan, Ann Arbor, MI 48109 USA (e-mail: jiasi@umich.edu).

Karthikeyan Sundaresan is with the School of Electrical and Computer Engineering, Georgia Tech Atlanta, GA 30332 USA (e-mail: karthik@ece.gatech.edu).

Srikanth V. Krishnamurthy is with the Department of Computer Science and Engineering, University of California at Riverside, Riverside, CA 92521 USA (e-mail: krish@cs.ucr.edu).

Color versions of one or more figures in this article are available at https://doi.org/10.1109/TWC.2024.3445720.

Digital Object Identifier 10.1109/TWC.2024.3445720

enterprise settings a non-trivial task. First, one must be careful not to significantly compromise reuse by creating very large DAS clusters; uninformed formation of DAS clusters can cause a degradation in throughput compared to when no DAS is implemented, due to poorly utilized wireless frequency bands. Second, because the PLC capacities to the different extenders can vary [3], [4] (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 vanilla case where users are attached to individual extenders. 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 real software-defined radios (SDRs) and PLC extenders as described in section § III. This study not only helps understand the aforementioned factors, but also sheds light on how the PLC capacities are spatially distributed; specifically, we see that PLC capacities cannot be inferred simply based on relative extender locations, such as proximity.

Guided by the understanding gained by our measurement study, we design a framework that we call Priza (electric outlet in Greek), to adaptively assign WiFi frequencies to WiFi-PLC extenders and subsequently group them into DAS cells with the goal of increasing the aggregate throughput. In brief, in a first step 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 band. In a subsequent second step, 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 fewer DAS transmitters), thereby boosting the PLC time share for each. An informed DAS cluster construction can drastically alleviate the inefficient sharing of the PLC backhaul capacity.[1] 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. With Priza, the DAS cells are 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 extenders with significantly different propagation delays to the user do not perform a joint DAS transmission. The algorithm within Priza, which is guided by key observations from our offline measurements, 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.

[1] Later in §VI-D 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.

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. These insights in turn guide the design of Priza.
- 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 via emulations on a real testbed with comparisons with three baselines; we show 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 a 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. We also show that Priza does not bring about any adverse effects on co-existing access points that are not part of the enterprise under consideration.

II. Related Work

In this section, we review relevant related work. We point out that there have been no efforts, to the best of our knowledge, that consider the use of DAS with PLC in enterprise settings as we do (the only prior work on DAS is our own work in a home setting [5], which we describe later).

A. PLC

The authors of [7] and [8] studied the PLC backhaul medium access control mechanism. Their focus was on modeling and understanding how the PLC medium is accessed by the PLC extenders. In contrast, our paper aims to maximize the end-to-end throughput over the concatenated WiFi-PLC link by cleverly clustering WiFi-PLC extenders to form DAS clusters. In [9], the authors consider multi-hop transmissions over power lines; mutliple repeaters use distributed space time codes to perform joint transmissions to boost combining performance. However, this work does not consider a hybrid PLC-wireless setting.

B. Hybrid WiFi-PLC

There is work on applications that use hybrid WiFi-PLC networks; the authors of [10] develop a synchronization mechanism that uses electrical wiring for MIMO (multi-input multi-output) access points. Other papers *e.g.*, [11] consider maximizing the throughput of hybrid WiFi-PLC networks by exploiting multipath routes in the PLC backhaul. Some papers such as [12] and [13] discuss architectures for hybrid heterogeneous networks that include PLC in the context of IoT and e-health, but do not consider DAS. Moreover, these efforts do not consider the impact of the PLC backhaul being the bottleneck. In [4], our prior work, 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.

C. Distributed Antenna Systems

There are efforts such as [14] and [15] that use DAS for increasing the throughput via dynamic clustering or transmission power allocations. In [16], the authors consider a shared UE equipment side DAS to boost the quality of downlink communications, by joint usage of licensed and unlicensed bands. In [17], the authors analytically compute the asymptotic throughput achievable in a massive DAS system. The authors of [18], use DAS as one of the modes in a system that uses multiple MIMO modes depending on user mobility (static versus dynamic), to give the best performance for that category of mobility. None of the above efforts however, consider a PLC backhaul. Similarly, the authors of [19] investigated the usage of DAS over fiber optic infrastructure. However, they do not account for backhaul links with different and often low capacities like PLC. In our prior work in [5], 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 i.e., there was a single client to which a DAS transmission was targeted. In contrast, in this work, we consider DAS in enterprise settings, and how such an approach can drastically decrease contention across frequency bands used by the extenders, thereby retaining the benefits of frequency reuse. Recent work [20], [21] explored PLC plus MIMO clients, which is different than the PLC plus DAS scenario we consider in this work.

D. WiFi Throughput Maximization

Papers like [22] and [23] (among others) propose algorithms to maximize throughput in WiFi networks. However, they all assume a high capacity backhaul such as Ethernet or fiber.

III. Measurement Study

In this section, we first ask if simply clustering nearby extenders can suffice in forming effective DAS clusters. To this end, we perform 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.

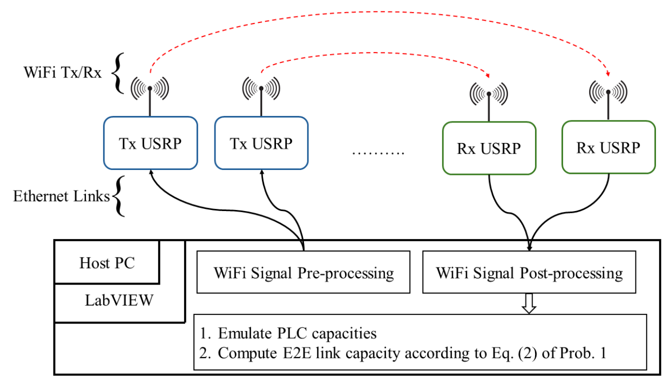

Fig. 1. LabVIEW Experimental Topology.

A. Experimental Setup

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

1) 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 primary 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. 2.

2) DAS Setup: For the DAS setup, we use six NI USRP-N210 radios [24] and an OctoClock CDA-2990 [25]. The radios are connected to the clock over SMA (SubMiniature version A) 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 [5]. 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 ($Tx = \{Tx1, Tx2, Tx3\}$) and the remaining three are the users or receivers ($U = \{U1, U2, U3\}$). 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 Lab 3 as to reflect the case of real WiFi-PLC extenders' locations (see the area shaded blue in Fig. 2). User antennas (U) are placed as shown. A Lenovo T460p running Windows 10 64 bit

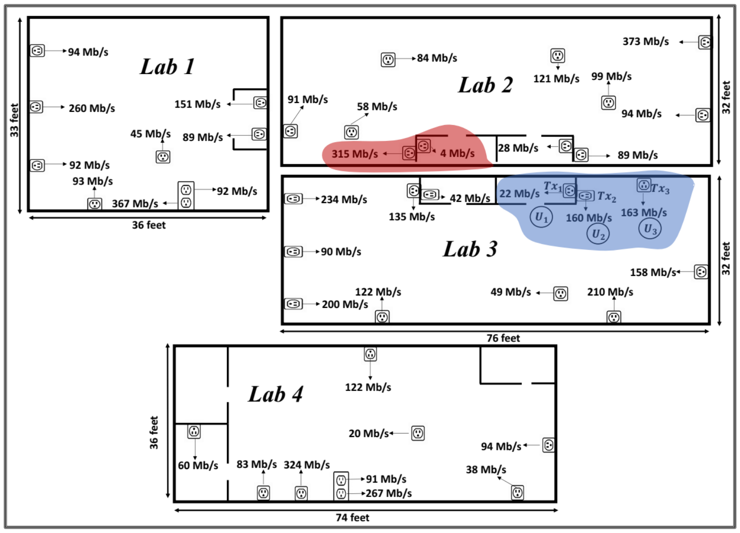

Fig. 2. PLC capacities in different labs in our university setting. The red and blue shaded areas show that a PLC extender with poor capacity can exist next to a PLC extender(s) with good capacity.

is used as the central server, connected to the same switch over an Ethernet cable; its role is to synthesize the signal and pass it to the Tx radios. The Tx radios transmit the signal (the transmissions are based on the 802.11n standard) 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 WiFi throughput for each user.

3) 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 [26], 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 as shown in Fig. 1. 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. 2.

B. Distribution of PLC Capacities

The PLC link capacities 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 [3], [27]. Unfortunately, these characteristics, and their impact on the end-to-end capacity of each PLC link are often not easily available and can be considered opaque. However, one might ask if these factors influence nearby PLC-extenders similarly. The goal of this experiment is to understand if this is the case, and if nearby WiFi-PLC extenders can be grouped into a common DAS cluster, trivially. When a DAS cluster is to be formed, all the extenders are treated as one multicast group [5], and the data source sends the data on the PLC backhaul at once, to all extenders in the associated DAS cluster. However, the data is delivered to the different extenders in a cluster at different times due to the discrepancies in the PLC capacities to those extenders. Thus, the extenders with the good PLC links will have to await laggard extenders with poor PLC links to receive the data on the PLC backhaul before performing their synchronized DAS transmissions. *As a consequence, if extenders with poor PLC capacities are clustered with extenders that have good PLC links, the backhaul throughput of the whole cluster degrades to the throughput of the poorest PLC extender in that cell.*

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* [28]. *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 [3], [4], [5], 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.

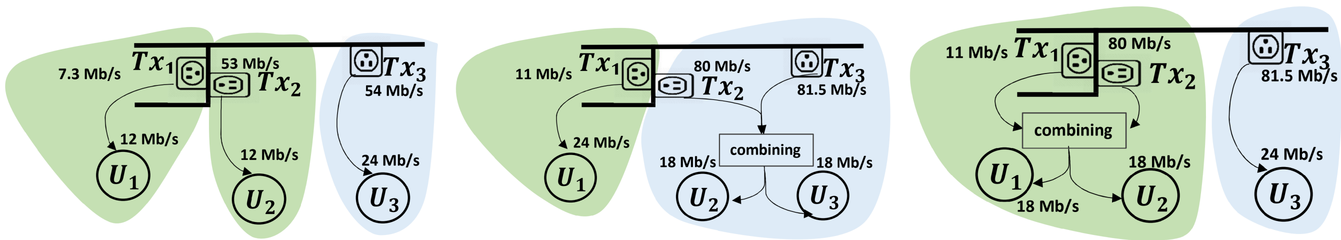

(a) The network setup in reuse mode. U_1 and U_2 could achieve WiFi data rate of 24 Mbps in isolation. Since Tx_1 and Tx_2 share the same frequency, U_1 and U_2's WiFi throughputs become 12 Mbps each.

(b) The network setup with an informed DAS clustering. U_2 and U_3 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.

Fig. 3. Snapshots of the area surrounded by solid line in lab 3 from Fig.2 with three different network setups.

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 those with good capacities as shown in Fig. 2 (e.g., the outlet with a PLC capacity of 315 Mbps is right next to one with 4 Mbps in Lab 2 in the red shaded area). 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.

C. The Need for Informed WiFi-PLC DAS Clustering

The observation that proximity does not translate to similar PLC capacities exemplifies the importance of making informed decisions when clustering WiFi-PLC extenders into a DAS cluster. While as shown previously [18], DAS can indeed improve the system capacity and WiFi link robustness, blindly grouping WiFi-PLC extenders can lead to poor network throughput. Specifically, we show with simple experiments that (a) DAS can *potentially* 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) to showcase two connection modes as shown in Fig. 3. 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. We consider the extenders shown in Fig. 3a from Lab 3. We focus on downloads only and heavy loads (saturated links) since these are the cases where handling excess contention will be an issue. Every lab we consider is very similar in terms of the setup (e.g., Lab 3 is no different than Lab 1). With each experiment, the throughputs of the connections are estimated over a period of time to increase the accuracy of our measurements. While we provide representative results from Lab 3, limited experiments conducted in different labs yielded similar results to those conducted in Lab 3 (thereby ratifying the results).

First, if DAS is not used i.e., all extenders use the reuse mode only (the default configurations of today). 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. 3c), 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. 3b), we see how DAS, if constructed in an informed way, can help increase the aggregate throughput. 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.

1) 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. 3). Specifically, we assign Tx_1 and Tx_2 the same frequency and, Tx_3 a different frequency. We set all the three antennas to the reuse mode; thus, the network contains three reuse cells as shown in Fig. 3a, 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 throughputs 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 [4], [5] [3] 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).

2) Uninformed DAS Clustering Can Hurt Capacity: 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. 3c (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 discussed 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), the end-to-end throughput for $U3$ is equal to its WiFi (bottleneck) 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)).

3) Reuse Compromise: One way to eliminate the PLC sharing entirely is to cluster all three extenders together (configuration not shown) as one (PLC) multicast group. In this scenario, only one frequency (out of two) is used and assigned to the cluster, and the other frequency is left unutilized. This compromises the benefits of having multiple simultaneous transmissions enabled by frequency 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 to realize the potential gains that DAS offers in practice. 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 may offset any DAS gains. Importantly, a naive, greedy clustering *of extenders into DAS clusters, would leave frequencies under-utilized and thus, compromise the acheivable overall capacity.*

IV. 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 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 §III. We find unfortunately that the resulting integer problem is generally NP-hard and given this, we design greedy solution guided by our measurement study, that is of polynomial time complexity.

A. 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 [29]. Each DAS cell (cluster) receives data from the central controller via a multicast. When all of a cell's extenders have the data, they perform a synchronized transmission to the users over the WiFi interface.

Our objective is to ensure that we maximize the achievable throughput by eliminating undesirable PLC contention that can arise in dense deployments of WiFi-PLC extenders. To ensure that we do not compromise reuse (guided by the last experiment in § III), 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 1 below. The notations used in our formulation are tabulated in Table I.

TABLE I
Table of Notation

Variable	Description
A	The set of WiFi-PLC extenders (and hence possible cells).
c_j	The PLC capacity of extender j.
$f(.)$	A function that takes the SNR value in dB and returns the corresponding WiFi modulation rate.
NC	The set of WiFi channels/frequencies.
snr^i_j	The SNR value experienced by user i from extender j.
p_{ik}	The throughput of the PLC backhaul for user i from cell k.
R	The number of cells with assigned extenders.
U	The set of users.
U_k	The set of users associated with cell k. U_k form a partition of U.
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.
$\bar{z}_{jj'}$	A binary variable that indicates whether extender j and j' interfere with each other.
$z_{kk'}$	A binary variable that indicates whether cell k and k' interfere with each other.
$\bar{\gamma}_{jj'}$	A binary variable that indicates whether the signal propagation delays between extenders j and j' is greater than 600 ns.
γ_k	A binary variable that indicates whether all extenders in cell k have delay less than 600 ns.
Ω_{qk}	A binary variable which indicates whether WiFi channel q is assigned to cell k.
v_{ik}	The achievable end-to-end throughput of user i from cell k.

Problem 1: WiFi-PLC Clustering

$$\max_{y_{jk},\Omega_{qk}} \sum_{i,k} v_{ik} \tag{1}$$

s.t.

$$v_{ik} = \min(w_{ik}, p_{ik}) x_{ik}, \quad \forall i \in U \tag{2}$$

$$p_{ik} = \frac{\min_{\{j:y_{jk}=1\}} c_j}{R \sum_{i' \in U_k} x_{i'k}} \quad \forall i \in U, \forall k \in A \tag{3}$$

$$y_{jk} = \begin{cases} 1 & \text{if extender } j \text{ is in cell } k \\ 0 & \text{otherwise} \end{cases} \quad \forall j,k \in A \tag{4}$$

$$\sum_{j=1}^{|A|} y_{jk} = 1 \quad \forall k \in A \tag{5}$$

$$w_{ik} = \begin{cases} \dfrac{\Omega_{qk}}{\sum_{i \in U_k} \frac{1}{f(snr^j_i)}}, & \text{if } \sum_{j=1}^{|A|} y_{jk} = 1, \forall i \in U, \forall k \in A \\ \dfrac{D(i,j,k,y)\Omega_{qk}}{\sum_{i \in U_k} x_{ik}} \gamma_k, & \text{if } \sum_{j=1}^{|A|} y_{jk} > 1, \ \forall i \in U, \forall k \in A \\ 0, & \text{otherwise} \end{cases} \tag{6}$$

$$\bar{\gamma}_{jj'} = \begin{cases} 1 & \text{if the delay between extenders } j \text{ and } j' > 600 \text{ ns} \\ 0 & \text{if the delay between extenders } j \text{ and } j' \leq 600 \text{ ns} \end{cases} \quad \forall j, j' \in A \tag{7}$$

$$\gamma_k = \begin{cases} 1 & \text{if } \sum_{j,j'} \bar{\gamma}_{jj'} y_{jk} y_{j'k} = 0 \\ 0 & \text{otherwise} \end{cases} \quad \forall k \in A \tag{8}$$

$$x_{ik} = \sum_j \bar{x}_{ij} y_{jk} \quad \forall i \in U, k \in A \tag{9}$$

$$R = \sum_k \mathbb{1}_{\sum_j y_{jk} \geq 1} \tag{10}$$

$$\bar{z}_{jj'} = \begin{cases} 1 & \text{if extenders } j \text{ and } j' \text{ are neighbors} \\ 0 & \text{otherwise} \end{cases} \quad \forall j, j' \in A \tag{11}$$

$$z_{kk'} = \begin{cases} 1 & \text{if } \sum_{j,j'} \bar{z}_{jj'} y_{jk} y_{j'k'} > 0 \\ 0 & \text{otherwise} \end{cases} \quad \forall k, k' \in A \tag{12}$$

$$\Omega_{qk} + \Omega_{qk'} \leq 1 \quad \forall \{k, k' \in A : z_{kk'} = 1\}, q \in NC \tag{13}$$

$$\Omega_{qk} = \begin{cases} 1 & \text{if frequency } q \text{ is assigned to cell } k. \\ 0 & \text{otherwise} \end{cases} \quad \forall k \in A \; \forall q \in NC \tag{14}$$

$$\sum_{q=1}^{|NC|} \Omega_{qk} = 1 \quad \forall k \in A \tag{15}$$

The objective (1) is to maximize the summation of all users' (i) throughputs across all cells (k) in the network. Constraint (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 (3) which is the minimum PLC link capacity in the cell (the bottleneck PLC capacity dictates the capacity of a DAS cluster) divided by the number of cells in the system and the number of users in that cell. 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 [3], [5]; thus, the PLC throughput achieved in each cell is inversely proportional to the number of the cells R in the system. In constraint (4), y_{jk} is a decision variable which is equal to one when extender j is assigned to cell k and 0 otherwise. Constraint (5) ensures that each extender is assigned to one and only one cell. Constraint (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 [30]. The second case defines the WiFi throughput user i enjoys if it is connected to DAS cell k. The function $D(.)$ in the numerator is given by:

$$D(i,j,k,y) = f(10 \log(\sum_{j=1}^{|A|} (10^{(snr_j^i/10)} y_{jk}))) \tag{16}$$

where, the $f(.)$ function in Eqn. (16) is one that takes the resulting DAS SNR as the input and converts it to the achievable throughput. Each DAS SNR value range corresponds to a modulation scheme and each modulation scheme encodes a specific number of bits in each OFDM symbol which is then used to compute the DAS throughput achievable upon combining the signals from the various transmitters. The mapping from the DAS SNR to a modulation scheme is consistent with the 802.11n standard. In our prior work [5], we have shown that this equation faithfully reflects what is achieved experimentally via comparisons with testbed experiments. We find that [31] reports the same observation as we do here. Constraint (7) is a system parameter that indicates whether the propagation delay difference between receptions from two extenders on the wireless medium, is less than or equal to 600 ns. Constraint (8) indicates whether all extenders in a cell have delay less than or equal to 600 ns; our measurements (also reported in [5]) show that if the propagation delay differences between transmissions from the antennas (extenders) in a DAS cluster, exceed this value, DAS combining will fail. Constraint (9) requires each user to associate with the cell to which its primary extender belongs. Constraint (10) quantifies the number of cells that time share the PLC backhaul. Constraint (11) defines a parameter to indicate whether two extenders would interfere in the WiFi domain if they were to be assigned the same frequency. Constraint (12) indicates whether two cells interfere with each other, i.e., if any two extenders across the cells interfere with each other. Constraint (13) ensures that no two interfering cells (in the WiFi domain) have the same frequency. Finally, constraint (14) specifies what frequency is assigned to each cell and (15) states that each cell must not have more than one assigned frequency. All flows in Problem 1 are assumed to be saturated to reflect the worst case scenario.

B. Time Complexity

Problem 1 is an integer program, which is generally NP-hard [32]. If P≠NP, then NP-hard problems cannot be optimally solved in polynomial time, motivating our heuristic solution below.

V. 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. The design of Prizais guided by our measurement study in § III. Specfcally, we seek to (a) maximize reuse without creating conflicts and (b) at the same time prevent forming DAS clusters wherein the extenders have diverse link qualities on the PLC backhaul (our experiments showed that this would cause a degradation in throughput). Priza's workflow is shown in Fig. 4 and runs in two stages: (i) it assigns frequencies to a subset of the extenders (antennas) such that there is no interference 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 (we leave frequencies unutilized). In §VI-D, we discuss two approaches where DAS clusters are formed first, based on naive policies, and show that they underperform Priza.

Priza groups PLC extenders based on their PLC capacities. Formally, let α be the number of available frequencies. Then, we set the number of these PLC groups to α. If c_j represents the PLC capacities to the various extenders, then, we compute

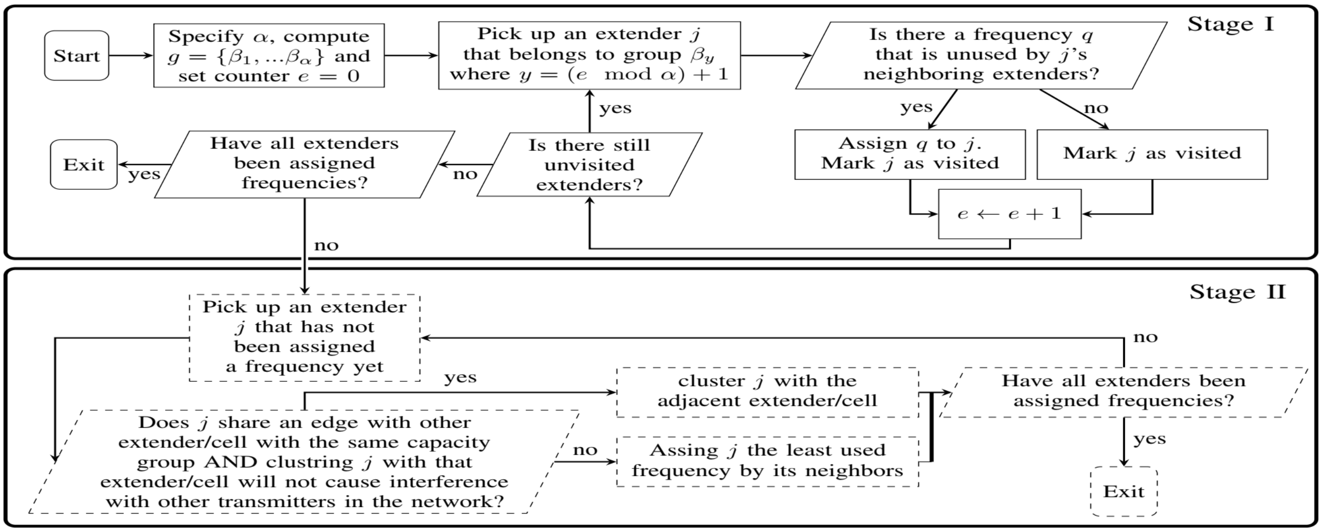

Fig. 4. A schematic showing our Priza approach for clustering the extenders into DAS clusters and assigning frequencies.

a set of intervals g, given by $g = \{\beta_1, \beta_1, \ldots, \beta_\alpha\}$. Each "capacity group interval," β_n is of width:

$$\beta_n = \left((n-1)\frac{\max(c_j)}{\alpha}, n\frac{\max(c_j)}{\alpha}\right], \forall n \in [1, 2, \ldots, \alpha]. \tag{17}$$

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 $c_i \in \beta_l$, ($\in \{1, \alpha\}$) then it is tagged as a part of the l^{th} group. The idea here is to group the extenders which have the similar PLC capacities within the same group. The extent of similarity (how close are the capacities?) is dictated by the parameter α, which is empirically chosen to trade-off reuse and DAS gains.

Assignment of frequencies to cells has been shown to be an NP-complete problem by itself [33]. Thus, to assign frequencies, Priza maps the extenders on to an interference graph (as is typically done [34]). Each extender is represented by a node in the graph and the edges are added between a 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 [34], then these two extenders are considered to interfere with each other and an edge is added to the interference graph between them. To this end, Priza is now able to make frequency assignment and clustering decisions. Next, we describe its two stages.

A. Stage I

First, Priza's algorithm, running on the central controller, chooses extenders greedily from the capacity groups, such that the chosen extender shares the minimum number of links (if possible no links) with any previously chosen extender. Since the number of frequencies is much smaller than the number of links, we find that it is possible to assign frequencies such that all of the extenders in this step have exclusive frequencies. Next, Priza greedily iterates over the capacity groups (in the same order in which the original selection was done, i.e., choose one extender from each capacity group) and tries to assign an exclusive frequency (one that is not assigned to any of its neighbors). The idea here is to assign unique frequencies to extenders that differ in terms of their PLC capacities (this policy is driven by our understanding from our measurements). If an extender cannot be assigned a frequency different from its neighboring extenders, then it is tagged as 'visited' and, subsequently, left to be handled in Stage II. The goal of Stage I is to achieve as high a spatial spectral reuse as possible (note the problem is NP-hard and thus, we cannot achieve optimal reuse). Stage I also ensures 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 of being clustered with extenders that are in the same capacity group (in Stage II), thereby alleviating PLC capacity discrepancies within each group. It is easy to see that without our policy, in an extreme case, if we were to assign exclusive frequencies to all extenders with capacities in β_n, then we may end up with a large set of extenders from the other (multiple) 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 β_n.

B. 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 an 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 a frequency) shares an edge with another cell which is in the same capacity group, and clustering j with that cell does not create interference with other cells in the system (due to DAS power pooling [18]), 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.

C. Algorithm Complexity

Stage I of the above algorithm runs in $\mathcal{O}(|A|)$ where $|A|$ is the number of extenders. This is because the algorithm visits each extender a constant number of times (once). The runtime of Stage II is $\mathcal{O}(|A| - m)$ where m is the number of extenders with assigned frequencies in Stage I. This again is a direct result of the fact that the algorithm only considers unassigned extenders, one at a time. Hence, the total runtime for Algorithm 1 is $\mathcal{O}(|A|)$, i.e., it has a linear time complexity in the number of extenders.

VI. Implementation and 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.

A. Implementation Details

1) Testbed Configuration and Equipment: Our testbed and configurations were described earlier in §III. To reiterate in brief, our experiments are with five radios as Tx antennas and three radios as users. These radios all perform transmissions as per the 802.11n standard [35]. Our testbed is configured to use the 2.4 GHz band. We believe that using 5GHz band would make no difference in the behavioral nature of the results for the following reasons: (a) while the 2.4 band supports lower data rates than the 5GHz, this will not affect our result qualitatively, since the increase in the data rate of the band will yield higher throughput for Priza and other baselines in a similar way. In other words, the performance of both Priza and the other baselines will increase due to the better data rates supported on the 5GHz band (b) the 5GHz band is more prone to interference and it has lower ability to penetrate walls; however, this will impact both Priza and the other baselines in similar ways. The approach with Priza will account for these when forming DAS clusters (i.e., if two nodes cannot reach each other due to walls, they will be able to reuse frequencies and thus not be part of the same DAS cluster) (c) the 2.4GHz is preferable in environments with obstacles such as walls, desks and chairs [36], [37]; thus, this is more likely to be used when both bands are available in enterprise settings, which are considered in this paper. The Tx antennas are assigned PLC capacities to emulate a real PLC backhaul. 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) to ensure that we have adequate reuse cases.

2) Implementation: Priza is implemented in Labview as a user space utility that runs on the server. Five USRP radios are designated as Tx antennas and the other three are users. Priza first builds the interference graph (as described in §V). 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 [34] then they are marked as interfering. The emulation of the PLC links and integration of WiFi-PLC link segments is as described in §III.

B. Simulation Platform

To evaluate Priza in larger settings, we implement Priza entirely in Matlab [38]. 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. To estimate the powerline capacity, we used iperf [28]. iperf is a network tool that can stream saturated traffic between two network endpoints and calculates the capacity of the link based on the amount of traffic received successfully by the receiving point. In our PLC estimations, we used TCP traffic and the process of the estimation lasted for 60 seconds for each PLC link, where the capacity of the PLC link was computed each second (for 60 seconds) and then averaged over those 60 measurement points. Both the extenders and the users are randomly and geographically distributed in a 2D-plane with an area of 108×148 square feet commensurate with the area of six densely packed real labs with several outlets (similar to that in Fig. 2). 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 [39]. The resulting DAS SNR is estimated using Eqn. (16) (also reported by [5] and [40]). Each SNR value is mapped to an appropriate modulation scheme as in [41]. Each modulation scheme is designed to encode the appropriate number of bits within each OFDM symbol [42]. To capture the impact that can be caused by the interfering frequencies on the WiFi domain, we implemented a carrier sensing mechanism consistent with 802.11n. Essentially, if two cells share the same frequency and they are in the same contention domain, then they will have to share the air. Two cells are considered to be in the same contention domain if they both have the same frequency and the signal sensed from each other is equal to or higher than 4 db [34]. When two cells are adjacent to each other (SNR >=4 db), in the case of reuse, the transmitters in the simulation will act similar to conventional APs and try to pick up different frequencies if it is possible. If finding different frequencies is impossible, then the simulated cells will try to find the least congested frequency and use it. A similar approach is followed for DAS, wherein even if one of the transmitters senses a specific frequency, the DAS cell is precluded from using that frequency. The simulation results reported in this work is averaged over 100 trials. Where in each trial, the users locations are randomized and the PLC capacities are randomly chosen from a pool of offline real estimated PLC capacities.

1) Scope: Because the USRP radios only support link layer functions, our studies are exclusively at that layer. For the PLC we estimate the capacity using iPerf, but then use those capacities in our emulation. All throughputs that are reported are link layer throughputs. Unfortunately our platforms do not support IP/transport layer. To be consistent with our experiments, we follow the same principle in our simulations to compute the link layer throughputs (error free packets (bits) delivered at the link layer per unit time). We believe that increased link layer throughputs will naturally manifest into higher network and transport layer throughputs.

C. Baselines & Performance Metrics

1) 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. The size of the cluster is chosen based on our empirical analysis. When the size of the cluster increases, we approach the case where the throughput of the users becomes similar to L-DAS. On contrary, reducing the size of the cluster leads to increasing the contention on the WiFi segment of the aggregated WiFi-PLC link and the throughput of the users start to drop to levels similar to the Reuse case. To benchmark Priza fairly against the B-DAS baseline, we find the cluster size that yields the highest end-to-end throughput, and use that size for B-DAS. 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. To the best of our knowledge, we are the first to consider deploying DAS for alleviating contention among densely deployed PLC extenders. Specifically, previous literature lacks any baselines to compare Priza against except reuse, and thus we form our own DAS baselines. These three baselines are representative of the likely spectrum of cluster sizes. The reuse mode is equivalent to using DAS but with only one antenna (or extender) in each cell. For B-DAS, we used different values of balanced assignments. We find that approximately 5 extenders in a group yields the best results. If we increase the size to beyond this, the throughput starts to drop drastically because reuse is affected (hence, we use this number in what we report). As one can see, the L-DAS throughput (which has the largest cluster size with 50 extenders) is drastically lower than that of B-DAS (as we will see later in section VI-E).

2) Performance Metrics: Our performance metrics of interest are (a) the aggregate throughput achieved by Priza, (b) the fairness of the users' individual throughputs and (c) the impact caused by Priza and other baselines on the throughput performance of external co-existing APs in the vicinity of the PLC extenders (in deployment area). We use the commonly used Jain's fairness index [43] to evaluate fairness.

D. Experimental Evaluations

1) Gains in Aggregate Throughputs: We first depict the throughputs with Priza and the baselines in Fig 5a and Fig. 5b. We see that Priza outperforms the other baselines and achieves higher aggregate throughputs. The average aggregate

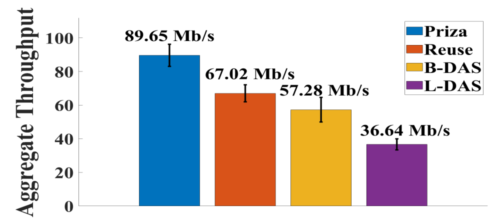

(a) Priza, Reuse, B-DAS and L-DAS comparison on testbed.

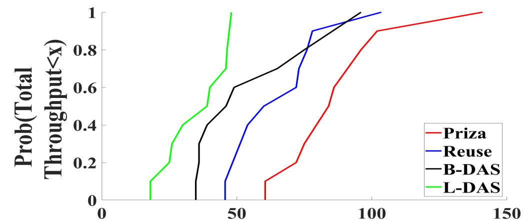

(b) CDF of the aggregate throughputs.

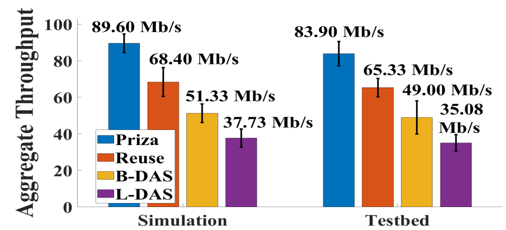

(c) Validating the fidelity of our simulations.

Fig. 5. Experimental results on our testbed to showcase the benefits of Priza, and the consistency of the experiments with our simulations.

(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, and harness the gains from DAS as discussed in § III, unlike the other schemes.

2) 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 simlar 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.

E. Simulation Results

1) Simulation Fidelity: We perform our simulations with ten topologies with the same users and antennas locations

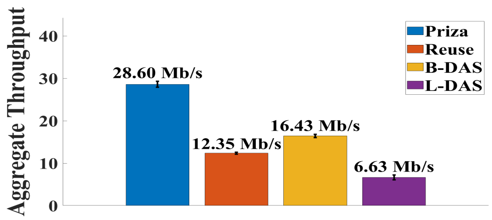

(a) Average aggregate throughputs.

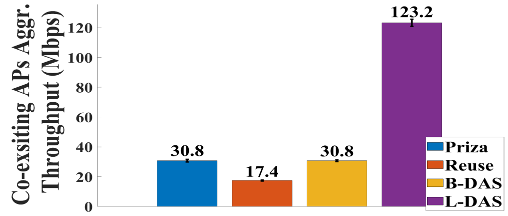

(b) Average aggregate throughput for co-existing APs.

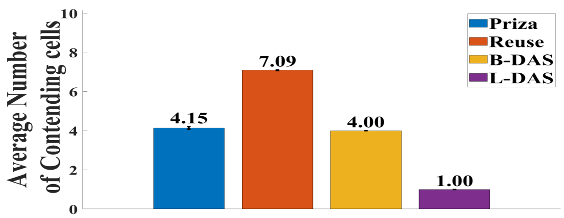

(c) Average Number of contending cells for co-existing APs.

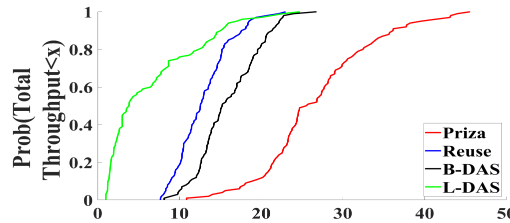

(d) CDF of aggregate throughput.

Fig. 6. Simulation Results showing (a) average aggregate throughputs with various schemes (b) average throughputs accomplished by co-existing APs when present (c) average numbers of cells contending with the co-existing APs and (d) CDF of the throughputs of the various schemes within the enterprise.

TABLE II
SIMULATION JAIN'S FAIRNESS INDEX

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

of our testbed experiments and compare them with our real testbed results. We see from Fig. 5c that the results are very similar, thus attesting to the high fidelity of our simulation models.

2) Aggregate Throughputs: This simulation experiment was as described in § VI-B with the number of available frequencies set to 11 so as to reflect the case of the actual number of WiFi channels in 802.11n [35]. The results of our simulations are in Fig. 6. Fig. 6a 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 outperforms Reuse. We attribute this to the significant increase in extender density and scale. That is, when the topology is small (same size as our testbed), Reuse outperforms B-DAS; but as the network size grows larger, B-DAS starts to improve compared to Reuse and starts to overtake Reuse (results omitted due to space constraints). 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).

3) Fairness: In Table II, 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 II we show the fairness index values for the same users when the PLC backhaul has a very high capacities (e.g, $w_{ik} = \min(w_{ik}, p_{ik})x_{ik}, \quad \forall i \in U$). 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 (6) of Problem 1).

4) Impact on Co-Existing APs: To quantify the impact of the proposed approach and other baselines on co-existing APs that are not part of the enterprise under consideration, two simulations were conducted. For both, we simulated three co-existing APs randomly distributed around the geographical area of interest where the PLC extenders exist. Each AP is then assigned two to four users in each trial (one hundred trials in total). The goal here is investigate the impact of Priza and other baselines on the co-existing APs in terms of (a) these APs' achieved throughputs and (b) the number of cells contending on the same frequency band as the APs. Fig. 6b shows that Priza is better than or at least as good as the best performing baselines, except for L-DAS. The reason the co-existing APs perform much better under L-DAS is that L-DAS groups all PLC extenders into very few cells (usually one large cell), which renders most of the frequency bands unused. These unused frequency bands can then be picked up by the co-existing APs without any need to share them with the deployment of the extenders' network. Furthermore, there is very little co-channel interference experienced, if at all. However, this causes a significant degradation in the throughputs achieved by L-DAS as shown in Fig. 6a. This is because of very low throughputs rendered to the PLC extenders' users due to the sharing of the backhaul and underutilizing the frequencies on the wireless part. In other words, Priza strikes a balance between the PLC extenders' and APs' users throughputs, and it increases the aggregate throughput of the whole heterogeneous network, while simultaneously preserving the throughputs of the co-existing APs as in a balanced DAS scheme. This is directly seen in Fig. 6c, which shows the number of contending channels available for the co-existing APs. The throughputs of the co-existing APs is significantly better with Priza compared to Reuse (Fig. 6b) because of the reduction in contention (which was the goal of Priza); B-DAS has a similar effect on the throughputs of the co-existing APs because it also reduces contention; however, as discussed earlier and shown in Fig. 6a, B-DAS achieves a significantly lower throughput compared to Priza in the DAS network of extenders.

VII. Conclusion and Possible Future Directions

In this paper we consider the dense deployment of plug-and-play PLC based WiFi extenders in enterprises. 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. Priza's approach retains the gains from frequency reuse, but forms DAS clusters of PLC extenders to mitigate the inefficiencies in sharing the PLC backhaul. It also provides the benefits of power pooling and diversity from DAS. 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 both in terms of providing higher throughputs, but also largely retaining the throughputs of co-existing APs in the vicinity.

Future directions. One can visualize extenders as being used to spatially multiplex streams to users instead of using DAS as we do in this work. In fact, DAS and spatial multiplexing could be intelligently used in conjunction in ways to further boost throughput beyond what we have accomplished here. Designing algorithms that bring the benefits of both these schemes to PLC could be a separate research direction. Finally, instead of considering homogeneous frequency bands, channel bonding [44] may be considered, and how frequency assignments to DAS clusters should occur in such cases is also an interesting future direction.

References

[1] *Private Tutoring Global Market Trajectory & Analytics*, Global Ind. Analysts, San Jose, CA, USA, 2020.

[2] *Power Line Communication Market—Global Industry Analysis, Size, Share, Growth, Trends and Forecast 2016–2024*, Transparency Market Res., Pune, India, 2016.

[3] K. Apicharttrisorn, A. O. F. Atya, J. Chen, K. Sundaresan, and S. V. Krishnamurthy, "Enhancing WiFi throughput with PLC extenders: A measurement study," in *Proc. PAM*. Sydney, NSW, Australia: Springer, 2017, pp. 257–269.

[4] H. A. Alhulayyil, K. Apicharttrisorn, J. Chen, K. Sundaresan, S. Oymak, and S. V. Krishnamurthy, "WOLT: Auto-configuration of integrated enterprise PLC-WiFi networks," in *Proc. IEEE ICDCS*, Nov. 2020, pp. 563–573.

[5] H. A. Alhulayyil, J. Chen, K. Sundaresan, and S. V. Krishnamurthy, "Boosting home WiFi throughputs via adaptive DAS clustering of PLC extenders," in *Proc. IEEE 18th Int. Conf. Mobile Ad Hoc Smart Syst. (MASS)*, Oct. 2021, pp. 499–509.

[6] Y. Zhang, H. Hu, and J. Luo, *Distributed Antenna Systems: Open Architecture for Future Wireless Communications*. Boca Raton, FL, USA: CRC Press, 2007.

[7] C. Vlachou, A. Banchs, P. Salvador, J. Herzen, and P. Thiran, "Analysis and enhancement of CSMA/CA with deferral in power-line communications," *IEEE J. Sel. Areas Commun.*, vol. 34, no. 7, pp. 1978–1991, Jul. 2016.

[8] C. Vlachou, A. Banchs, J. Herzen, and P. Thiran, "On the MAC for power-line communications: Modeling assumptions and performance tradeoffs," in *Proc. IEEE 22nd Int. Conf. Netw. Protocols*, Oct. 2014, pp. 456–467.

[9] L. Lampe, R. Schober, and S. Yiu, "Distributed space-time coding for multihop transmission in power line communication networks," *IEEE J. Sel. Areas Commun.*, vol. 24, no. 7, pp. 1389–1400, Jul. 2006.

[10] V. Yenamandra and K. Srinivasan, "Vidyut: Exploiting power line infrastructure for enterprise wireless networks," in *Proc. ACM Conf. SIGCOMM*, Aug. 2014, pp. 595–606.

[11] S. Henri, C. Vlachou, J. Herzen, and P. Thiran, "Empower hybrid networks: Exploiting multiple paths over wireless and electrical mediums," in *Proc. 12th Int. Conf. Emerg. Netw. Exp. Technol.*, Dec. 2016, pp. 51–65.

[12] M. Kashyap, Z. Naaz, N. Bansal, and V. Sharma, "Performance evaluation of IoT architecture for heterogeneous networks," in *Proc. 14th Int. Conf. Comput. Commun. Netw. Technol. (ICCCNT)*, Jul. 2023, pp. 1–8.

[13] S. U. Khan, S. U. Jan, T. Hwang, and I.-S. Koo, "Hybrid Wi-Fi and PLC network for efficient e-health communication in hospitals: A prototype," *Bull. Elect. Eng. Inform.*, vol. 13, no. 2, pp. 1400–1410, 2024.

[14] J. A. Sheikh, H. Haseena, and F. Mustafa, "A new dynamic cooperation cluster empowered cell-free massive MIMO architecture network for future networks," Res. Square, Durham, NC, USA, Tech. Rep., 2023. [Online]. Available: https://www.researchsquare.com/article/rs-3444498/v1

[15] B. Banerjee, R. C. Elliott, W. A. Krzymień, and M. Medra, "Access point clustering in cell-free massive MIMO using conventional and federated multi-agent reinforcement learning," *IEEE Trans. Mach. Learn. Commun. Netw.*, vol. 1, pp. 107–123, 2023.

[16] M. Breiling, D. W. K. Ng, C. Rohde, F. Burkhardt, and R. Schober, "Resource allocation for outdoor-to-indoor multicarrier transmission with shared UE-side distributed antenna systems," in *Proc. IEEE 81st Veh. Technol. Conf. (VTC Spring)*, May 2015, pp. 1–7.

[17] Q. Wang, H. Xiong, S. Yu, and Y. Liu, "Throughput of wireless-powered massive distributed antenna systems over composite fading channels," *J. Commun.*, vol. 12, no. 1, pp. 24–31, Jan. 2017.

[18] S. Singh, K. Sundaresan, S. V. Krishnamurthy, X. Zhang, A. Khojastepour, and S. Rangarajan, "TRINITY: A practical transmitter cooperation framework to handle heterogeneous user profiles in wireless networks," in *Proc. 16th ACM Int. Symp. Mobile Ad Hoc Netw. Comput.*, Jun. 2015, pp. 297–306.

[19] S. Pato, J. Pedro, J. Santos, A. Arsénio, P. Inácio, and P. Monteiro, "On building a distributed antenna system with joint signal processing for next generation wireless access networks: The futon approach," in *Proc. 7th Conf. Telecommun., Portugal*, 2008.

[20] W. Sun, B. Chen, V. Yenamandra, and K. Srinivasan, "Grid MIMO: Exploiting spatiality of power line infrastructure for MIMO," in *Proc. IEEE Int. Conf. Pervasive Comput. Commun. Workshops Affiliated Events (PerCom Workshops)*, Mar. 2023, pp. 416–419.

[21] M. B. Mohajer et al., "On the achievable rate of MIMO narrowband PLC with spatio-temporal correlated noise," 2023, *arXiv:2308.14645*.

[22] E. Khorov, A. Kiryanov, and A. Krotov, "Joint power control and time division to improve spectral efficiency in dense Wi-Fi networks," in *Proc. IEEE Int. Black Sea Conf. Commun. Netw. (BlackSeaCom)*, Jun. 2018, pp. 1–5.

[23] R. Atawia and H. Gacanin, "Self-deployment of future indoor Wi-Fi networks: An artificial intelligence approach," in *Proc. IEEE Global Commun. Conf.*, Dec. 2017, pp. 1–6.

[24] *USRP N210*. Accessed: Jun. 2021. [Online]. Available: https://www.ettus.com/all-products/un210-kit/

[25] *OctoClock CDA-2990*. Accessed: Jun. 2021. [Online]. Available: https://www.ettus.com/all-products/octoclock/

[26] *USRP Hardware Driver and USRP Manual*. [Online]. Available: https://files.ettus.com/manual/page_usrp2.html#usrp2_network

[27] M. Zimmermann and K. Dostert, "A multipath model for the powerline channel," *IEEE Trans. Commun.*, vol. 50, no. 4, pp. 553–559, Apr. 2002.

[28] *TCP/UDP Bandwidth Measurement Tool*. Accessed: Aug. 2021. [Online]. Available: https://iperf.fr/

[29] B. Song, R. L. Cruz, and B. D. Rao, "Downlink optimization of indoor wireless networks using multiple antenna systems," in *Proc. IEEE INFOCOM*, vol. 4, Mar. 2004, pp. 2778–2789.

[30] M. Heusse, F. Rousseau, G. Berger-Sabbatel, and A. Duda, "Performance anomaly of 802.11b," in *Proc. IEEE INFOCOM*, Apr. 2003, pp. 836–843.

[31] D. Brennan, "Linear diversity combining techniques," *Proc. IEEE*, vol. 91, no. 2, pp. 331–356, Feb. 2003.

[32] R. Kannan and C. L. Monma, "On the computational complexity of integer programming problems," in *Optimization and Operations Research*. Bonn, Germany: Univ. of Bonn, 1978, pp. 161–172.

[33] I. Broustis, K. Papagiannaki, S. V. Krishnamurthy, M. Faloutsos, and V. P. Mhatre, "Measurement-driven guidelines for 802.11 WLAN design," *IEEE/ACM Trans. Netw.*, vol. 18, no. 3, pp. 722–735, Jun. 2010.

[34] I. Ramachandran and S. Roy, "WLC46–2: On the impact of clear channel assessment on MAC performance," in *Proc. IEEE Globecom*, Nov. 2006, pp. 1–5.

[35] E. Perahia, "IEEE 802.11n development: History, process, and technology," *IEEE Commun. Mag.*, vol. 46, no. 7, pp. 48–55, Jul. 2008.

[36] A. Qaddus, "An evaluation of 2.4 GHz and 5 GHz ISM radio bands utilization in backhaul IP microwave wireless networks," in *Proc. Int. Conf. Inf. Sci. Commun. Technol. (ICISCT)*, Nov. 2019, pp. 1–5.

[37] A. S. Haron, Z. Mansor, I. Ahmad, and S. M. M. Maharum, "The performance of 2.4 GHz and 5 GHz Wi-Fi router placement for signal strength optimization using altair WinProp," in *Proc. IEEE 7th Int. Conf. Smart Instrum., Meas. Appl. (ICSIMA)*, Aug. 2021, pp. 25–29.

[38] *MATLAB Version 9.6.0.1072779 (R2019a)*, MathWorks, Natick, MA, USA, 2019.

[39] J. Lei, R. Yates, L. Greenstein, and H. Liu, "Wireless link SNR mapping onto an indoor testbed," in *Proc. 1st Int. Conf. Testbeds Res. Infrastruct. Develop. Netw. Communities*, 2005, pp. 130–135.

[40] D. G. Brennan, "Linear diversity combining techniques," *IEEE Proc. IRE*, vol. 47, no. 6, pp. 1075–1102, 1959.

[41] M. Dangi and M. K. Porwal, "Analyses of SNR threshold for minimum BER in various modulations schemes and development of an adaptive modulation scheme," *Int. J. Innov. Sci. Eng. Technol*, vol. 2, no. 3, pp. 139–142, 2015.

[42] A. M. Abdelgader and W. Lenan, "The physical layer of the IEEE 802.11p WAVE communication standard: The specifications and challenges," in *Proc. WCECS*, 2014, pp. 22–24.

[43] R. Jain, A. Durresi, and G. Babic, "Throughput fairness index: An explanation," in *Proc. ATM Forum Contribution*, 1999, pp. 1–13.

[44] M. Y. Arslan et al., "ACORN: An auto-configuration framework for 802.11n WLANs," *IEEE/ACM Trans. Netw.*, vol. 21, no. 3, pp. 896–909, Jun. 2013.

Hisham A. Alhulayyil received the master's degree in information security from the University of Pittsburgh in 2016 and the M.Sc. and Ph.D. degrees in computer science from the University of California at Riverside in 2018 and 2022, respectively. He works as the Vice Dean of information technology and e-learning at Imam Muhammed Ibn Saud Islamic University (IMSIU). Also, he was a Consultant of planning and development for the higher management with IMSIU and an IT Consultant for multiple companies in the industrial sector. He is currently an Assistant Professor with the College of Computer and Information Sciences, IMSIU. He has published multiple articles in different venues in the areas of cyber-security and computer networks.

Jiasi Chen received the B.S. degree from Columbia University and the Ph.D. degree from Princeton University. She is currently an Associate Professor with the Electrical and Computer Engineering (ECE), University of Michigan, Ann Arbor. Previously, she was with the University of California at Riverside. Her projects typically involve some combination of mathematical optimization combined with systems experimentation. Her research interests include wireless networks, multimedia systems, mobile computing, and augmented and virtual reality. She was a recipient of the NSF CAREER Award and the Industry Research Awards from Meta and Adobe.

Karthikeyan (Karthik) Sundaresan (Fellow, IEEE) is currently a Professor with the School of ECE, Georgia Tech. He received business contribution awards for bringing research technology to commercialization in industry at NEC. He also led the spin-out efforts of two innovative, laboratory-grown research technologies for infrastructure-free tracking of first responders in GPS-denied environments and sustainable, and massive scale product tracking in supply chains. He has participated in various organizational and editorial roles for IEEE and ACM conferences and journals. He holds over 60 patents. His research interests include wireless networking and mobile computing, and span both algorithm design and system prototyping. He is a fellow of the National Academy of Inventors and an ACM Distinguished Scientist. He was a recipient of the ACM Sigmobile's Rockstar Award in 2016 for early career contributions to mobile computing and wireless networking and several best paper awards at prestigious ACM and IEEE conferences. He served as the PC Co-Chair for ACM MobiCom'16.

Srikanth V. Krishnamurthy (Fellow, IEEE) received the Ph.D. degree in electrical and computer engineering from the University of California at San Diego in 1997. From 1998 to 2000, he was a Research Staff Scientist with the Information Sciences Laboratory, HRL Laboratories, LLC, Malibu, CA, USA. Currently, he is a Professor of computer science with the University of California at Riverside. His research interests include wireless networks, computer security, and machine learning for security and networking. He is a fellow of AAAS. He was a recipient of the NSF CAREER Award in 2003. He has been the Program Committee Chair for several conferences, including IEEE Infocom, IEEE SECON, and IEEE ICNP. He was the Vice-Chair for ACM MobiCom in 2007. He was the Editor-in-Chief of ACM MC2R and the Associate Editor-in-Chief of IEEE TRANSACTIONS ON MOBILE COMPUTING.