

**PRACTICAL INTERFERENCE TOMOGRAPHY FOR UNLICENSED
SPECTRUM ACCESS: APPLICATIONS IN RESOURCE MANAGEMENT AND
JAMMER LOCALIZATION**

A Dissertation
Presented to
The Academic Faculty

By

Aadesh Madnaik

In Partial Fulfillment
of the Requirements for the Degree
Master of Science in the
School of Electrical and Computer Engineering
College of Engineering

Georgia Institute of Technology

May 2024

© Aadesh Madnaik 2024

**PRACTICAL INTERFERENCE TOMOGRAPHY FOR UNLICENSED
SPECTRUM ACCESS: APPLICATIONS IN RESOURCE MANAGEMENT AND
JAMMER LOCALIZATION**

Thesis committee:

Dr. Karthikeyan Sundaresan
Electrical and Computer Engineering
Georgia Institute of Technology

Dr. Raghupathy Sivakumar
Electrical and Computer Engineering
Georgia Institute of Technology

Dr. Gordon L. Stüber
Electrical and Computer Engineering
Georgia Institute of Technology

Date approved: April 26, 2024

ACKNOWLEDGMENTS

I would like to express my deepest gratitude to my thesis advisor, Prof. Karthikeyan Sundaresan, for his unwavering support and guidance. His expertise, insights, and constructive feedback have been invaluable in shaping this work, and making it possible. Further, I would like to thank the members of my thesis committee, Prof. Gordon L. Stüber and Prof. Raghupathy Sivakumar, for their help in preparation of this work.

Special thanks are due to N. Cameron Matson without whom this work would not have been achievable. Additionally, I would like to thank my fellow labmates, Yu-Tai Lin, and Alan Liu, for their helpful discussions and advice. I would also like to thank all my friends for their moral support and words of encouragement in times of doubt. And most importantly, I am deeply indebted to my family for their unwavering love, encouragement, and understanding throughout this journey. Their patience, support, and belief in me have been my source of strength.

TABLE OF CONTENTS

Acknowledgments	iii
List of Tables	vi
List of Figures	vii
Summary	viii
Chapter 1: Introduction	1
Chapter 2: Background and Prior Art	5
2.1 5G Access	5
2.2 Unlicensed Spectrum	6
2.3 Network Tomography	7
Chapter 3: The Problem of Interference	8
3.1 Impact of Interference	8
3.2 Challenge of Scalable Measurements	10
Chapter 4: Interference Tomography	11
4.1 Direct Estimation Phase	12
4.2 Scaling to High Order Distributions	14

Chapter 5: Applications	22
5.1 Resource Management	22
5.2 Jammer Localization	25
Chapter 6: Evaluation	27
6.1 Evaluation Methodology	27
6.2 Evaluations for Resource Management	28
6.3 Evaluation for Jammer Localization	32
Chapter 7: Conclusions	34
References	36

LIST OF TABLES

6.1	Classification accuracy of the jammer localization system	32
-----	---------------------------------------------------------------------	----

LIST OF FIGURES

3.1	Illustration of hidden terminal interference	9
3.2	HOD prevents drop in performance with a) increased interference, b) larger number of antennas	9
4.1	System design and information flow	11
4.2	Impact of hidden terminals across three channels	13
4.3	Mean squared error between true and estimated HOD from pair-wise marginal inputs vs. latent variable size	18
4.4	Blue-printing interference and localizing jammers: a) True interference map, b) HOD clustering, c) Graph creation, d) ID HTs, e) Estimated zones, f) Refine zones.	20
6.1	a) RB utilization on a single channel SISO system, b) Throughput gains over PF with varying HTs and c) clients.	28
6.2	Resource block utilization per channel and measurement overhead of SISO schedulers.	30
6.3	Normalized total throughput and overhead of MU-MIMO vs. number of antennas.	30
6.4	Effect of the number of clusters used by the system on the performance and measurement overhead.	31
6.5	Accuracy (left) and Precision (right) of the system in localizing jammers with 40 clients and 3 hidden terminals	33

SUMMARY

Mobile networks have increased spectral efficiency through advanced multiplexing strategies that are coordinated by base stations (BS) in licensed spectrum. However, external interference on clients leads to significant performance degradation during dynamic (unlicensed) spectrum access (DSA). We introduce the notion of network tomography for DSA, whereby clients are transformed into spectrum sensors, whose joint access statistics are measured and used to account for interfering sources. Albeit promising, performing such tomography naively incurs an impractical overhead that scales exponentially with the multiplexing order of the strategies deployed – which will only continue to grow with 5G/6G technologies.

To this end, this work proposes a scalable network tomography framework that estimates joint client access statistics with just linear overhead, and forms a blue-print of the interference, thus enabling efficient DSA for future networks. The system’s design incorporates intelligent algorithms that leverage multi-channel diversity and the spatial locality of interference impact on clients to accurately estimate the desired interference statistics from just pair-wise measurements of its clients. The merits of its framework are showcased in the context of resource management and jammer localization applications, where its performance significantly outperforms baseline approaches and closely approximates optimal performance at a scalable overhead.

CHAPTER 1

INTRODUCTION

The evolution of next-generation mobile networks is driven by the escalating demand for data and the need for larger bandwidths to accommodate this exponential growth in traffic. While higher frequencies (millimeter-wave) have the potential for larger bandwidths and higher data rates, they come with several limitations. Millimeter waves are challenged by shorter propagation distances and susceptibility to environmental factors like obstacles and atmospheric conditions.

Despite the advantages of higher frequencies, the spectral bands below 6 GHz, often referred to as sub-6 GHz, remain highly coveted for mobile access. These bands provide broader coverage and penetration characteristics, making them essential for providing reliable connectivity, especially in urban areas. Further, mobile networks (LTE/5G) have evolved to accommodate several sophisticated technologies (e.g. carrier aggregation [1], multi-user MIMO [2]) that increase spectral efficiency by leveraging synchronous and licensed spectrum access. Recognizing the significance of these lower bands, regulatory bodies like the Federal Communications Commission (FCC) have been exploring ways to make more efficient use of them. One approach involves re-purposing underutilized bands, such as federal-owned bands used for naval radars, for mobile access. The Citizens Broadband Radio Service (CBRS) band [3], spanning from 3.5 to 3.7 GHz, is a prime example of a spectrum range being considered for such repurposing efforts.

To enable more equitable use of these lower bands, various flexible models of *spectrum sharing* have been proposed. Fundamentally, the same frequency resource can be used by several competing agents under some regulatory specifications. Based on these model specifics, spectrum sharing can range from lightly-licensed to unlicensed.

As spectrum sharing evolves towards more flexible and dynamic models, interference

becomes a leading factor contributing to a drop in performance. The multiplexing features of mobile networks are highly vulnerable to unknown, external interference (e.g. interference from WiFi and other cellular operators), which is the norm in unlicensed spectrum. It can result in substantial performance degradation (as high as 80% seen in section 3.1). Therefore, for mobile networks to be able to justify democratized access, the exorbitant cost of spectrum ownership needs to be tackled. This makes dynamic spectrum access (DSA) becomes increasingly important. DSA ensures efficient and adaptive allocation of spectral resources in real time, maximizing utilization while minimizing interference and ensuring fair access for all users. A fundamental challenge facing the future of open spectrum access is: *Can we continue to realize mobile network's multiplexing features, maintaining performance comparable to licensed spectrum access, even in unlicensed spectrum and in the presence of unknown interference?*

An important step in addressing such unknown interference is to first estimate it accurately. Spectrum sensing has understandably become a key component of DSA research. Through the help of a sophisticated spectrum scanning device, the base station (BS) aims to obtain a real-time view of spectrum usage. While a spectrum scanning device reveals interference information at the location of the base station, it is unable to provide a comprehensive view of the interference its clients experience. So, estimating interference is not just a spectrum-dependent process; it is equally a (receiver) *location*-dependent one. This has led to the rise of network tomography for DSA whereby the client devices' existing communication interfaces are transformed to double-up as virtual spectrum sensors. Through intelligent scheduling of the clients' channel access and measuring the corresponding outcomes, the BS can decipher the impact of interference on the clients, accounting for both its spectrum and location dependence.

However, measuring the impact of interference on groups of clients is a non-trivial task. To deliver the gains of multiplexing techniques like multi-user MIMO in the presence of interference, one must estimate the joint access distributions of clients whose order de-

depends on the number of clients operating simultaneously. The number of channels (C) and antennas (M) have increased in 5G and 6G, posing a huge scalability challenge—the measurements needed to estimate higher order distributions (HODs) for clients grows exponentially as $O(C \cdot 2^M)$.

To this end, we propose a *scalable network tomography* framework that helps address the scalability challenge in accurately estimating HODs across multiple channels, antennas, and a large number of clients. The proposed system estimates HODs accurately, at just a linear overhead, and generates a blue-print of the interferers, to enable efficient DSA for 5G/6G networks. The HOD serves as a key building block for multiple applications like (i) resource management—enabling mobile networks to retain the benefits of their multiplexing techniques, and (ii) security—allowing defense and private networks to localize potential interfering/jamming sources.

The system’s design incorporates three key elements. First, the system clusters clients based on interference. It leverages the spatial consistency and channel diversity of the interference experienced by clients to cluster clients and choose representatives. A single representative client from a cluster is sufficient for measurements, which helps to significantly reduce the measurement overhead. Second, the system generates HODs from pairwise measurements using recent results in latent variable decomposition. Further, we propose an algorithm to estimate the HODs of arbitrary groups of clients, from just pairwise measurements of the clusters’ representatives. Lastly, we propose a novel algorithm that accurately estimates the number of interfering sources and their corresponding dependencies on clients based on the latent variable decomposition of the HODs.

The proposed system makes two key applications of the interference dependencies of the clients. The first is resource management where HODs are incorporated into mobile network’s scheduling framework to enable interference-aware multi-user access, thereby maximizing the multiplexing gain even in the presence of interference. Further, the system is implemented in NS3 and supplemented with numerical evaluations to better understand

the systems ability to deliver gains in multiplexing (higher order MU-MIMO and carrier aggregation) schemes despite the presence of external interference. The second application is jammer localization. Leveraging the blue-printed interference, the system is able to identify a small set of key clients, whose locations in turn allow us to maximally narrow down the location of the interfering sources to a relatively small area of under $200 m^2$.

CHAPTER 2

BACKGROUND AND PRIOR ART

2.1 5G Access

Cellular networks, such as LTE in 4G and the upcoming NR in 5G, operate as synchronous, scheduled access systems within licensed spectrum bands. In these networks, the base station (BS) plays a central role in orchestrating communication by scheduling both downlink (DL) and uplink (UL) transmissions for clients within its coverage area.

The communication resources in these networks are organized into sub-frames, each comprising two-dimensional resource elements spanning time (symbols) and frequency (sub-carriers). These resource elements are known as resource blocks (RBs) [4]. LTE and 5G-NR utilize Orthogonal Frequency Division Multiple Access (OFDMA) [5] as the access scheme, allowing for efficient allocation of resources to multiple users within each sub-frame. In OFDMA, multiple clients are typically scheduled on different RBs within a sub-frame. However, with the advent of multi-user [6, 7] and massive Multiple Input Multiple Output (MIMO) [8, 9] technologies in 5G systems, multiple clients can also be scheduled on the *same* RB, albeit spatially separated by the use of multiple antennas at the base station. This enables enhanced spectral efficiency and improved user experience through better spatial multiplexing.

The scheduling of DL and UL transmissions is communicated to the clients through the control portion of the DL sub-frames. This ensures that clients are aware of their allocated resources and can efficiently transmit and receive data within the designated time slots and frequency bands. Overall, this synchronous and scheduled approach to access management forms the backbone of efficient and reliable communication in LTE and upcoming 5G networks.

2.2 Unlicensed Spectrum

Operating in unlicensed spectrum poses unique challenges for 5G networks compared to traditional licensed spectrum systems. Unlike the always-on mode characteristic of licensed spectrum operation, 5G in unlicensed spectrum necessitates the adoption of asynchronous access principles, such as clear-channel assessment (CCA) through energy sensing, a mechanism known as Listen-Before-Talk (LBT) [10]. Additionally, back-off mechanisms are employed to ensure co-existence with other technologies like WiFi or other cellular networks that may be operating in the same unlicensed bands. and back-off for co-existence with other technologies (e.g. WiFi or other cellular networks).

A prevalent solution to address these challenges is license-assisted access (LAA), exemplified by LTE-U [11] and LTE-LAA/eLAA [12, 13, 10]. These approaches, along with research efforts studying LTE-WiFi coexistence [14, 15, 16, 17, 18, 19, 20], aggregate unlicensed carriers with existing licensed carriers. By utilizing licensed carriers as anchors, the impact of asynchronous interference is contained within the unlicensed carriers, minimizing overall disruption.

Another approach involves spectrum access servers, particularly in the Citizens Broadband Radio Service (CBRS) [21], which facilitate coarse timescale coordination for exclusive spectrum use. However, operating 5G entirely in unlicensed spectrum, as seen in initiatives like MulteFire [22] (in CBRS bands in GAA mode), exposes the vulnerability of capacity-enhancing multi-user schemes to asynchronous interference. For instance, WiFi access points may interfere with clients without being sensed by the base station (BS), highlighting challenges in readiness for practical deployments [23].

While [23] has showed promise for how 4G LTE systems can cope with such asynchronous interference, it left open several fundamental algorithmic and system level challenges that must be solved for a practical 5G solution, especially scaling to larger multiplexing gains (from larger bandwidths and antenna arrays in 5G) during real-time opera-

tions. Addressing these challenges is crucial for achieving the objectives of deploying 5G effectively in unlicensed spectrum, forming a key focus of ongoing research efforts.

2.3 Network Tomography

Spectrum sensing has emerged as a critical element in Dynamic Spectrum Access (DSA) research, as evidenced by numerous studies [24, 25, 26, 19, 14, 27, 13]. Spectrum sensing plays a crucial role in DSA by enabling unlicensed users to detect and utilize vacant spectrum bands efficiently. Spectrum sensing involves the detection and characterization of unused spectrum bands in a particular geographical area.

There are various techniques used for spectrum sensing including energy detection [28, 29], matched filtering [30, 31], and machine learning-based sensing [32, 33]. However, this process typically requires sophisticated algorithms and hardware capable of detecting signals within the spectrum of interest. While these methods have improved to be useful, it is important to note that estimating interference is not solely dependent on spectrum characteristics. We are interested in the interference characteristics not only at the location of the base station, where it is most likely that an advanced spectrum sensing tool will be present, but also at the locations of the clients.

Concurrently, network tomography has gained traction in wired networks [34, 35, 36, 37]. In this context, network tomography involves leveraging measurements between end-hosts to unveil significant network properties. A similar concept was introduced for DSA in [23], where cellular network clients were transformed into spectrum sensors. Their spectrum access measurements were then utilized to discern important interference patterns.

This concept capitalizes on the Orthogonal Frequency Division Multiple Access (OFDMA) technology, enabling base stations to orchestrate measurements by scheduling clients either individually or simultaneously with others, and subsequently estimating the resulting access distributions.

CHAPTER 3

THE PROBLEM OF INTERFERENCE

3.1 Impact of Interference

Unlike an asynchronous WiFi network, a synchronous cellular network enables BSs to centrally schedule both its downlink and uplink transmissions. This tight control enables the high multiplexing gains offered by 5G antenna technologies. However, in an open/unlicensed spectrum, the presence of hidden terminals (HTs) causes external interference, which affects clients, but cannot be sensed by the base station. This causes centralized scheduling to become a major drawback, and severely affects the network's multiplexing ability. We conduct network simulations with an LTE base station, several tens of clients, and a variable number of interfering sources, to highlight this impact as illustrated in Figure 3.1. The conventional proportional fair (PF) scheduler [38] is deployed alongside *mandated* unlicensed spectrum access protocol [10] at the BS. The results presented in Figure 3.2 reveal three key observations:

(a) *Interference can be modelled probabilistically:* Hidden terminal interference is commonly modelled as a Poisson Point Process [39, 40, 41]. As observed through our simulations, and verified by [42], this interference can be closely approximated by a Bernoulli Process with discretized time intervals.

(b) *Interference impact is significant:* As seen in Figure 3.2a, with an increase in the number of HTs, the ability of the BS to deliver or receive traffic from the affected clients diminishes, thus leading to a rapid drop (as high as 80%) in resource utilization. This reveals existing networks' acute vulnerability to interference despite sophisticated access protocols.

(c) *Limited interference information is useful, but not sufficient:* In contrast to PF, which




Figure 3.1: Illustration of hidden terminal interference

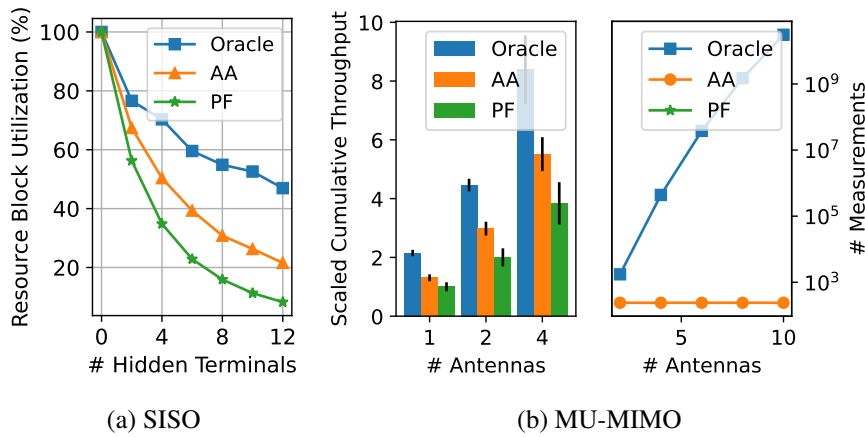


Figure 3.2: HOD prevents drop in performance with a) increased interference, b) larger number of antennas

is interference-agnostic, we consider an alternate *access-aware* (AA) scheme. It incorporates some knowledge of the interference—namely the probability that individual clients can access the channel—in the scheduling decisions. The AA scheme performs 50-100% better than PF, which is promising. However, as more antennas M are introduced in our 5G systems, the throughput depends on the specific set of users that are jointly scheduled on the same uplink resource. Hence, we also consider an Oracle that has access to all higher-order ‘joint’ access distributions (HODs) of all combinations of clients. Such an ideal scheduler benefits from interference-aware ‘group’ selection that is critical for MU-MIMO, and made possible by HODs. Figure 3.2 reveals that while AA improves performance, there is still a wide gap to the theoretical performance that is possible with HODs.

3.2 Challenge of Scalable Measurements

While HODs enables a network to operate in a smarter, more efficient way, the challenge lies in acquiring them in the first place. Since the interference is hidden from the BS, we rely on the clients to act as sensors and estimate their local interference. However, such an approach would reveal only the first-order *marginal* probabilities, i.e. “What is the probability that client i is interfered with?” We have seen (Figure 3.2) that we require the HOD to operate the network efficiently. The HOD answer questions such as “What is the probability that client i is interfered with but clients j and k are not?” We could continue to measure the HOD directly by scheduling each possible combination of clients and observing the results, but in an M -antenna MU-MIMO system with N total clients, this results in an exponential measurement overhead ($O(2^M)$) which affects the resources available for data transfer. Note that, if the system employs multiple channels (C), as is the common case today with carrier aggregation [43], then this overhead further scales as $O(C \cdot 2^M)$. As each of these parameters increases, so to does the number of measurements required to accurately estimate the desired probability.

Figure 3.2b demonstrates this challenge by plotting the number of measurements required to estimate the desired HOD for an increasing number of antennas. While the AA scheme does not rely on HODs, the number of measurements needed by the *oracle* to estimate and benefit from HOD increases exponentially, making it impractical.

To address this fundamental challenge, our proposed a network tomography approach for DSA, the system , allows us to leverage the benefits of HODs, not just for scalable performance in higher-order antenna and channel systems, but also in other interference-analysis applications (e.g. jammer localization)—all at a practical overhead that scales polynomially with the number of channels employed and the clients ($O(C \cdot \binom{K}{2})$).

CHAPTER 4

INTERFERENCE TOMOGRAPHY

As indicated in section 3.1, possessing comprehensive information about the joint-access distribution (HODs) proves advantageous for both resource management efficiency and pinpointing interfering sources. Accordingly, the objective is to precisely estimate HODs using minimal measurements and without disrupting the network’s regular operation. We do this by treating the clients as a distributed set of sensors that allow the BS to learn information about the interference environment and intelligently scaling those measurements to estimate the HOD.

The system consists of several steps summarized in Figure 4.1. The system’s operation is divided into two phases: a *direct measurement* phase followed by a *scaling* phase. The measurement phase consists of three steps: measuring first-order marginals per client, channel-wise clustering of clients, and measuring second-order marginals per cluster. First, the first-order marginals are measured during clients’ ongoing data transmissions by observing the interference they experience on each channel (subsection 4.1.1). Next, to reduce the amount of measurements required in estimating second-order marginals, the sys-

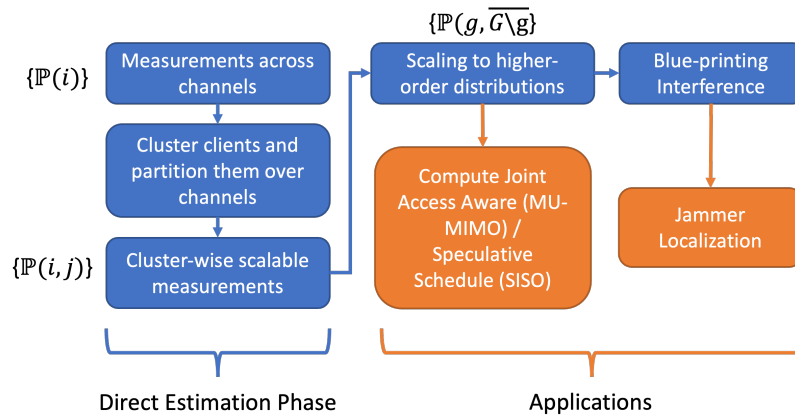


Figure 4.1: System design and information flow

tem uses a clustering algorithm to group clients with similar interference patterns such that interference dependencies can be captured at a cluster- (instead of client-) level without losing information (subsection 4.1.2). Finally, second-order marginals are measured by scheduling pairs of clients from unique clusters (subsection 4.1.3).

In the second phase (section 4.2), we use latent variable decomposition to scale these estimates up to the HOD. Once the HODs have been estimated, we can incorporate the probabilistic model into the schedulers. An additional step (subsection 4.2.4), which we call interference blue-printing, uses the HOD to reveal more information about the interference itself, allowing us to uniquely identify the sources of interference. This information, along with knowledge of a few client locations, can be used to localize the interferers.

Notation: Vectors, matrices and tensors are represented as \mathbf{a} , \mathbf{A} , $\underline{\mathbf{X}}$, respectively. The i -th element of a vector is given by $\mathbf{a}(i)$. Similarly, f -th column and i -th row of a matrix \mathbf{A} are notated $\mathbf{A}(:, f)$ and $\mathbf{A}(i, :)$, respectively.

4.1 Direct Estimation Phase

4.1.1 Channel Sampling

The first phase consists of estimating the channel access probability for each client individually across all of the channels. Clients are indexed with $i \in [1, \dots, N]$, and channels are indexed with $c \in [1, \dots, C]$. The probability that a client, i , is able to access channel, c , is described by a Bernoulli random variable with parameter a_i^c . We estimate these individual channel accesses probabilities empirically by counting the number of times a scheduled client is able to access the channel.

$$\mathbb{P}(\text{client } i \text{ can access channel } c) = \frac{\# \text{ times } i \text{ accesses } c}{\# \text{ times } i \text{ scheduled on } c} = a_i^c$$

By stacking each of the channel access probabilities for a given client, we create a channel access vector $\mathbf{a}_i = [a_i^1, a_i^2, \dots, a_i^C]$. This process requires scheduling each of the N clients

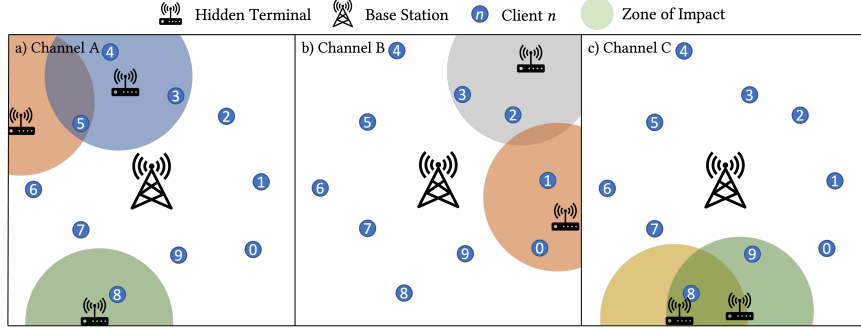


Figure 4.2: Impact of hidden terminals across three channels

across each of the C channels to take measurements, i.e. total overhead of $O(CN)$. Measurements are collected through uplink grants, and the clients send uplink packets as part of the normal operation of the network.

4.1.2 Client Clustering

The direct phase continues by measuring pair-wise access probabilities. Doing so naively would require us to schedule every pair of clients (jointly) on each of the channels, resulting in an overhead of $O(C \cdot \binom{N}{2})$ measurements. To reduce this number, we leverage the following observation: if two clients are physically located close enough such that they are affected by the same set of HTs, they will likely have similar first-degree marginal channel access probabilities over all of the channels. To illustrate this, consider Figure 4.2. In this topology there are ten clients, seven HTs, and three channels. For example, in channel A , client 5 is impacted independently by two HTs. However, in channel C , the two HTs are close enough such that their transmissions can be sensed, through some form of carrier sensing like CSMA, by one another. So, when the yellow HT is transmitting, the green HT cannot use the channel (which frees client 9 to transmit). In this scenario, client 8 is impacted by a *dependent* set of HTs. Now, consider clients 0 and 1. After the first phase of measurements, we would observe that they have similar channel access probabilities on all channels—no interference on channels A and C , and same probability of interference on channel B , i.e. $\mathbf{a}_0 \approx \mathbf{a}_1$.

We group the clients into clusters via K-means using the squared $L2$ norm of the difference between their channel access vectors, \mathbf{a} , to measure the distance. Note that clustering based on identifiers like RSSI has previously been used to spatially locate clients [44]. After clustering it is sufficient to perform pair-wise measurements with a single representative client from each cluster. Intuitively, this is because the clusters group clients that experience similar interference; therefore, the joint distributions of the representative clients across the clusters will be the same as those between the remaining, non-representative clients from those same clusters. The clustering phase reduces the size of measurement set from N clients to K clusters which does not depend directly on N . We note that K is a design parameter of the system. An approximation is sufficient, since K should simply capture number of the contention (interference) regions in the network. We investigate the impact of incorrect clustering as well as the choice of K on the accuracy of the eventual HOD and the measurement overhead reduction in subsection 6.2.

4.1.3 Pairwise Probability Estimation

The final step of the direct estimation phase is to measure the pair-wise channel access probabilities which can be computed empirically as before with the single client access. The BS simultaneously schedules a representative client from distinct clusters i and j and counts the number of opportunities one, both, or neither are able to access the channel. This incurs $O(C \binom{K}{2})$ measurements, after which we have an estimate of all (cluster) pair-wise channel access probabilities, $\mathbb{P}(i, j)$.

4.2 Scaling to High Order Distributions

Clustering helps to keep the measurement set manageable for estimating the pair-wise densities, but the number of measurements required to estimate HODs directly—even between clusters—quickly becomes unmanageable. In order to scale the pair-wise marginals up to higher order marginals of arbitrary size (including full joint distribution) we leverage latent variable decomposition techniques from [45, 46].

4.2.1 Joint Probability Representation

Consider a set of discrete, finite-alphabet random variables Z_1, Z_2, \dots, Z_N . We use $\mathbb{P}(z_1, z_2, \dots, z_N)$ as a shorthand notation to represent $\mathbb{P}(Z_1 = z_1, Z_2 = z_2, \dots, Z_N = z_N)$. Consider the scenario in which we do not have access to the joint probability distribution for $\{Z_n\}_{n=1}^N$. Rather, we only have access to one- and two-degree marginals, $\mathbb{P}(z_i, z_j)$ over all pairs $\{i, j\}$. This leads to an important question:

Can we recover the joint probability distribution given we have access to solely low-dimensional marginals without any structural assumptions on the random variables?

We can use the fact that any joint probability mass function (PMF) admits a Naive Bayes' model representation. This implies that any joint PMF can be generated from a latent variable model with just one hidden variable (H) given a sufficiently rich alphabet (F , [46] Theorem 1). Thus, the joint PMF of $\{Z_n\}_{n=1}^N$ can always be decomposed as,

$$\mathbb{P}(z_1, \dots, z_N) = \sum_{f=1}^F \mathbb{P}(H = f) \prod_{n=1}^N \mathbb{P}(Z_n = z_n | H = f). \quad (4.1)$$

We exploit this property to express degree-2 marginals as

$$\mathbb{P}_{ab}(z_a, z_b) = \sum_{f=1}^F \lambda(f) \mathbb{P}(Z_a = z_a | f) \mathbb{P}(Z_b = z_b | f). \quad (4.2)$$

The correctness of Equation 4.1 is related to canonical polyadic decomposition (CPD). Given an N -way tensor $\underline{\mathbf{X}} \in \mathbb{R}^{I_1 \times \dots \times I_N}$ with CP rank F , it admits a decomposition of the form:

$$\underline{\mathbf{X}} = \sum_{f=1}^F \lambda(f) \mathbf{A}_1(:, f) \otimes \mathbf{A}_2(:, f) \otimes \dots \otimes \mathbf{A}_N(:, f) \quad (4.3)$$

where \otimes is the tensor outer product operator, $\mathbf{A}_n \in \mathbb{R}^{I_n \times F}$ is called the mode- n latent factor, and $\|\lambda\|_0 = F$ is used to normalize the columns of \mathbf{A}_n . Mapping Equation 4.3 onto Equation 4.1, as explored in [46, 45], indicates that the random variables Z_1, \dots, Z_N are conditionally independent given a latent variable H .

4.2.2 From Pairwise to Higher Order Distribution

In the case of channel access probability, each of the random variables is binary ($z_n \in \{0, 1\}, \forall n$), so it suffices to write the degree-2 marginal for a pair of clients i and j as $\mathbb{P}_{ij}(\cdot, \cdot)$. Using Equation 4.2 we represent the joint access probability of clients i and j using the conditionals on the latent variable H . As shorthand notation, we use p_i^f for $\mathbb{P}(Z_i = 1 | H = f)$. There are $N \times F$ independent parameters for a given set of clients and latent variable H . We stack them to create the matrix \mathbf{P} . The probability that clients i and j will be able to access the channel jointly can now be written as

$$\mathbb{P}_{ij}(1, 1) = \sum_{f=1}^F \lambda(f) p_i^f p_j^f. \quad (4.4)$$

Similarly, the probability that client i can access the channel, while j cannot is

$$\mathbb{P}_{ij}(1, 0) = \sum_{f=1}^F \lambda(f) p_i^f (1 - p_j^f). \quad (4.5)$$

We observe that the probability measure \mathbb{P}_{ij} is parameterized by p_i^f, p_j^f , and the vector λ . Next, we minimize the KL-divergence between the measured pair-wise marginals $\hat{\mathbb{P}}(i, j)$ and the model $\mathbb{P}(i, j)$ for all pairs $\{i, j\}$. The loss function is

$$\mathcal{L} = \sum_{\text{all pairs } \{i, j\}} \mathbb{D}_{KL}(\hat{\mathbb{P}}_{ij}, \mathbb{P}_{ij}). \quad (4.6)$$

Under the naive Bayes' model assumption, the problem is a convex optimization problem with one hyper-parameter, the alphabet size F of the hidden variable H . Thus, we can use gradient descent to find λ^* and \mathbf{P}^* .

The power of this technique is that it allows us to write any marginal distribution of arbitrary degree using Equation 4.1. Consider we have a group of clients G and we want to know the probability that a subset of that group $g \subset G$ will be able to access the channel

while the remainder of G will not. This can be scalably found by:

$$\mathbb{P}(g, \overline{G \setminus g}) = \sum_{f=1}^F \lambda(f) \prod_{i \in g} p_i^f \prod_{j \in G \setminus g} (1 - p_j^f) \quad (4.7)$$

4.2.3 High Order Distribution Accuracy Analysis

This technique attempts to strike a balance between measurement overhead and the accuracy of the estimated HOD. It would simply not be possible to collect enough measurements to estimate the full distribution, nor would we be able to store the 2^N values for any moderate size value of N . On the other hand, if we start with a higher low-order marginal (e.g. three-way marginals instead of two) we could in principle increase the accuracy of the estimated HOD [45]. In practice, we find that using the pair-wise marginals as inputs is sufficient to achieve accurate HOD. Crucially, the modest gains provided by a higher order direct estimation do not outweigh the measurement overhead cost associated with them.

Figure 4.3 shows the mean squared error between the true and estimated HOD, for $N = 20$ over 50 randomized seed values, with various latent variable sizes F . The estimated HODs are calculated using the measured pair-wise marginals. On increasing the alphabet size of the latent/hidden variable, we see that the higher-order marginals are estimated with lower error. However, the error does not reduce substantially when we increase the alphabet size from N to $2N$. Thus, in the results that follow in chapter 6, we use a latent variable of size $F = N$.

4.2.4 Blue-Printing Interference

The estimate of the HOD is sufficient for scheduling as we will see in section 5.1, but it also provides the basis for generating a more accurate joint description of clients, and consequently interferers, in the physical space. Instead of treating the clients as users that are subject to interference, we change our perspective and use them as anchors/sensors with a finite search/impact radius. If we can successfully identify and assign interferers to




Figure 4.3: Mean squared error between true and estimated HOD from pair-wise marginal inputs vs. latent variable size

groups of clients based on their impact, we can create an interference graph. Leveraging the interference graph along with location information of a limited set of clients, we can narrow down and locate HTs to a small, useful area. This has implications for jammer localization, as we will see in section 5.2. We break down the process of generating the map of HTs into two steps, and illustrate using Figure 4.4.

Step 1: HT-focused client clustering: Define $\mathbf{p}_i = \boldsymbol{\lambda}^* \circ \mathbf{P}^*(i, :)$ where \circ is the Hadamard product. The vector $\mathbf{p}_i \in [0, 1]^F$ is a representation of the joint access probability distribution of client i and the realizations of the latent variable. Joint access vector of two clients with high correlation over the realizations of H , implies that the clients are affected by the same sets of HTs. On each channel, we form clusters of clients based on these vectors \mathbf{p}_i using DBSCAN [47]. In the associated example, this results in the clusters $\{0\}$, $\{1, 2, 8\}$, $\{3,4\}$, $\{6,7\}$, $\{9\}$, each colored differently, on the selected channel (Figure 4.4b). We then identify a representative for each of the clusters at random. The number of clusters obtained is of the order of HTs active on the channel.

Step 2: Estimating and mapping HTs: Next, we characterize how groups of clients are jointly impacted. We first generate a dependency graph $G = (V, E)$ where each representative client $i \in V$ and $E = \{(i, j) | \mathbb{P}(i, j) \neq \mathbb{P}(i)\mathbb{P}(j)\}$. That is, an edge exists between two representative clients if the product of their first-order marginals is not equal to the

pair-wise marginal, indicating that the channel access probabilities are not independent. A dependency graph is created based on these dependence relationships. In this example, the edges of the graph are $\{3,5\}$, $\{5,7\}$ and $\{0,9\}$ (Figure 4.4c).

We then take the following approach to approximate the number of HTs active on a channel. We note that

$$\mathbb{P}(\bar{g}|V\setminus g) = f(\text{HT impacts at least one of } g \text{ and not } V\setminus g) \quad (4.8)$$

This equation encodes the relationship between a HT and the likelihood it uses the same channel as the set of clients in g . The function $f(\cdot)$ evaluates to 0 if and only if there is no such HT. This is the key to identifying the smallest descriptive set of HTs. Using the relationship established in Equation 4.8, we obtain all such sets of clients affected *jointly* with non-zero probability. We do this by iterating over cliques g generated by graph G . We first consider the smallest size cliques and increment the size in each iteration—individuals, pairs, triples, etc.—and add a HT affecting the clique when the following is satisfied:

$$\mathbb{P}(\bar{g}|V\setminus g) > \prod_i \mathbb{P}(\bar{g}_i|V\setminus g) \quad \forall \text{ tuples } \{g_i\} \quad (4.9)$$

where $\cup_i g_i = g$ and g_i 's are sets found to be affected by HTs satisfying Equation 4.8 in a previous iteration. Intuitively, this process identifies whether the clique g is affected by a *new* HT rather than an independent combination of the *existing* ones. This procedure is summarized in algorithm 1.

In the accompanying example, we start with a clique size of 1. The only probability greater than 0 is the probability that client 0 is blocked conditioned on all others transmitting. Hence, we introduce a HT impacting client 0 (Figure 4.4d HT furthest to the right). In the subsequent iteration (clique size 2), we consider the cliques $\{3,5\}$, $\{5,7\}$ and $\{0,9\}$. For each, the probability of being blocked while other clients can transmit is greater than zero. Hence, we introduce three additional HTs to explain the additional information (Figure 4.4d remaining HTs).

Algorithm 1: Estimating Hidden Terminals

Data: Graph $G = (V, E)$, λ^* , \mathbf{P}^*
Result: Set of tuples (g) affected by unique hidden terminals
 $Q = \text{formCliques}(V, E)$;
 $g = \text{emptyList}()$; /*Initializing g^* */
 $\text{tuples} = \text{formTuples}(g)$; /*set of all comb. of g^* */
for q **in** Q **do**
 for tuple **in** $\text{sortBySize}(\text{tuples})$ **do**
 if $q == \text{Union}(\text{tuple})$ **then**
 if $\mathbb{P}(\bar{q} | V \setminus q) > \prod_i \mathbb{P}(\text{tuple}_i | V \setminus q)$ **then**
 $g.append(q)$;
 $\text{tuples} = \text{formTuples}(g)$;
 end
 end
 end
end

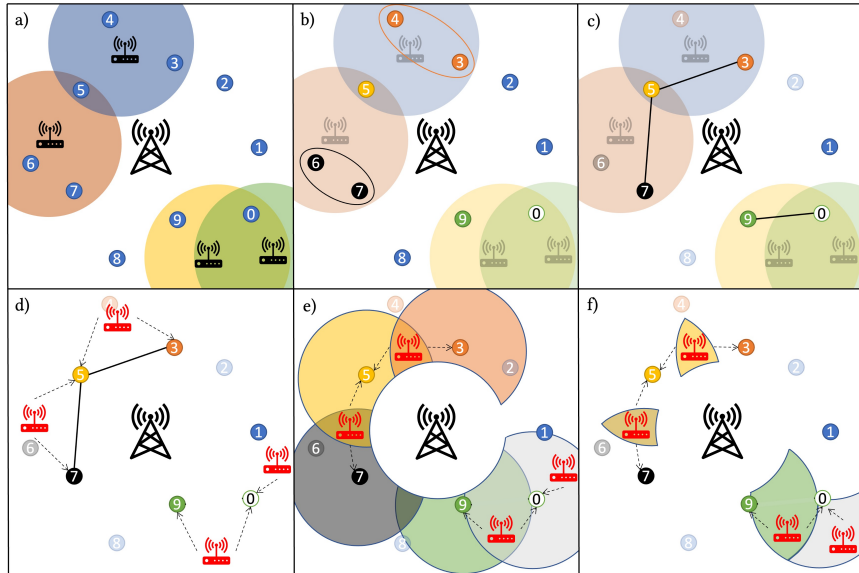


Figure 4.4: Blue-printing interference and localizing jammers: a) True interference map, b) HOD clustering, c) Graph creation, d) ID HTs, e) Estimated zones, f) Refine zones.

The result is to produce a topology like the one in Figure 4.4d which maps four HTs to their respective impacted clients. In section 5.2, we show how the estimated interference graph can be leveraged along with a limited amount of knowledge of the clients location to assist in localizing the HTs (shown in Figure 4.4e and f).

CHAPTER 5

APPLICATIONS

In this chapter, we describe two use cases in which knowledge of interference can significantly benefit the network: resource management and localization of adversarial jammers.

5.1 Resource Management

5.1.1 Current LTE/NR Schedulers

We briefly capture the essence of existing cellular schedulers where the BS is responsible for allocating time-frequency resource blocks (RBs) to clients on both DL and UL. Scheduling policies that balance throughput and fairness, are typically derived using a utility-optimization approach. We consider the widely-adopted proportional fair (PF) scheduling policy [38].

For SISO systems with B RBs and N single-antenna clients, the schedule is a matrix $\mathbf{X} \in \{0, 1\}^{N \times B}$ that assigns the RBs to users in each time slot. If the set of all possible schedules is S , the optimal schedule aims to maximize the aggregate marginal utility:

$$\mathbf{X}^*(t) = \arg \max_{\mathbf{X} \in S} \left\{ \sum_{b=1}^B \sum_{i=1}^N \frac{x_{i,b} r_{i,b}(t)}{R_i(t-1)} \right\}, \sum_{i=1}^N x_{i,b} = 1, \forall b \quad (5.1)$$

where $r_{i,b}(t)$ and $R_i(t-1)$ are the instantaneous rate (on resource block b) and average throughput of client i , respectively, and the constraint on RB assignment $x_{i,b}$, ensures that only one user is scheduled on each RB.

For MU-MIMO systems, where the BS has several antennas (M), a group of users is selected, $G \subset \{1, \dots, N\}$, $|G| < M$, and resource blocks are allocated to the group. Here,

the optimal schedule is:

$$\mathbf{Y}^*(t) = \arg \max_{\mathbf{Y} \in S} \left\{ \sum_{b=1}^B \sum_{i=1}^N \frac{y_{i,b} r_{i,b}(t, G)}{R_i(t-1)} \right\}, \quad (5.2)$$

where $G = \{i | y_{i,b} > 0\}$ and $\sum_{i=1}^N y_{i,b} \leq M, \forall b$. It is important to note that the MU-MIMO rate of a client $r_{i,b}(t, G)$ varies based on the other clients in its MU-MIMO group. At the end of every schedule, the average throughput of all clients is updated through a weighted moving average, to accommodate dynamic traffic conditions.

5.1.2 Interference-aware Scheduling

The existing schedulers, originally designed for licensed spectrum, do not incorporate interference information. In the following simulations, we progressively add interference information to the scheduling process. We consider the case of uplink, where the impact of HT interference is more pronounced. When each device (including both clients and HTs) employ LBT [10] before transmission in unlicensed spectrum, it is possible several of the BS's scheduled UL grants will not be utilized by the clients because of HT interference.

Recall that we can measure the probability of a client i utilizing its allocated grant, $\mathbb{P}(i)$. This indirect interference information can be incorporated to benefit the PF scheduler (Equation 5.1, Equation 5.2), so as to maximize the *expected* marginal utility. We term this optimal scheduler *Access-Aware* (AA). For MU-MIMO the scheduler is:

$$\mathbf{Y}_{AA}^*(t) \stackrel{MU-MIMO}{=} \arg \max_{\mathbf{Y} \in S} \left\{ \sum_{b=1}^B \sum_{i=1}^N \frac{\mathbb{P}(i) y_{i,b} r_{i,b}(t, G)}{R_i(t-1)} \right\} \quad (5.3)$$

where the same constraints on $y_{i,b}$ from Equation 5.2 apply. The equation for SISO can be recovered by setting $M = 1$.

5.1.3 Joint Access-Aware Scheduling using HOD for MU-MIMO

The aggregate rate of an MU-MIMO group not only depends on the set of users in the group, but also the ability of all users to transmit (subject to interference)—fewer partic-

ipating streams results in higher SINR for the active streams in the group. Therefore, a more accurate model of rate is $r_i(t, g, G)$, which is a function of both the group G , and the *successful* transmissions $g \subseteq G$. We can compute $\mathbb{P}(g, \overline{G \setminus g})$, i.e. the probability that all clients in g can transmit and the remaining clients, $G \setminus g$ cannot because they are blocked by interference using the parameters λ, \mathbf{P} found in subsection 4.2.1. The resulting scheduler is termed *Joint-Access-Aware* (JAA):

$$\mathbf{Y}_{JAA}^*(t) = \arg \max_{\mathbf{Y} \in \mathcal{S}} \left\{ \sum_{b=1}^B \sum_{i=1}^N \sum_{g \subseteq G} \frac{\mathbb{P}(g, \overline{G \setminus g}) y_{i,b} r_{i,b}(t, g, G)}{R_i(t-1)} \right\} \quad (5.4)$$

Choosing a user-group is non-trivial to solve computationally. We adopt a greedy approach to approximating the solution in a tractable time complexity (refer Equation 5.7). The difference between AA (Equation 5.3) and JAA (Equation 5.4) is that AA assumes that all client transmissions in the group G are successful, which is sub-optimal.

5.1.4 Speculative Scheduling using HOD for SISO

With the JAA reducing to AA in a SISO scenario, one might question the value for HOD in SISO systems. Note that while AA incorporates access information to cope better with interference, it is not able to efficiently address it and maximize usage of its resource blocks. Observe that in a single-user scenario, the *Joint-Access-Aware* scheduler reduces to *Access-Aware*. While the *Access-Aware* scheduler is more optimal than proportional fair, it is not able to maximize the usage of resource blocks. The challenges of interference management highlight the need to deal with scheduling in a more creative way. Through a policy termed *Speculative Scheduling* (SP) (introduced in [23]), we aim to leverage the additional information on these interferers to adaptively *over-schedule* clients on a RB. Overscheduling is analogous to a popular phenomenon of airlines overbooking flights in expectation that a certain number of people will not make it, leaving the exact number of seats occupied.

AA limits its vision to single-user scheduling, which naturally restricts the upper bound

on RB utilization, determined by interference. We observe that different clients in the same cell can be interfered by different sets of HTs, leading to *interference diversity*. This creates opportunities to schedule multiple clients on the same RB with the expectation that only one of them will be able to transmit, preventing collisions. We primarily consider SISO since the opportunities for over-scheduling diminish with higher multiplexing (MU-MIMO).

For instance, consider a client $i = \arg \max_k \left\{ \frac{\mathbb{P}^{(k)} r_{k,b}}{R_k} \right\}$ that was selected by AA. The speculative scheduler tries to select an additional client j such that scheduling it on the same resource block will likely increase the utility because they are impacted by different HTs. Specifically:

$$j = \arg \max_{j \neq i} \left\{ \mathbb{P}(i, \bar{j}) \frac{r_{i,b}}{R_i} + \mathbb{P}(\bar{i}, j) \frac{r_{j,b}}{R_j} \right\}. \quad (5.5)$$

Note that we want *exactly* one of them to be able to transmit to allow for SISO decoding (lest a collision). For a given set of over-scheduled clients on an RB, its expected utility depends on the joint access probability distribution of the set/group

$$E(G) = \sum_{i \in G} \mathbb{P}(i, \overline{G \setminus i}) \frac{r_{i,b}(t)}{R_i(t-1)}. \quad (5.6)$$

The optimal scheduler for a SISO system is $\mathbf{G}^*(t) = \arg \max_G E(G)$. It is difficult to efficiently compute such groupings, and we resort to a greedy strategy described by Equation 5.7. We start from a baseline G of two users and updating $G \leftarrow G \cup l^*$ after each iteration.

$$l^* = \arg \max_{l \notin G} (E(G \cup l) - E(G)) \quad (5.7)$$

5.2 Jammer Localization

Our framework for *blue-printing* the network interference from subsection 4.2.4, helped identify the HTs, the sets of clients they impact, and the magnitude of impact. An immediate use-case of such HT-client mapping is to identify the *location* of the HTs. This is of

special interest when the HTs are adversarial and the interference is an active, yet subtle (not always ON) attempt to impact the network.

We accomplish this by changing the anchoring of the interference from clients to physical space. Specifically, we request the coordinates of representatives $\{(x_i, y_i)\}_{i \in R}$, one each from the clusters identified through algorithm 1. Now, the region of impact around a particular client i is an area called $A_i = \{(x, y) | \text{distance}((x, y), (x_i, y_i)) < D\}$, where D is an estimate of the maximum range between a client and a HT, given the client's energy detection threshold (and maximum transmit power in the band). A HT located in the area A_i is assumed to impact the access of client i , located at (x_i, y_i) , while a HT outside it does not. Now, given a HT obtained using algorithm 1 and the group of clients it impacts (g), we can narrow down its location to the following area,

$$A_{HT} = \{\cap_{i \in g} A_i\} \cap \{\cap_{i \in R \setminus g} A_i^c\} \cap A_{BS}^c \quad (5.8)$$

where, A_i^c is the complementary area to A_i and A_{BS} is the area around the base station with range D . In other words, the HT is restricted to the common area determined by its affected clients and outside the area of both its unaffected clients as well as the BS (by definition of a HT). This process is captured in Figure 4.4e and f. We take the centroid of the resulting area to be the HT's location. Thus, with the locations of a few clients, we are able to construct an approximate physical map of the interferers. While this indirect localization is not exact, it is nonetheless impressive (median accuracy under 10m, section 6.3) and sufficient in our application.

CHAPTER 6

EVALUATION

6.1 Evaluation Methodology

We evaluate the system comprehensively in both resource management and jammer localization applications using a combination of NS3 [48] and numerical simulations. While we are able to instrument NS3 for DSA on a single channel SISO system using the LAA-LBT implementation [49], it is not conducive for multi-channel and multi-antenna scenarios, wherein we resort to numerical simulations.

We are able to instrument NS3 for DSA on a single channel SISO system using the LAA-LBT implementation [49]. In these simulations, we consider 100 topologies (similar to Figure 4.2). The HTs are distributed uniformly, but to ensure they are “hidden” from the BS, they are placed at least $70m$ away, where $50m$ is approximately the radius of the HT’s zone of impact. There are 2 to 8 HTs (Wi-Fi APs) randomly distributed amongst 8 to 20 clients. The BS is centrally located and the Wi-Fi nodes are out of its range, but impact the clients. NS3 currently supports LBT only at the BS (not at clients). Hence, we incorporate the impact of collisions using the physical layer traces. While this works for SISO, it is not conducive for multi-channel and MU-MIMO scenarios with many clients, wherein we supplement with numerical simulations. Our topologies for numerical evaluations mimic our NS3 network set-up. In all our simulations, clients are uniformly distributed in a circular area bounded by the BS’s transmission range.

To estimate the HOD, the BS follows the procedure in chapter 4. The first- and second-order marginals are estimated over 1000 frames. Unless specified, the number of clusters is N/C , for N clients and C channels. Each HT has a fixed probability of transmission drawn uniformly from the interval $[0.2, 0.8]$. If any node, either a client or another HT, is within

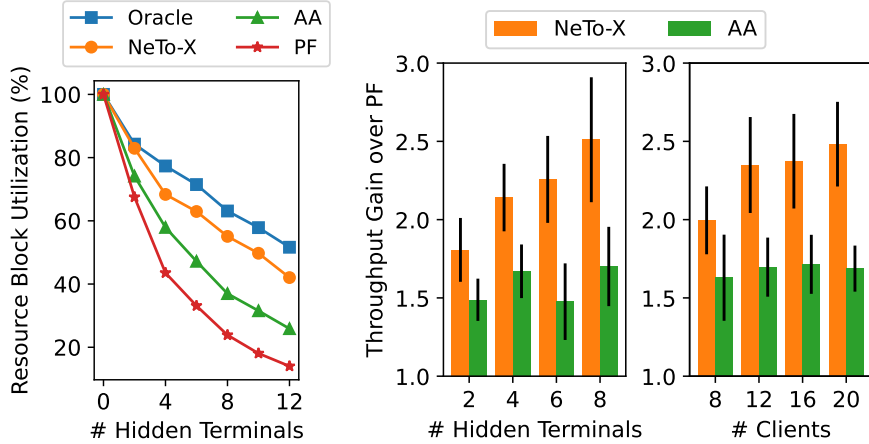


Figure 6.1: a) RB utilization on a single channel SISO system, b) Throughput gains over PF with varying HTs and c) clients.

the HT's zone of impact, the probability that they will not be able to transmit at a given frame is dependent on the HT transmit probability. Simulations consist of scheduling 1,500 LTE/5G frames. In each frame the BS attempts to schedule the N clients on $B = 10$ available RBs according to the schedulers presented in section 5.1.

6.2 Evaluations for Resource Management

First, we look at how the estimated HOD can be used to optimize client scheduling. We compare the system to PF and AA schedulers as described in section 5.1 as well as an optimal *oracle* scheduler. Our metrics include resource block utilization (RBU), which captures how efficient the scheduler is. We also report a measure of the cumulative throughput a schedule can achieve which is simply the sum of the long-term rates, and is normalized to that achieved by PF. Note that fairness is automatically incorporated by the schedulers in their policy. Additionally, we investigate the trade-off between performance and the measurement overhead.

Impact of Interference

We aim to understand the scheduler performance with increasing levels of interference. Figure 6.1a and b capture the effect of adding hidden terminals on the same channel for 20 clients using NS3. Figure 6.1a shows the percentage of resource blocks successfully utilized in a single-channel SISO system for 20 clients. Notice that the resource block utilization of all schedulers, including the *oracle* decrease as the level of interference increases, highlighting the inherent challenges of DSA. The PF scheduler performs worse since it is oblivious to the interference. Beyond eight HTs, its efficiency falls below 25%. The AA scheduler, which simply weights the PF schedule with the measured first-order marginal channel access probabilities, performs only slightly better. However, by estimating the HOD with the system, and performing speculative scheduling, we are almost able to achieve the performance of the *oracle*, differing by only a few percent. In Figure 6.1b, with an increase in the number of hidden terminals, the cumulative throughput falls drastically for PF (from (a)), the system minimizes losses through speculative scheduling to yield an appreciable gain of 50% over AA and 2x over PF.

Impact of Client Density

The LAA-LBT simulations in NS3 support a limited number of clients. In Figure 6.1c, we observe that the throughput gains over PF increase substantially with an increasing number of clients, wherein the diversity of interference impact is more. the system is able to leverage this diversity to speculatively schedule more groups of clients together. These gains are also observed for a much larger client density in our numerical evaluations.

Multi-Channel Performance

Next, we would like to investigate how the number of channels impact scheduling performance. Recall that modern LTE/5G networks use multiple channels through carrier aggregation to increase overall throughput, but this increases resource management complexity




Figure 6.2: Resource block utilization per channel and measurement overhead of SISO schedulers.




Figure 6.3: Normalized total throughput and overhead of MU-MIMO vs. number of antennas.

especially in the interference-rich environment of DSA.

The performance and measurement overhead of a SISO system with 30 clients vs. number of channels is shown in Figure 6.2. To isolate the effect of the channels from the level of interference, the number of hidden terminals per channel is fixed at 4. The performance of *PF* and *AA* is independent of the number of channels as they do not utilize cross-channel information. In contrast, the system clusters clients affected together, enhancing joint access probability estimation across channels for more efficient channel allocation. This allows the system to approach the performance of the *oracle* scheduler with fewer measurements (a factor of 10^{-6}) as the number of channels increases.

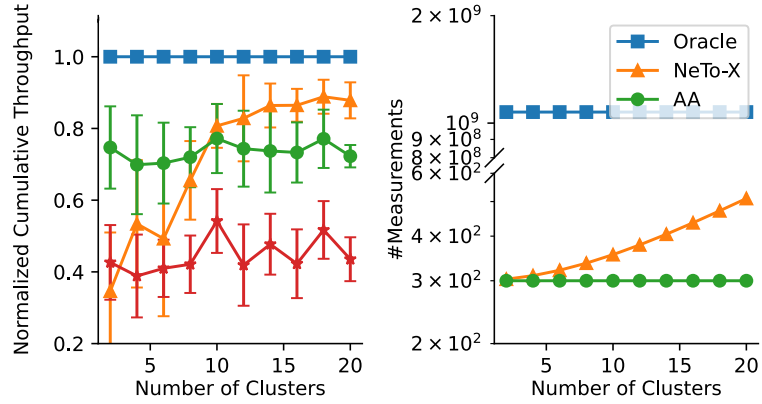


Figure 6.4: Effect of the number of clusters used by the system on the performance and measurement overhead.

Multi-Antenna Systems

We now investigate the benefits of the system for scheduling in MU-MIMO systems. the system outperforms both PF and AA schemes for client scheduling in MU-MIMO systems as seen in Figure 6.3. Indeed, the cumulative throughput of the system grows linearly with the number of antennas, while the PF scheduler starts exhibiting saturation with more antennas and is unable to deliver the promised gains. For a limited number of antennas we are able to compare the system to the *oracle*; however, as the number of antennas increases, the computational and storage overhead, let alone the required number of measurements, increases to a point where even simulation becomes infeasible. On the other hand, the system has a significantly smaller measurement overhead which *does not* scale with the number of antennas and clients making its implementation computationally feasible.

Clustering Analysis

One of the design parameters is the number of clusters, K , formed after obtaining the first-degree marginals as described in chapter 4. By grouping clients into clusters, we are making the assumption that they are impacted by the same HTs. A smaller number of clusters requires fewer measurements to estimate the pair-wise marginals, but it may mean

Table 6.1: Classification accuracy of the jammer localization system

Number of HTs	2	3	4	6	8
Accuracy (%)	99.3	90.5	81.7	73.5	70.6

that we have made incorrect assumptions about the nature of the interference. The trade off between scheduling performance and measurement overhead as a function of the number of clusters is captured in Figure 6.4. Aggressive clustering ($K < 10$ for 150 clients) results in the least amount of measurements but also causes the system to underperform compared to AA and potentially even PF. Increasing the number of clusters provides a better estimation of the joint access distribution at the cost of an increase in the number of measurements which scales as $\binom{K}{2}$, though this overhead is still significantly less than that of the *oracle*.

6.3 Evaluation for Jammer Localization

To evaluate how well the system can be used to locate jammers, we simulate the presence of 2-8 HTs with 40 clients over 200 randomized topologies. For each topology, we try to localize jammers using the method described in section 5.2. We perform the same analysis using the locations of (1) all UEs and (2) only the representatives described in subsection 4.2.4. The number of representatives ranges from 3 to 10 clients depending on the topology. After narrowing down the HTs to a candidate zone, we measure the metrics

- *Accuracy*: the distance between the ground truth location of the HT and the centroid of the candidate zone,
- *Precision*: the total area of the candidate zone.

Figure 6.5 depicts the accuracy and precision. When using all clients' location as anchors, the median accuracy is about 5 m. Using only the representatives' locations, the median accuracy reduces to 10 m. This granularity is still sufficient to infer the location of the jammers with high likelihood.




Figure 6.5: Accuracy (left) and Precision (right) of the system in localizing jammers with 40 clients and 3 hidden terminals

However, using the locations of only the representatives makes pinpointing the location of jammers precisely more challenging. The client placement is sparse and narrowing down the candidate zone is difficult. While the median precision is 200 m^2 , there is a long tail which extends beyond 400 m^2 in the 80^{th} percentile. When we have complete information of the locations, the system is able to determine an appreciably smaller candidate zone.

Table 6.1 shows the accuracy with which the system identifies the number of HTs present on a channel over 200 sample topologies with 40 clients. the system accurately identifies the number of jammers on the network ($> 90\%$ for up to 3 HTs). However, for a larger number of HTs, the accuracy dips to 70% as some HTs are very closely located to each other and tend to present themselves as a single, consolidated HT. However, this might still be satisfactory to localize the area of aggregate HT presence.

CHAPTER 7

CONCLUSIONS

As future wireless networks explore flexible spectrum sharing models, understanding the impact of external interference has become central to delivering the increased spectral efficiencies offered by their multi-antenna, multi-channel solutions. To this end, this work has made the following contributions.

This work proposed a scalable network tomography framework that enables efficient multi-user, multi-channel access for 5G/6G systems in unlicensed spectrum at low overhead. Further, the system performs interference-aware channel allocation and resource management mechanism that helps deliver the benefits of carrier aggregation and higher-order MU-MIMO even in the presence of interference. Lastly, this work proposed a novel mechanism that further leverages the hidden dependencies between interfering sources and clients, to localize the interfering sources from a carefully identified, small set of clients and their positions. All in all, the system significantly outperforms baseline approaches, and closely approximates optimal performance at significantly (several orders of magnitude) lower overhead.

While this work provides a concrete understanding of the interference environment with static clients, the robustness with client mobility remains unexplored. A future direction to explore would be the integration of machine learning algorithms to enhance interference prediction, localization, and mitigation capabilities in dynamic environments. Machine learning techniques could be used to analyze large volumes of data generated by the network and extract valuable insights to improve interference management strategies.

For further experimental validation and deployment, we would conduct real-world experiments and deployment trials to verify the effectiveness and practical feasibility of the proposed framework in diverse operating environments. This could involve deploying a

testbed consisting of a 5G base station and several other nodes in unlicensed spectrum to evaluate the performance of both, resource management and jammer localization techniques, under real-world conditions. Testbed deployments would serve as a stepping stone towards standardization efforts, i.e., to integrate the proposed interference management framework into existing and future wireless network standards.

REFERENCES

- [1] R. Ratasuk, D. Toli, and A. Ghosh, “Carrier aggregation in lte-advanced,” in *2010 IEEE 71st Vehicular Technology Conference*, 2010, pp. 1–5.
- [2] T. L. Marzetta, “Noncooperative cellular wireless with unlimited numbers of base station antennas,” *IEEE Transactions on Wireless Communications*, vol. 9, no. 11, pp. 3590–3600, 2010.
- [3] OnGo Alliance, *Citizen broadband radio service*, <https://ongoalliance.org/>, 2023.
- [4] G. L. Stuber, *Principles of Mobile Communication*, 4th ed. Switzerland : Springer, c2017.
- [5] H. Yin and S. Alamouti, “Ofdma: A broadband wireless access technology,” in *2006 IEEE Sarnoff Symposium*, 2006, pp. 1–4.
- [6] Q. Spencer, C. Peel, A. Swindlehurst, and M. Haardt, “An introduction to the multi-user mimo downlink,” *IEEE Communications Magazine*, vol. 42, no. 10, pp. 60–67, 2004.
- [7] E. Castañeda, A. Silva, A. Gameiro, and M. Kountouris, “An overview on resource allocation techniques for multi-user mimo systems,” *IEEE Communications Surveys Tutorials*, vol. 19, no. 1, pp. 239–284, 2017.
- [8] E. G. Larsson, O. Edfors, F. Tufvesson, and T. L. Marzetta, “Massive mimo for next generation wireless systems,” *IEEE Communications Magazine*, vol. 52, no. 2, pp. 186–195, 2014.
- [9] L. Lu, G. Y. Li, A. L. Swindlehurst, A. Ashikhmin, and R. Zhang, “An overview of massive mimo: Benefits and challenges,” *IEEE Journal of Selected Topics in Signal Processing*, vol. 8, no. 5, pp. 742–758, 2014.
- [10] 3GPP Release 14, 36.213, *Channel Access Procedures for LAA*, 2016.
- [11] LTE-U Forum White Paper, *Lte-u*, 2014.
- [12] 3GPP Technical Report 36.889, *Feasibility Study on Licensed-Assisted Access to Unlicensed Spectrum*, 2015.
- [13] Q. W. Paper, *License Assisted Access (LAA)*, <https://www.qualcomm.com/invention/technologies/lte>

- [14] M. Hirzallah, M. Krunz, B. Kecicioglu, and B. Hamzeh, “5g new radio unlicensed: Challenges and evaluation,” *IEEE Transactions on Cognitive Communications and Networking*, vol. 7, no. 3, 2021.
- [15] G. Naik, J.-M. Park, J. Ashdown, and W. Lehr, “Next generation wi-fi and 5g nr-u in the 6 ghz bands: Opportunities and challenges,” *IEEE Access*, vol. 8, pp. 153 027–153 056, 2020.
- [16] N. Rupasinghe and İ. Güvenç, “Reinforcement learning for licensed-assisted access of lte in the unlicensed spectrum,” in *2015 IEEE Wireless Communications and Networking Conference (WCNC)*, 2015, pp. 1279–1284.
- [17] A. Mukherjee *et al.*, “System architecture and coexistence evaluation of licensed-assisted access lte with ieee 802.11,” in *Communication Workshop (ICCW), 2015 IEEE International Conference on*, 2015.
- [18] S. Sagari, S. Baysting, D. Saha, I. Seskar, W. Trappe, and D. Raychaudhuri, “Coordinated dynamic spectrum management of lte-u and wi-fi networks,” in *2015 IEEE International Symposium on Dynamic Spectrum Access Networks (DySPAN)*, 2015, pp. 209–220.
- [19] Q. Chen, G. Yu, H. Shan, A. Maaref, G. Y. Li, and A. Huang, “An opportunistic unlicensed spectrum utilization method for lte and wifi coexistence system,” in *2015 IEEE Global Communications Conference (GLOBECOM)*, 2015, pp. 1–6.
- [20] S. Yun and L. Qiu, “Supporting wifi and lte co-existence,” in *2015 IEEE Conference on Computer Communications (INFOCOM)*, 2015, pp. 810–818.
- [21] S. Sarkar, M. Buddhikot, A. Baset, and S. K. Kasera, “Deepradar: A deep-learning-based environmental sensing capability sensor design for cbrs,” in *Proceedings of the 27th Annual International Conference on Mobile Computing and Networking*, ser. MobiCom ’21, New Orleans, Louisiana: Association for Computing Machinery, 2021, pp. 56–68, ISBN: 9781450383424.
- [22] Qualcomm White Paper, *MulteFire Specification*, www.multefire.org/specification, 2016.
- [23] R. K. Sheshadri, K. Sundaresan, E. Chai, A. Khojastepour, S. Rangarajan, and D. Koutsonikolas, “Blu: Blue-printing interference for robust lte access in unlicensed spectrum,” in *Proceedings of the 13th International Conference on Emerging Networking EXperiments and Technologies*, ser. CoNEXT ’17, Incheon, Republic of Korea: Association for Computing Machinery, 2017, pp. 15–27, ISBN: 9781450354226.

- [24] Y. Zeng, Y.-C. Liang, A. T. Hoang, and R. Zhang, “A review on spectrum sensing for cognitive radio: Challenges and solutions,” *EURASIP journal on advances in signal processing*, vol. 2010, pp. 1–15, 2010.
- [25] T. Yucek and H. Arslan, “A survey of spectrum sensing algorithms for cognitive radio applications,” *IEEE Communications Surveys Tutorials*, vol. 11, no. 1, pp. 116–130, 2009.
- [26] W. S. H. M. W. Ahmad *et al.*, “5g technology: Towards dynamic spectrum sharing using cognitive radio networks,” *IEEE Access*, vol. 8, pp. 14 460–14 488, 2020.
- [27] “LTE in unlicensed spectrum: Harmonious coexistence with Wi-Fi,” *Qualcomm Research*, Jun. 2014.
- [28] I. Sobron, P. S. Diniz, W. A. Martins, and M. Velez, “Energy detection technique for adaptive spectrum sensing,” *IEEE Transactions on Communications*, vol. 63, no. 3, pp. 617–627, 2015.
- [29] M. A. Abdulsattar and Z. A. Hussein, “Energy detection technique for spectrum sensing in cognitive radio: A survey,” *International Journal of Computer Networks & Communications*, vol. 4, no. 5, p. 223, 2012.
- [30] S. Kapoor, S. Rao, and G. Singh, “Opportunistic spectrum sensing by employing matched filter in cognitive radio network,” in *2011 International Conference on Communication Systems and Network Technologies*, IEEE, 2011, pp. 580–583.
- [31] A. Brito, P. Sebastião, and F. J. Velez, “Hybrid matched filter detection spectrum sensing,” *IEEE Access*, vol. 9, pp. 165 504–165 516, 2021.
- [32] D. Janu, K. Singh, and S. Kumar, “Machine learning for cooperative spectrum sensing and sharing: A survey,” *Transactions on Emerging Telecommunications Technologies*, vol. 33, no. 1, e4352, 2022.
- [33] K. M. Thilina, K. W. Choi, N. Saquib, and E. Hossain, “Machine learning techniques for cooperative spectrum sensing in cognitive radio networks,” *IEEE Journal on selected areas in communications*, vol. 31, no. 11, pp. 2209–2221, 2013.
- [34] N. Yaseen, J. Sonchack, and V. Liu, “Synchronized network snapshots,” ser. SIGCOMM ’18, Budapest, Hungary: Association for Computing Machinery, 2018, pp. 402–416, ISBN: 9781450355674.
- [35] R. Ben Basat, S. Ramanathan, Y. Li, G. Antichi, M. Yu, and M. Mitzenmacher, “Pint: Probabilistic in-band network telemetry,” in *Proceedings of the ACM Special Interest Group on Data Communication on the Applications, Technologies, Architectures,*

and Protocols for Computer Communication, ser. SIGCOMM '20, Virtual Event, USA, 2020, pp. 662–680, ISBN: 9781450379557.

- [36] I. Brugere, B. Gallagher, and T. Y. Berger-Wolf, “Network Structure Inference, A Survey: Motivations, Methods, and Applications,” *ACM Comput. Surv.*, vol. 51, no. 2, 24:1–24:39, Apr. 2018.
- [37] T. Bu, N. Duffield, F. L. Presti, and D. Towsley, “Network tomography on general topologies,” in *Proceedings of the 2002 ACM SIGMETRICS international conference on Measurement and modeling of computer systems*, ser. SIGMETRICS '02, New York, NY, USA: ACM, pp. 21–30, ISBN: 978-1-58113-531-2.
- [38] H. Kushner and P. Whiting, “Convergence of proportional-fair sharing algorithms under general conditions,” *IEEE Transactions on Wireless Communications*, vol. 3, no. 4, pp. 1250–1259, 2004.
- [39] R. W. Heath, M. Kountouris, and T. Bai, “Modeling heterogeneous network interference using poisson point processes,” *IEEE Transactions on Signal Processing*, vol. 61, no. 16, pp. 4114–4126, 2013.
- [40] K.-C. Huang and K.-C. Chen, “Interference analysis of nonpersistent csma with hidden terminals in multicell wireless data networks,” in *Proceedings of 6th International Symposium on Personal, Indoor and Mobile Radio Communications*, vol. 2, 1995, 907–911 vol.2.
- [41] L. Hu and Z. Dai, “Performance and reliability analysis of prioritized safety messages broadcasting in dsrc with hidden terminals,” *IEEE Access*, vol. 8, pp. 177 112–177 124, 2020.
- [42] A. Xia, “On the rate of poisson process approximation to a bernoulli process,” *Journal of Applied Probability*, vol. 34, no. 4, pp. 898–907, 1997.
- [43] 3GPP, “Evolved universal terrestrial radio acces (e-utra); medium access control (mac) protocol specification (release 10),” in *TS 36.321*.
- [44] B. Pinto, R. Barreto, E. Souto, and H. Oliveira, “Robust rssi-based indoor positioning system using k-means clustering and bayesian estimation,” *IEEE Sensors Journal*, vol. 21, no. 21, pp. 24 462–24 470, 2021.
- [45] N. Kargas, N. D. Sidiropoulos, and X. Fu, “Tensors, learning, and “kolmogorov extension” for finite-alphabet random vectors,” *IEEE Transactions on Signal Processing*, vol. 66, no. 18, pp. 4854–4868, 2018.

- [46] S. Ibrahim and X. Fu, “Recovering joint probability of discrete random variables from pairwise marginals,” *IEEE Transactions on Signal Processing*, vol. 69, pp. 4116–4131, 2021.
- [47] M. Ester, H.-P. Kriegel, J. Sander, and X. Xu, “A density-based algorithm for discovering clusters in large spatial databases with noise,” in *Proceedings of the Second International Conference on Knowledge Discovery and Data Mining*, ser. KDD’96, Portland, Oregon: AAAI Press, 1996, pp. 226–231.
- [48] *Network simulator (ns-3)*, <https://www.nsnam.org/>.
- [49] L. Giupponi, T. Henderson, B. Bojovic, and M. Miozzo, *Simulating lte and wi-fi coexistence in unlicensed spectrum with ns-3*, 2016.