The Spaces Between Us: Setting and Maintaining Boundaries in Wireless Spectrum Access

Lei Yang
Department of Computer Science, University of California, Santa Barbara
CA 93106, USA
leiyang@cs.ucsb.edu

Ben Y. Zhao
Department of Computer Science, University of California, Santa Barbara
CA 93106, USA
ravenben@cs.ucsb.edu

Haitao Zheng
Department of Computer Science, University of California, Santa Barbara
CA 93106, USA
htzheng@cs.ucsb.edu

ABSTRACT

Guardbands are designed to insulate transmissions on adjacent frequencies from mutual interference. As more devices in a given area are packed into orthogonal wireless channels, choosing the right guardband size to minimize cross-channel interference becomes critical to network performance. Using both WiFi and GNU radio experiments, we show that the traditional "one-size-fits-all" approach to guardband assignment is ineffective, and can produce throughput degradation up to 80%. We find that ideal guardband values vary across different network configurations, and across different links in the same network. We argue that guardband values should be set based on network conditions and adapt to changes over time.

We propose Ganache, an intelligent guardband configuration system that dynamically sets and adapts guardbands based on local topology and propagation conditions. Ganache includes three key mechanisms: an empirical model of guardband sizes based on power heterogeneity of adjacent links, network-wide frequency and guardband assignment, and local guardband adaptation triggered by real-time detection of cross-band interference. We deploy a Ganache prototype on a local 8-node GNU radio testbed. Detailed experiments on different topologies show that to minimize interference, traditional fixed-size configurations allocate more than 40% of available spectrum to guardbands, while Ganache does the same using only 10% of the spectrum, leading to a 150% gain in throughput.

Categories and Subject Descriptors

C.2 [**Computer-Communication Networks**]: Network Architecture and Design

General Terms

Design, Experimentation, Performance

Keywords

Dynamic Spectrum Access, Out-of-band Emissions, Guard-band Configuration

MobiCom'10, September 20–24, 2010, Chicago, Illinois, USA.

1. INTRODUCTION

Wireless is ubiquitous. At work or at home, we can find a high density of wireless devices competing for available spectrum. In a typical living room of 250 square feet (23 square meters), multiple wireless devices such as wireless displays, gaming consoles, media centers, and WiFi APs, are all within a few meters of each other. To transmit without mutual interference in high density environments, wireless devices can spread out onto orthogonal frequency ranges. Given the ever-increasing demands on wireless spectrum, it is fortunate that a new generation of frequency-agile wireless devices can dynamically access and share spectrum [1]. For example, existing WiFi devices can change their operating channels and the width of the channels to avoid interference and support varying traffic demands [5]. New devices based on software defined radios (SDR) can intelligently sense locally available spectrum and coordinate with other communication endpoints to occupy specific frequency ranges for reliable, high throughput communication [13, 30, 31, 33, 34].

As more devices in a given area are packed into orthogonal wireless channels, minimizing interference across channels becomes a critical problem. Artifacts such as hardware filter nonlinearity and radio propagation can cause transmissions on one frequency range to "spill" energy into adjacent ranges, creating undesirable cross-band interference. Wireless systems use "guardbands" to separate neighboring link frequencies, effectively acting as inert buffers that help protect each channel from energy spillover. For example, 802.11a makes nearby frequency channels (20MHz in size) orthogonal to each other by placing a 3.4MHz guardband between neighboring channels [17]. Since these guardbands are not usable for data transmission, this leads to a 17% overhead.

We ask the question: *"what is the best way to configure guardbands for today's high density networks?"* An ideal configuration not only buffers channels from their neighbors' transmissions, but does so using the smallest frequency range possible, leaving the rest for data transmissions. We perform experiments using commodity 802.11a devices, and find that the simple, fix-sized guardband configuration used by 802.11a fails to minimize cross-band interference. Our experiments emulate typical dense networks in residential or enterprise environments, and show that spillover interference can lead to as much as 80% throughput degradation. The impact is felt most strongly when a strong "interfering" transmitter sits on an adjacent but "orthogonal" channel. We also use USRP GNU Radios to confirm these results across different frequencies and power configurations.

A simple explanation of our findings is that 802.11a has chosen an overly aggressive guardband size. To verify this hypothesis, we configure our GNU Radio experiment to test different scenarios, each involving four nearby links occupying adjacent channels, and test the impact of different guardband sizes on performance. We make two observations from our results: First, we find that no single chosen value was appropriate for all links. In each case, some links were not buffered from adjacent transmissions and saw significant cross-band interference. Others were over-protected and significant frequency bands were wasted as guardbands. Second, different network topologies and power configurations also had significant impact on which guardband values worked best.

The conclusion from our experiments is that finding the best guardband size for a given network is very challenging, and finding a value that works for different network configurations is nearly impossible. The fixed-size guardband configuration, or "one-size fits all" approach does not work. To prevent cross-band interference without wasting spectrum to excessively large guardbands, we must configure guardbands based on local network topology and propagation conditions.

Given the wide variety of possible network configurations, we believe an effective and efficient (*i.e.* low overhead) guardband configuration involves both a static and dynamic component. For a given network configuration, static analysis can provide estimates of "good" guardband values. But as traffic load and propagation effects change, guardbands will need to adapt in time to remain effective. Thus we identify three key questions for guardband configuration. First, how can devices occupying adjacent frequencies decide the necessary amount of guardband to support interference-free transmissions? Clearly, a trial-and-error approach would lead to significant overhead and disrupt other transmissions. A more intelligent approach is necessary. Second, guardband sizes are a function of the network topology. Can we plan spectrum usage across the network to minimize the total overhead of guardbands? And given a planned network, how do we assign guardbands? Finally, devices should adapt guardband settings over time, based on observations of cross-band interference. But can devices distinguish cross-band interference from conventional channel loss as the real source of observed packet losses?

Ganache. Our solution to these questions is *Ganache*, an intelligent guardband configuration system. Ganache applies both centralized planning and dynamic per-link tuning to protect links against cross-band interference with minimum overhead.

First, a Ganache server builds and calibrates an empirical model to estimate required guardband sizes from measurements of power levels over frequency-adjacent links. Using this model, the server performs network-level frequency planning to allocate frequency usage to links and configure an effective set of guardbands that eliminate the bulk of cross-band interference. Because our results show that frequency-adjacent links with higher power heterogeneity require larger guardbands, the server can organize link frequency usage to minimize power heterogeneity, thus reducing network-wide guardband overhead.

After their frequency and guardbands are configured, individual Ganache links monitor physical distortion of their signals to detect residual cross-band interference, and adjust guardband locally to compensate. Together, these two techniques allow Ganache to configure guardbands for efficacy and minimal overhead.

We implement, deploy and evaluate a Ganache prototype on our 8-node USRP GNU radio testbed. Measurements show that Ganache can effectively suppress cross-band interference, thus improving link throughput by 30–150%. We also find that dynamic guardband configuration contributes to 50% of our performance gain, and frequency planning is particularly effective for reducing overhead near weak links. We also find Ganache's cross-band interference detection to have better than 91% accuracy. Overall, Ganache offers an effective and efficient way to tackle cross-band interference, thus restoring frequency orthogonality to support reliable concurrent wireless transmissions.

2. THE IMPACT OF GUARDBANDS

Our work begins with a detailed experimental study of guardbands and their impact on performance of commodity wireless systems. We are particularly interested in high density wireless settings, where nearby devices operate on adjacent channels to avoid interference. We perform two groups of experiments. One group uses commodity WiFi cards, while the other uses USRP GNU radios for more fine grain results. In both cases, we examine the impact of cross-band interference on link throughput as a function of guardband size.

2.1 WiFi Experiments

Our WiFi experiment seeks to measure the real impact of cross-band interference on today's commodity WiFi networks. 802.11a WiFi partitions its available spectrum into 20MHz channels, placing 3.4MHz guardbands* between adjacent frequencies to create "orthogonal" channels.

To study the efficacy of WiFi guardband configuration, we perform link-level experiments using four laptops with Linksys WiFi cards configured to 802.11a ad-hoc mode at 6Mbps. As shown in Figure 1a, we place the four laptops approximately 4 meters apart and form two links, each with a transmit power of 17dbm. We use iperf to generate UDP traffic and measure the link throughput of both links when they operate on the same or different channels. Each experiment runs for 1 minute, and we show the average result across 10 runs.

Because of the limited configurability of these devices, we can only gather coarse grain results for two different guardband sizes. We measure the throughput of each link (A and B) when they are operating on adjacent channels 1 and 2 (3.125MHz guardband), non-adjacent channels 1 and 3 (approximating a 23.125MHz guardband), and when both links are on channel 1 (as a point of reference). Figure 1b plots throughput normalized by the "ideal" throughput obtained in the absence of the second link.

In theory, channels 1 and 2 in 802.11a are non-overlapping and "orthogonal." Our results show that this does not hold in high-density environments. When the links are using channels 1 and 2, both links suffer significant loss in throughput, 20% for link A and 75% for link B. Link B suffers more loss because its signal strength at its receiver is weaker relative to that of link A, making it more sensitive to cross-band interference. The same near-far problem was also observed by [28]. This performance degradation is severe: total throughput is actually similar to the throughput in the overlapping scenario, where both links are on the same channel! In contrast, emulating a very large guardband by spreading links across channels 1 and 3 eliminates the bulk of the impact from cross-channel interference.

We then perform another set of experiments after disabling carrier sensing, from which we seek to understand whether the heavy degradation seen above comes from "amplification" by protocol features like carrier sense. That is, when sensing some energy spilled from "orthogonal" channels, a link could backoff "unnecessarily." Figure 1c plots the normalized link throughput using the 3.4MHz guardband size, by varying the physical distance between the two links. Compared to Figure 1b, the degradation (at 4m) re-

*Each 802.11a channel operates on 64 OFDM subcarriers, and uses the first 6 and last 5 subcarriers as guard bands. Thus the total guardband size is 11 subcarriers or $11 \times 20/64 = 3.4375$MHz.

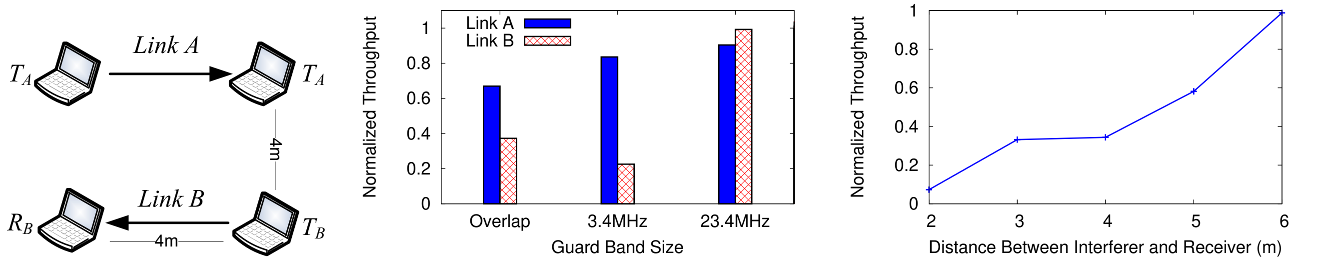

(a) Topology (b) Carrier Sensing On (c) Carrier Sensing Off

Figure 1: WiFi experiments: (a) 4 laptops with Linksys WiFi cards form 2 links. (b) Measured link throughput as a function of the guardband size. Current WiFi setting (3.4MHz guardband) is insufficient. (c) Measured link throughput after disabling carrier sensing but maintaining the 3.4MHz guardband. The impact of cross-band interference depends heavily on the network topology.

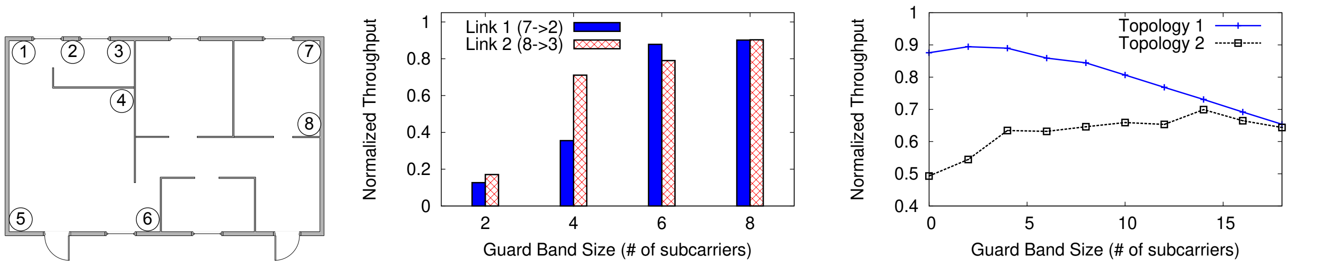

(a) Topology (b) Link Level Results (2 links) (c) Network Level Results (4 links)

Figure 2: GNU radio experiments: (a) The 8-node testbed in a 12m $\times$ 7m room with walls and furniture. (b) Measured link throughput as a function of the guardband size, with 2 links. (c) Measured network-level performance with 4 links and two different network configurations. The best guardband size differs significantly across network configurations. For topology 1, a guardband of 2 subcarriers produces the highest overall throughput. The same guardband configuration in topology 2 leads to 45% packet losses.

duces from 80% to 65%. However, the impact of cross-band interference is still dramatic, leading to 40–90% of performance degradation when the interferer is within 5 meters from the receiver. The impact depends heavily on network topology.

We also perform the same set of experiments using WiFi cards from two other vendors, Netgear and Wistron, and arrive at similar conclusions. The impact of cross-band interference is dramatic, with and without carrier sense. We note that our observation of strong cross-band interference between non-overlapping channels is different from those of an earlier study [24]. This is because [24] targets large-scale WiFi networks where APs are well separated. In this case, the impact of cross-band interference is significantly smaller than the dense environment targeted by our work.

2.2 USRP GNU Radio Experiments

To perform more fine-grain experiments on guardband size, we use a testbed of USRP GNU Radio devices. Unlike WiFi radios, GNU radios can be programmed to access various frequency ranges with fine grain control. We use an indoor testbed of eight GNU Radio nodes (see Figure 2a). Each node operates like a 802.11 device with carrier sense disabled: they use Orthogonal Frequency Division Multiplexing (OFDM) based modulation in the 2.4GHz range, with a total of 64 subcarriers, 52 of which are used for data transmission. Processing overheads on the radio platform limit each transmission to a smaller transmission bandwidth of 500KHz, and a subcarrier width of 7.8KHz (similar to those used in WiMAX, 10KHz). By changing each link's operating central frequency, we can effectively determine the number of subcarriers in each guardband.

We first seek to verify our WiFi results, by configuring four GNU radio nodes in a two-link topology (Figure 1a), and measuring the impact of different guardband sizes on each link's throughput. We show results in Figure 2b with different guardband sizes. The results are consistent with our WiFi results: both links suffer up to 80% throughput degradation when using a small guardband size (2 subcarriers). Compared to WiFi, the GNU Radios use smaller frequency guardbands because it uses a compact filter design.

Next, we want to understand the impact of different network configurations on the optimal choices for guardband sizes. We position eight GNU Radio nodes as illustrated in Figure 2a, and use them to form four links, each operating on a different frequency range. Between each frequency range, we place a guardband of size G subcarriers. As we vary G, we measure the total network throughput aggregated over all four links. Figure 2c plots this total throughput for two different network configurations, normalizing the result in each case against an ideal scenario with zero cross-band interference, where each link operates in isolation.

We see that optimizing G for each configuration can produce significantly different results for guardband size. In our case, topology 1 maximizes its network-wide throughput with a small guardband size of 2 subcarriers. Choosing a value of $G > 2$ degrades the throughput due to higher overhead. The same guardband size in topology 2 would produce 45% packet loss, compared to the maximum throughput of only 30% loss at $G = 14$. Our experiments with other random topologies produced similarly divergent results. Clearly, network configuration has a significant impact on the optimal guardband size, and no single value works for all network topologies.

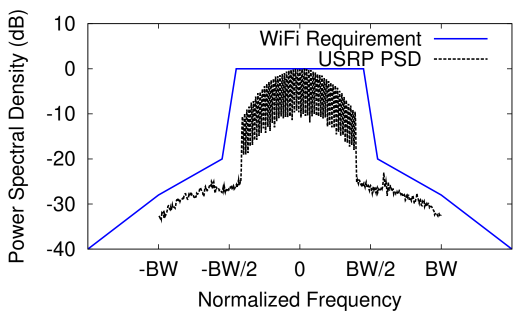

Figure 3: USRP's power spectral density is within the WiFi spectral mask requirement defined by IEEE standards [17].

Finally, we note that even after optimizing G for network-wide throughput, using a uniform guardband size across the network is suboptimal. A closer look at our traces reveals that for most values of G, there are always some network links that suffer significant cross-band interference, and other links for which the guardband is excessively large.

2.3 Discussions

Are Packet Losses due to the Capture Effect? When examining the impact of cross-band interference, we also distinguish it from packet losses caused by other factors. The most likely candidate is *the capture effect*, where a stronger interfering signal overwhelms the intended signal at a receiver, preventing the receiver from detecting/decoding any of its packets [20]. We found that this is not the case for our experiments. In our WiFi experiments, the interference power at a receiver is always weaker than the intended signal power. In our USRP experiments, more than 80% packet losses are directly caused by bit errors (and thus cross-band interference).

Is USRP a Representative Case? We choose USPR GNU Radios for our fine-grain experiments because of their wide availability and flexibility. While USRP radios do not have built-in hardware filters, we implement software digital filters to effectively reduce their out-of-band emission [34]. The resulting power spectral density, shown in Figure 3, indicates that the software filters effectively limit the out-of-band emission within the normalized WiFi spectrum mask requirement [17]. Therefore, we believe that the impact of cross-band interference on USRP radios is comparable to that of WiFi.

2.4 Summary

Overall, our experiments using both WiFi and USRP GNU radio devices lead to two key findings.

Cross-band Interference is Harmful. Our experiments show that in high density environments, cross-band interference can have a drastic impact on wireless throughput. Badly configured guardbands can fail to protect links from transmissions on adjacent frequencies, resulting in throughput degradation up to 80%.

Fixed-sized Guardband Placement is Ineffective. Measurements of 802.11a and GNU radio networks show that the "right" guardband size depends on a number of factors, including the radio hardware and the locations of wireless links. Not only is the optimal guardband size different for different network topologies, but a single guardband size can fail to protect some links while wasting valuable spectrum for others that require less protection.

3. THE SOLUTION SPACE

Initial conclusions from our experiments motivate us to identify new solutions for tackling cross-band interference. In this section, we discuss different potential solutions and their feasibility. More specifically, we consider two general types of solutions: using alternative mechanisms to mitigate cross-band interference while relying on simple fixed-size guardband configurations, and taking an adaptive approach to configuring guardbands. We argue that adaptive guardband configuration offers the most direct, efficient and effective solution. We then identify design challenges of this approach, and outline our proposed solution.

3.1 General Interference Mitigation

We explore potential solutions that would enable wireless links to reduce cross-band interference without dynamically configuring guardband sizes. Many existing mechanisms addressing channel loss fall into this category. We consider the most common solutions and discuss their suitability.

⋆ *Link Adaptation.* Links can use lower-order modulation or stronger coding schemes to allow successful decoding even in the presence of cross-band interference. This can be done at the packet level, or for each frequency subcarrier used. This approach can improve robustness against low levels of interference. The disadvantages of this approach are potentially significant reduction in power efficiency, and ineffectiveness in heavy interference scenarios such as Figure 1b-c.

⋆ *Carrier Sensing.* Alternatively, senders can delay their transmissions when detecting spillover energy on their frequency bands. This effectively performs time multiplexing between interfering links. Effectiveness is highly sensitive to the choice of sensing threshold, and is dependent on the power output of each transmitter. In addition, devices may detect interference on only a subset of their transmission subcarriers. Delaying transmissions would effectively waste a portion of their transmission spectrum.

⋆ *Power Control.* Links can reduce transmit power at boundary frequencies to reduce energy spillover. This is ineffective, however, since the lower power links themselves become vulnerable to neighboring transmissions. In addition, heterogeneous power levels are hard to avoid, since it is difficult in practice for devices from heterogeneous networks to synchronize power levels.

⋆ *Interference Cancellation.* A number of recent proposals describe receiver-side mechanisms that cancel interference at the packet-level [12, 15]. Adapting them to address only cross-band interference would require additional complexity and/or specialized hardware. Other solutions are designed specifically to reduce cross-band interference [4, 23]. These systems assume tight time or frequency synchronization, which is difficult and costly to implement in uncontrolled environments such as residential areas. Without tight synchronization, these mechanisms have been shown to be ineffective [16].

3.2 A Case for Adaptive Configuration

Prior work in this area has shown that correctly applying frequency guardbands can be more effective than alternative techniques at eliminating cross-band interference [16]. As a solution, increasing guardband sizes has several distinctive advantages, including simplicity, no additional hardware support, and independence with respect to network architectures and configurations.

Unfortunately, the results of our experiments show that that fixed-size guardband configurations are ineffective in real settings. In the absence of effective alternatives, we want to examine whether adaptively configuring guardbands based on local conditions can overcome limitation of the current "one-size fits all" approach. Awareness of local conditions means that links experiencing heavy cross-band interference can protect themselves using larger guardbands, while other links use smaller guardbands, leaving more spectrum available for data transmission.

To get an initial understanding of the potential benefits of this approach, we performed some simple GNU radio experiments, based on the topology in Figure 2. We used simple trial and error to tune the size of each link's guardbands, and found that this non-uniform approach produced per-link throughput improvements of up to 70%. Clearly, this direction is worthy of further exploration.

Solution Framework and Challenges. A reasonable system for determining local guardband sizes must satisfy two requirements. First, it must be efficient, *i.e.* it must be able to quickly produce an effective guardband configuration given information about individual link conditions in a specific network setting. Our trial-and-error straw-man solution is clearly too slow and disruptive to ongoing transmissions. Second, it must be adaptive, *i.e.* it must be able to tune itself as network conditions change over time.

The requirements naturally call for the design of a two-part system. This system includes a static, centralized component, which given data about a snapshot of link positions, transmission powers and transmission bands, produces an effective set of guardband configurations. The system also includes a per-link, dynamic component, which allows individual devices to make real-time corrections to guardband configurations by detecting the net impact of cross-band interference on transmissions. In the context of this general framework, we identify three key questions we must address.

⋆ *Static Guardband Configuration.* To build a system for configuring guardbands, we first need to understand this basic question: "what is the best way for devices on adjacent frequencies to determine the minimum guardband size necessary to minimize cross-band interference?" We can set estimates for guardband sizes once we have a strong model for understanding the relationship between local link metrics and guardband sizes.

⋆ *Frequency Planning.* Guardband sizes are a function of link frequency layout, *i.e.* who is adjacent to whom. Thus, guardband configuration can be integrated with network-level frequency allocation to reduce interference and guardband overhead. Thus we ask the question: "How can we use intelligent spectrum usage planning across the network to minimize the network-wide guardband overhead?"

⋆ *Online Adaptation.* Given network dynamics, devices must be able to adapt guardband settings based on real-time observations of cross-band interference. But how can devices distinguish cross-band interference from conventional channel loss as the source of observed packet losses?

3.3 Ganache

To overcome these challenges, we propose *Ganache*, an intelligent guardband configuration system that protects wireless links against cross-band interference while minimizing guardband overheads. Specifically, Ganache addresses the above three challenges using two complementary components, as shown in Figure 4.

Centralized Frequency Planning. In Ganache, a centralized server obtains measured power heterogeneity between any two links in its network, and uses a guardband model to estimate the required guardband if they were frequency-adjacent. The Ganache server then computes frequency usage and guardband configurations for all links in the network.

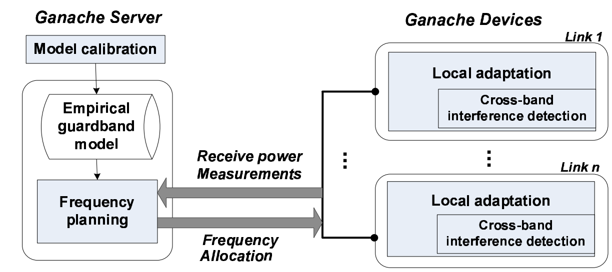

Figure 4: Ganache system architecture.

Local Guardband Adaptation. After receiving their configurations, links perform periodic observations of cross-band interference. Based on observed interference, each link uses a dynamic adaptation component to adjust guardband sizes locally. Cross-band interference is detected by measuring physical layer distortion in a link's received signals.

These two components are fully complementary. The centralized component controls guardband overhead and reduces cross-band interference to a minimum level, while the per-link dynamic component makes further local adjustments and adapts to time-varying dynamics. If and when local guardband adaptations begin to negatively impact data transmission, *i.e.* there are not enough subcarriers to maintain the desired data rate, a link can signal the central server to recompute the network-wide frequency and guardband configurations. We describe further details of these two components in Section 4 and Section 5.

4. GUARDBAND CONFIGURATION

Configuring guardbands across a network is a complex optimization problem. We take a two-step, model-driven approach in Ganache. We first use measurement experiments to capture the relationship between local network conditions and the guardband size required to block interference. To determine the size of guardband required at each link frequency boundary, we hypothesize that power levels of adjacent links is a critical factor in determining the level of cross-band interference, and consequently the guardband size necessary to block it. Using GNU radio measurements, we propose, verify and calibrate a model of guardband size based on power heterogeneity of frequency-adjacent links.

Using this model, a central Ganache server can take in link measurements from its network, and perform network planning, *i.e.* determine frequency and guardband usage for each link in the network. A centralized approach to guardband configuration offers significant advantages over local per-link decisions, because the latter may produce suboptimal results based on limited information gathered from nearby links.

4.1 An Empirical Model

Cross-band interference is the direct result of out-of-band emissions, where a transmission leaks energy immediately outside its own frequency range. The strength of cross-band interference created by an external interferer[†] i on a link s depends on i's signal strength seen at s's receiver, the out-of-band emission pattern, and the frequency separation between the two links, shown in Figure 5a. In the following, we seek to analytically model the relationship between the strength of cross-band interference and the size of guardbands used.

[†] We hereby refer to any transmitter who produces cross-band interference as an external interferer.

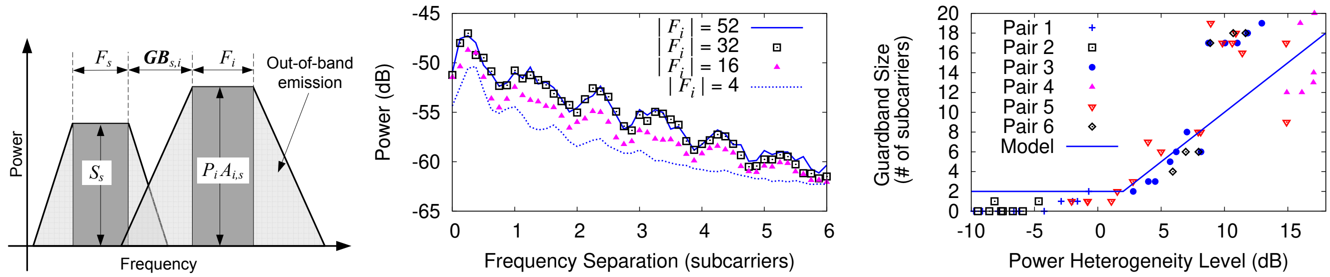

(a) An Abstract Representation (b) Measured Cross-band Interference (c) Mapping Required $\mathbf{GB}$ to Power Heterogeneity.

Figure 5: Understanding cross-band interference using GNU radio measurements. (a) An abstract representation in the frequency domain. (b) The strength of cross-band interference degrades exponentially with the frequency separation, and is relatively independent of the interferer's frequency usage $|F_i|$ if it is large enough. (c) The required guardband size depends heavily on the power heterogeneity level, from which we built a measurement-calibrated empirical model.

Our analysis assumes every node uses OFDM, a prevalent mechanism used by many standards to form transmissions. Every node uses the same subcarrier size and distributes transmit power uniformly across its subcarriers. Let P_i be a link i's per-subcarrier transmit power and F_i be its set of subcarriers in use. Let $\mathbf{GB}_{i\to s}$ be the required guardband size for link s to resist link i's cross-band interference.

Let $\mathbf{I}^{cross}_{i\to s}(f)$ represent the strength of cross-band interference produced by link i to the receiver of link s, at a s's subcarrier that is $f + \mathbf{GB}_{i\to s}$ away from i's frequency usage. We can estimate $\mathbf{I}^{cross}_{i\to s}(f)$ as:

$$\mathbf{I}^{cross}_{i\to s}(f) \approx \sum_{k\in F_i} P_i \cdot A_{i\to s}(f) \cdot \Omega(k, f, \mathbf{GB}_{i\to s}) \qquad (1)$$

where $A_{i\to s}(f)$ is the channel attenuation from i to s on f, and $\Omega(k, f, \mathbf{GB}_{i\to s})$ is the out-of-band emission pattern created by i's subcarrier k. Figure 5a illustrates an abstract representation of the cross-band interference, where the triangle marks the interference, and $\Omega(k, f, \mathbf{GB}_{i\to s})$ defines its shape. If we further assume that the channel impairment is frequency-flat, (1) reduces to $\mathbf{I}^{cross}_{i\to s}(f) \approx P_i \cdot A_{i\to s} \cdot \sum_{k\in F_i} \Omega(k, f, \mathbf{GB}_{i\to s})$.

Hou et al. examined out-of-band emission patterns in the context of decentralized OFDM transmissions [16]. They show that $\Omega(k, f, \mathbf{GB}_{i\to s})$ decreases exponentially with frequency separation, *i.e.* $f + \mathbf{GB}_{i\to s} + k$. Using GNU radio experiments, we measure the strength of cross-band interference generated by a link i. Results in Figure 5b confirm their analytical derivations in [16], and lead to two key findings:

- The cross-band interference produced by interferer i is most harmful to link s's subcarriers that are close to the frequency boundary.
- When the total number of subcarriers used by i's transmission ($|F_i|$) is above a certain threshold, it is no longer a factor in the amount of cross-band interference. Using experiments, we find that this threshold is 32 subcarriers for our current GNU radio settings.

From these findings, we found it reasonable to remove the f and k in the above equation, and instead use an abstract metric $I^{cross}_{i\to s}(F_s)$ to model the maximum cross-band interference produced by link i, seen by s's receiver on any of its subcarriers in use:

$$I^{cross}_{i\to s}(F_s) = P_i \cdot A_{i\to s} \cdot \hat{\Omega}(\mathbf{GB}_{i\to s}) \qquad (2)$$

where $\hat{\Omega}(\mathbf{GB}_{i\to s})$ is the out-of-band emission pattern. Because $\hat{\Omega}(\mathbf{GB}_{i\to s})$ relates to $\mathbf{GB}_{i\to s}$ in the log-scale [16], we propose the following linear model to relate them:

$$\begin{aligned} I^{cross}_{i\to s}(F_s)_{dB} &= (P_i \cdot A_{i\to s})_{dB} + (b - a \cdot \mathbf{GB}_{i\to s}) \\ &= I_{i\to s}(F_i)_{dB} + (b - a \cdot \mathbf{GB}_{i\to s}) \end{aligned} \qquad (3)$$

where $I_{i\to s}(F_i)_{dB}$ is the per-subcarrier power level observed by receiver s on link i's target frequency range F_i. The model parameters $a(a > 0)$ and b depend on the hardware configuration, including the precision of RF filters.

To minimize the impact of cross-band interference, the signal received on s's target frequency range must be stronger than its observed cross-band interference. Let $S_s(F_s)$ be link s's per-subcarrier power level received on its frequency range F_s. We have $S_s(F_s)_{dB} - I^{cross}_{i\to s}(F_s)_{dB} \geq \gamma$. With (3), this constraint implies that $\mathbf{GB}_{i\to s}$ must be large enough:

$$\begin{aligned} \mathbf{GB}_{i\to s} &\geq \frac{I_{i\to s}(F_i)_{dB} - S_s(F_s)_{dB}}{a} + \frac{b+\gamma}{a} \\ &= a' \cdot \mathbf{H}_{i\to s} + b' \end{aligned} \qquad (4)$$

This model maps link s's required guardband size into a linear function of $\mathbf{H}_{i\to s} = I_{i\to s}(F_i)_{dB} - S_s(F_s)_{dB}$, which we refer to as the level of power heterogeneity between s and i, as seen by s's receiver. Thus by observing the per-subcarrier signal strengths on both its frequency range (F_s) and on its adjacent interferer i's frequency range (F_i), we can estimate the guardband size required for link s to suppress the cross-band interference from link i.

4.2 Model Verification and Calibration

This empirical model indicates a strong *linear* relationship between $\mathbf{H}_{i\to s}$ and required $\mathbf{GB}_{i\to s}$. Using network measurements, we now verify this linear relationship and calibrate the model for our testbed by finding a' and b'.

To verify the model, we must measure for each frequency-adjacent link pair (s, i), the receive power levels $I_{i\to s}(F_i)_{dB}$ and $S_s(F_s)_{dB}$. To determine $\mathbf{GB}_{i\to s}$, we must sample different guardband sizes and find the minimum that suppresses the impact of interference. Obviously this requires fine grain control of guardbands, so we use the GNU radio testbed described in Section 2 and Figure 2a.

For each selected link pair (s, i) and their power settings, we perform the following experiment. We first turn on each sender separately and measure its link's packet loss rate (without any cross-band interference). We also measure $I_{i\to s}(F_i)_{dB}$, $S_s(F_s)_{dB}$ at s's receiver, and $I_{s\to i}(F_s)_{dB}$, $S_i(F_i)_{dB}$ at i's receiver, and compute

$\mathbf{H}_{i\to s}$ and $\mathbf{H}_{s\to i}$. We then turn on both links for 40 minutes and examine their packet loss rates under 20 different guardband sizes (1–20 subcarriers). We record $\mathbf{GB}_{i\to s}$ as the minimum guardband size required by link s to keep its packet loss rate below the original value plus 5%. Thus often $\mathbf{GB}_{i\to s} \neq \mathbf{GB}_{s\to i}$ because $\mathbf{H}_{i\to s} \neq \mathbf{H}_{s\to i}$, and we record them separately.

In total, we examined 100+ different link and power combinations for both links with line-of-sight (LOS) and without (NLOS). We also examined different frequency usages, and arrived at consistent findings. Figure 5c plots the required $\mathbf{GB}_{i\to s}$ as a function of the measured $\mathbf{H}_{i\to s}$ (and $\mathbf{GB}_{s\to i}$ vs. $\mathbf{H}_{s\to i}$). When $\mathbf{H}_{i\to s}$ is larger than 2, $\mathbf{GB}_{i\to s}$ is approximately a linear function of $\mathbf{H}_{i\to s}$, which confirms the trend predicted by our empirical model. When $\mathbf{H}_{i\to s}$ exceeds 10, however, we see a sudden rise in $\mathbf{GB}_{i\to s}$. This is due to an unexpected filter artifact, which generates extra side lobes outside the 64-subcarrier range.

An important observation is that the relationship between $\mathbf{H}_{i\to s}$ and $\mathbf{GB}_{i\to s}$ is independent of link formations and power settings. This motivates us to build a link-independent guardband model from measurement calibration:

$$\mathbf{GB}_{i\to s} = g(\mathbf{H}_{i\to s}) = \begin{cases} \mathbf{H}_{i\to s}, & \mathbf{H}_{i\to s} \geq 2 \\ 2, & \mathbf{H}_{i\to s} < 2 \end{cases}$$

To control cross-band interference, we intentionally make the model slightly conservative by using one more subcarrier than the best fit model.

There are noticeable differences between the model and the real measurement results, especially for large values of $\mathbf{H}_{i\to s}$. This can be attributed to artifacts of non-flat frequency fading, dynamic channel impairment, or the result of other simplified assumptions. Nevertheless, the empirical model provides a reasonable estimate of the required guardband size from local signal measurements.

4.3 Key Observations

The empirical model and measurement results produce two key observations on guardband configurations.

Local Information Is Not Enough. The required guardband size between links i and s depends on both $\mathbf{H}_{i\to s}$ and $\mathbf{H}_{s\to i}$. This is because $\mathbf{GB}_{i\to s}$ only captures the guardband size required to protect s from i's interference. To protect i and s against each other, the guardband between them needs to be of size $\mathbf{GB}_{s,i} = \max(\mathbf{GB}_{i\to s}, \mathbf{GB}_{s\to i})$. Thus the configuration depends on measurements from receivers of both links, and per-link local configuration may produce suboptimal results.

A Case for Reducing Power Heterogeneity. Both our model and measurements show that minimum guardband size scales with the level of power heterogeneity between adjacent frequencies. Therefore, careful assignment of frequencies to links can control the level of power heterogeneity between adjacent links, thereby reducing the aggregate size of guardbands network-wide. Figure 6 shows two assignments with three links s, $i1$ and $i2$. On the left, s is adjacent to $i1$ with $H_{i1\to s}$=2, and a guardband size of 2. On the right, s is adjacent to $i2$ with $H_{i2\to s}$=4, and a guardband of size 4. This shows we can reduce guardband overhead by organizing frequency usage such that, the perceived power of a receiver's own link is as close as possible to that of signals from adjacent frequencies.

Once we recognize that different frequency layouts across a network can impact overall guardband usage, it is clear that a centralized approach to network configuration offers significant advantages over local per-link configuration. Thus we propose for Ganache a network-wide, centralized approach to network planning, and describe it in detail below.

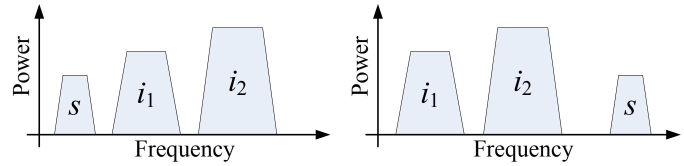

(a) Good Placement (b) Bad Placement

Figure 6: The placement of links' frequency usage is important. A bad placement (b) generates larger power heterogeneity across links, leading to additional guardband overhead than a good placement (a).

4.4 Centralized Frequency Planning

Given our calibrated guardband model $g(.)$, Ganache takes a centralized approach to network-wide frequency planning, *i.e.* assigning specific spectrum ranges to links to maximize efficiency and minimize guardband overhead. It includes two phases, *signal measurements* and *frequency planning*.

We first consider the case where link attenuation/fading is frequency flat, *i.e.* for each link pair (s,i), S_s, S_i, $I_{i\to s}$ and $I_{s\to i}$ do not depend on the transmission frequency f. When i and s are frequency-adjacent, they require a guardband of size $\mathbf{GB}_{s,i} = g(\max(\mathbf{H}_{i\to s}, \mathbf{H}_{s\to i}))$, independent of the actual frequency used.

Phase 1. Signal Measurements. Our central server requires each combination of transmitter and receiver to measure and report its local S_s and $I_{i\to s}$ values. The central server coordinates with all devices to transmit sequentially at their desired frequency width while all other devices listen. With good time synchronization [29], each device needs to only transmit a small number (< 10) of packets. Even with the relatively slow transmit rates on our current GNU radio hardware, each device can finish its measurement transmission in $< 50ms$.

Our signal measurement is quite reliable. We compute S_s and $I_{i\to s}$ directly from physical layer symbols across each packet [26], and average them over multiple packets, leading to a more stable estimate of $\mathbf{H}_{i\to s}$ than those from RSSI values. Our indoor GNU radio measurements show that most links' signals are stable on the order of hours (less than 4dB variance). In rare cases of heavy foot traffic near the devices, we see several seconds of significant variation. If necessary, the server can periodically trigger measurements and recalibrate frequency assignments to match variations over time.

Phase 2. Frequency Planning. The central server uses measurement results to compute a frequency allocation and guardband configuration that provides interference-free transmissions with minimal overhead. This optimal link frequency placement problem is NP-complete, because it can be reduced to the well-known NP-complete Traveling Salesman Problem (TSP) problem. Given the distance between any two cities, the TSP problem finds the shortest tour to visit a set of cities. For our frequency planning problem, we can map links to cities. Then we map the amount of guardband required between any two links to the distance between those the analogous cities in TSP. Thus, the problem of finding the minimum overall guardband reduces to the TSP problem of finding the minimum traveling distance to visit all cities.

Because these two problems are equivalent, we leverage existing TSP solutions to solve the central planning problem [11, 14]. For small networks like our 4-link topology, we can enumerate all possible combinations to identify the best placement. For larger networks, we can use greedy algorithms like nearest neighbor [14]

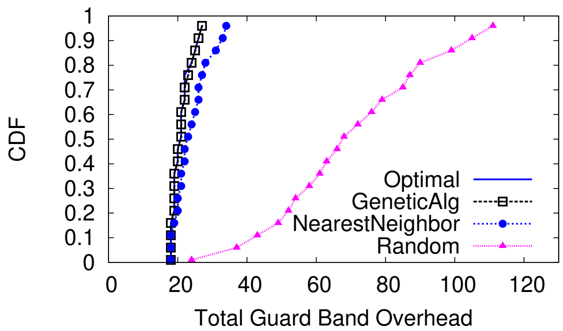

Figure 7: Comparing four frequency allocation algorithms using trace-driven simulations on 100 random topologies of 10 links.

($O(N^2)$ complexity), or genetic algorithms [11]. Both are lightweight, and have been shown to produce good approximations for TSP [11, 14].

We modified these algorithms for our use, and evaluate their efficacy using trace-driven simulations. We use measured attenuation results from our GNU radio testbed to calibrate a simulated 20x20m area with 20 devices using the same power level. We generate 100 random topologies, each with 10 links connecting randomly placed nodes. We run 4 planning algorithms, Optimal (brute force), GeneticAlgorithm [11], NearestNeighbor [14], and Random on each topology, and plot the total guardband required in Figure 7. GeneticAlgorithm mirrors Optimal, and NearestNeighbor is within 5%. More importantly, compared to Random, they reduce guardband overhead on most topologies by a factor of 3 to 5!

Addressing Frequency-Selective Fading. When links experience frequency selective fading, S_s or $I_{i \to s}$ are frequency-dependent. This creates new challenges for both the guardband model and network-wide frequency planning. For the guardband model, the selectivity could be partially compensated by using s's weakest subcarrier close to the boundary as S_s and i's strongest edge subcarrier as $I_{i \to s}$. Assuming the modulation and coding are tightly related to receive power, the model can still capture the impact of cross-band interference and the required guardband size. For frequency planning, the actual frequency location f matters. The server must adjust its network measurement and guardband computation to capture signal's variation across f.

5. ADAPTING GUARDBAND USAGE

While our model provides a good initial estimate of the guardband size that will suppress the bulk of cross-band interference, it may not eliminate the interference completely. In addition, a static guardband size will not work well over time as new links arrive and old ones disconnect. We now study the problem of dynamic guardband optimizations, and show how individual links can improve performance by iteratively adapting their guardband configurations based on real-time observations of cross-band interference.

Our proposed local adaptation consists of two phases: *detecting cross-band interference* and *local adjustments*.

Detecting Cross-band Interference. Reliable detection of cross-band interference is necessary for links to optimize their guardband configurations. It is also a difficult challenge, because the most obvious sign of cross-band interference is packet loss, which can be caused by multiple factors such as co-channel interference and fading. Thus the critical question is: how can links determine if observed packet losses are due to cross-band interference or conventional channel loss?

Our solution is to exploit information at the physical layer, relying on the fact that cross-band interference is particularly dominant on subcarriers near link frequency boundaries (Figure 5-b). In the presence of cross-band interference, data bits carried by these boundary subcarriers deliver lower channel quality than others. Since receivers themselves are unaware of errors at specific subcarriers, we use *symbol-level signal distortion* observed at different subcarriers as a reliable indicator of signal quality. The distortion measures the distance between a received symbol and its closest "constellation point" during the demodulation process. Figure 8-a shows a constellation map of BPSK modulation that plots two modulation points (1,0) and (0,1) as a cross, and all received symbols as dots. The demodulator decodes each received symbol as its closest modulation point: 1 if it is close to (1,0), and 0 otherwise. In general, signal distortion is inversely proportional to the confidence of symbol decoding: a symbol that experiences heavy impairment is likely to be further away from any modulation point.

Figure 8b and 8c illustrate the measured symbol distortion among subcarriers in use, with and without cross-band interference. We see that the distortion varies across subcarriers due to random channel impairments. But in the presence of cross-band interference, edge subcarriers suffer extremely high distortion that can be easily detected. Thus a receiver can reliably detect the presence of cross-band interference by examining symbol distortion levels across subcarriers.

This symbol distortion approach may fail to recognize cross-band interference if the interference disrupts the packet detection /synchronization process, which would produce errors across a large number of subcarriers. From empirical measurements on GNU radio platforms, we find these disruptions occur rarely in practice. Overall, our signal distortion based solution is highly effective, and is more than 91% accurate in our experiments.

Local Adjustments. Once a receiver detects the presence of cross-band interference, it coordinates with its transmitter to increase the link guardband size by one additional subcarrier. In most cases, a single additional subcarrier will dramatically reduce cross-band interference. Simultaneous decisions to increase guardbands by adjacent links can lead to unnecessary guardband overhead. However, this will only happen if both links detect cross-band interference simultaneously. In most near-far cases, the two links will experience different levels of link quality and interference, thus the weaker link will likely first detect the cross-band interference and increase its guardband. To prevent concurrent (and redundant) adjustments, each Ganache link waits for a small random delay before making adjustments.

It is tempting to also consider reducing guardbands in the absence of cross-band interference. However, such reductions can produce negative results. While it may free up more subcarriers for data transmission, it may create cross-hand interference for frequency-adjacent links. Since power heterogeneity is receiver-dependent, one link A might decrease its guardband while unaware of the interference it causes to an adjacent link B, forcing B to shrink its frequency usage to pad its side of the guardband. This effect is difficult to detect, and can propagate through the network. We believe this potential negative impact of reducing guardbands outweighs the benefit of freeing more subcarriers for transmission. In Ganache, links do not try to reduce guardbands. Instead, a Ganache server can periodically reconfigure network-wide guardbands based on updated network conditions, or upon detecting severe interference.

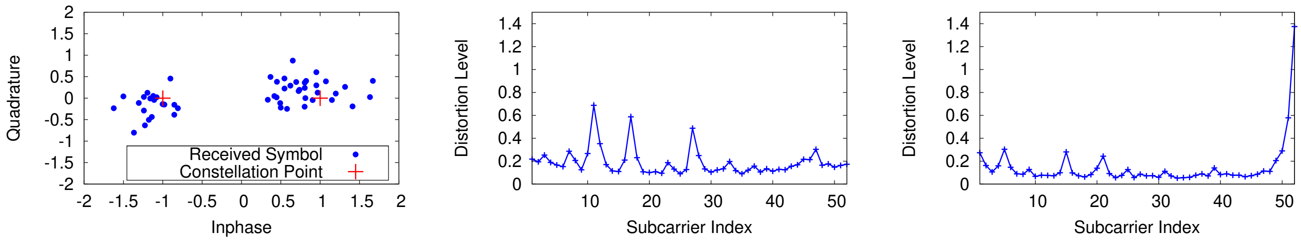

(a) Constellation Map (b) No cross-band interference (c) A cross-band interferer at the right side

Figure 8: Detecting cross-band interference (USRP GNU radio measurements): (a) The received BPSK symbols on the constellation map, decoded by the nearest constellation point (red +). The distortion is the distance between the symbol and its decoded constellation point. (b) The distortion across subcarriers when there is no interference. (c) The distortion across subcarriers when there is a cross-band interferer next to subcarrier 52.

6. A GANACHE PROTOTYPE

As a proof of concept, we implement a prototype of the Ganache system on top of USRP GNU radios, a widely available, reconfigurable software defined radio platform. While we chose GNU radios for their availability, our design can be ported to other platforms [13, 25, 27, 30, 33] for improved frequency bandwidth and processing speed.

To implement Ganache, we made modifications to the GNU radio software at both the physical and access layers.

Physical Layer. We configure GNU radios to operate on decentralized OFDM, each radio using a 500KHz band in the 2.38GHz range. The 500KHz is divided into 64 subcarriers with at most 52 subcarriers used for data transmission (406.25KHz). We adjust guardbands by changing each radio's central carrier frequency and its subcarrier usage. We modify the GNU radio software to expose the built-in signal distortion computation from the demodulation path. We add code to compute the per-subcarrier distortion averaged over a packet duration (63 symbols), and feed this back to the access layer for interference detection.

Access Layer. We implement the centralized planning component on a server connected to all GNU radios via Ethernet. We choose this option to enable all 8 radios for data transmissions. For link measurements, we compute the link attenuation S_s and $I_{i\to s}$ by averaging the signal strength over 300 consecutive symbols. We also implement local adaptation mechanisms on each radio, including a sender/receiver handshaking protocol to synchronize their frequency usage. Due to GNU radio's large processing delay [10], our current implementation does not perform local adaptation on a per-packet basis, but instead adapts every 500 packets.

To detect cross-band interference, a receiver extracts the per-subcarrier signal distortion for each corrupted packet, and measures the minimum distortion among 20 corrupted packets. If an edge subcarrier's distortion exceeds 3 times those averaged on subcarriers in the middle, we assert that cross-band interference is present. These parameters were chosen since they worked the best in our experiments.

Limitations. The current Ganache design focuses on exploiting the benefit of dynamic guardband configuration. It requires a central server for frequency planning, and targets static devices with similar OFDM configurations. One of the limitations of our current design is that it can tolerate only limited device mobility by using the local guardband adaptation to deal with occasional changes in link conditions. As future work, we plan to extend Ganache to support heterogeneous devices and explore decentralized designs.

7. EVALUATION

In this section, we evaluate our Ganache prototype using experiments on our eight-node GNU Radio testbed, running four wireless links concurrently. We evaluate the impact of network topologies by evaluating performance on both representative and random topologies. To emulate the different power levels used by different types of devices, *e.g.* wireless headphones or 802.11 access points, we assign different transmit power to different links. Each experiment runs for 2 minutes, and our results are each average values of 10 runs. We compare five systems for guardband configuration.

⋆ *Uni-Cons.* A conservative guardband scheme that uses a fixed value of 22-subcarriers for each guardband. it provides adequate protection for most links.

⋆ *Uni-Aggr.* An aggressive scheme using a fixed value of 2-subcarriers. It protects links with stronger signal strength than their external interferers.

⋆ *Model.* A basic version of Ganache with guardband values set using our guardband model, but no centralized link planning.

⋆ *C-Planning.* Ganache with centralized planning, but no local adaptation.

⋆ *Ganache.* The full version of Ganache.

Performance Metrics. To make a fair comparison between Ganache and its alternatives, we use a fixed spectrum range of size equal to 4×52 subcarriers (4×406.25=1.625 MHz). We allow 4 links to share this spectrum range while configuring their frequency usage and guardbands using different mechanisms. As our performance metric, we measure per-link throughput and compute their normalized ratio to the ideal throughput that each link can obtain if it operates in isolation using all 52 subcarriers. That is, if y is the performance in the presence of cross-band interference and x is the ideal throughput when there is no cross-band interference (using 52 subcarriers), we define **normalized impact** as $1 - y/x$. This metric captures the total impact of cross-band interference under each scheme, including throughput lost to guardband overhead and packets lost to insufficient guardbands. Thus lower values are better, and 0 is ideal.

7.1 Ganache vs. Fixed-size Configuration

We consider two representative network scenarios.

Topology 1 (Heterogeneous transmit power). We configure four links (node 1→5, 2→3, 4→7, 6→8) shown in Figure 9a. Link 2 (node 2→3) is short and transmits at power 10dB lower than others, but has a strong adjacent cross-band interferer (node 4). Figure 9b compares the normalized impact of Ganache and the fixed-

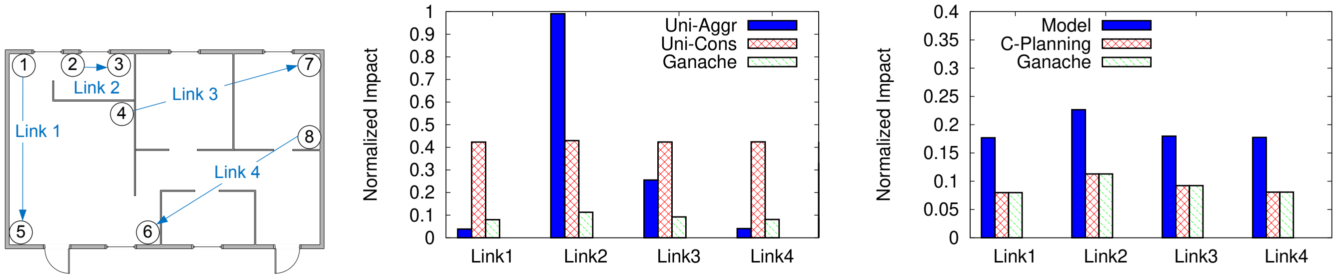

(a) Experiment Topology (b) Ganache vs. Fixed-size Approach (c) Ganache Performance Breakdown

Figure 9: Topology 1: four links (1→5, 2→3, 4→7, 6→8). Link 2 (2→3) is short and thus uses a 10dB lower transmit power than the rest, but suffers from a strong adjacent cross-band interferer (node 4). Ganache performs within a 10% distance from the idealized environment where there is no cross-band interference. (c) Both dynamic guardband configuration and centralized frequency planning contribute to the performance gain.

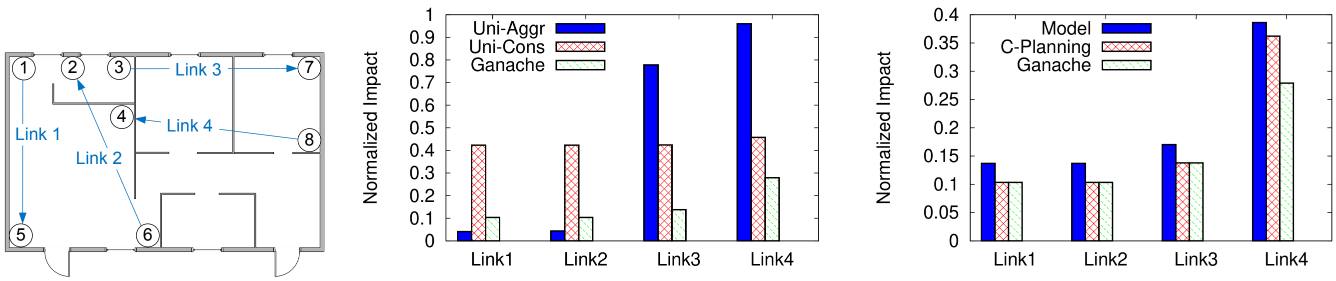

(a) Experiment Topology (b) Ganache vs. Fixed-size Approach (c) Ganache Performance Breakdown

Figure 10: Topology 2: four links (1→5, 6→2, 3→7, 8→4), all using the same transmit power. Both links 3 and 4 are weak while having a strong cross-band "interferer" nearby. Link 4 observes a large power heterogeneity(>10) where the model estimated guardband is less accurate. Ganache's local adaptation overcomes such error and improves link 4's throughput by adjusting the guardband usage.

size schemes. Being overly aggressive, Uni-Aggr fails to protect links 2 and 3 from cross-band interference, leaving link 2 a complete failure. Uni-Cons provides sufficient guardband protection but wastes frequency for unnecessary protection at link 1 and 4. In contrast, Ganache achieves throughput within 5-10% of the ideal, which is a 150+% throughput improvement over Uni-Cons, demonstrating its efficiency and effectiveness.

Topology 2 (Heterogeneous link attenuation). In this configuration, all four links use the same transmit power but links 3 and 4 are weaker compared to their external interferers due to signal attenuation from a room divider. Results from Figure 10b lead to similar conclusions: Uni-Aggr leads to severe packet losses for links 3 and 4 (70-90%); Uni-Cons consumes 40+% guardband overhead; and Ganache is within 13% distance to the ideal case for links 1, 2, 3 and 28% for link 4. Performance is worse for link 4 because it is a weak link and faces several strong interferers.

7.2 Impact of Individual Components

Ganache's performance gain can be attributed to its three components: model-based guardband estimation, centralized planning, and local adaptation. We now evaluate the contribution of each component.

Model-based Guardband Estimations. Figure 9c and Figure 10c compare the performance of different versions of Ganache. Comparing the basic "Model" scheme against Uni-Cons and Uni-Aggr (Figure 9b and 10b), we observe 50+% improvement in link performance for both topologies. This shows that using our guardband model to perform dynamic guardband configuration produces significant performance benefits.

Frequency Planning. To examine the benefits of centralized frequency planning, we compare the results of "C-Planning" and "Model" in Figure 9c and Figure 10c. We see that the inclusion of centralized planning boosts performance by another 50% in topology 1 and 20% in topology 2. To better understand the source of the benefits, we look at frequency boundaries of topologies with and without centralized planning, and plot in Figure 11 the maximum power heterogeneity observed. Centralized planning clearly reduces the power heterogeneity level for both topologies by moving vulnerable links away from their strongest interferer. Recall that negative power heterogeneity is desirable, *i.e.* the receiver's own signals dominate those of the interferer.

Local Adaptation. By comparing "C-Planning" with "Ganache" in Figures 9c and 10c, we see that local adaptation is only helpful in topology 2, where link 4 experiences large power heterogeneity (13dB) (see Figure 11, boundary #3). Our measurements in Section 4 indicate that the guardband model is less accurate when the power heterogeneity is beyond 10dB. In these cases, Ganache's local adaptation quickly expands the guardband size to suppress the remaining cross-band interference.

Cross-band Interference Detection. We now study the accuracy of Ganache's cross-band interference detection, using two groups

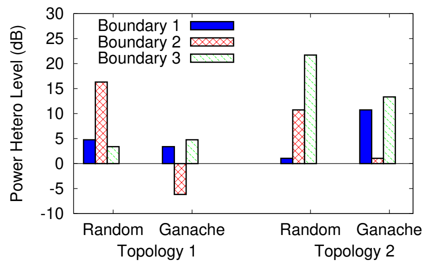

Figure 11: $H_{i \to s}$ using random and Ganache's central planning.

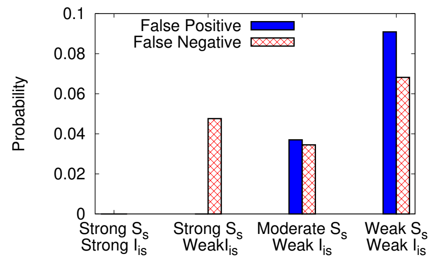

Figure 12: The accuracy of Ganache's cross-band interference detection.

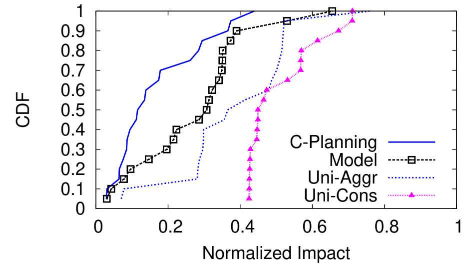

Figure 13: Comparing various schemes over 20 random topologies.

of controlled experiments to measure its false positive and false negative rates. To measure false positives (detecting normal impairments as cross-band interference), we turn on one single link, vary its transmit power to create packet losses, and examine detection results. To measure false negatives (failing to detect cross-band interference), we use a 2-subcarrier guardband to separate two frequency-adjacent links, making cross-band interference a major cause of packet losses. We run multiple experiments with different link combinations and transmit power settings. Figure 12 shows representative snapshots of four scenarios, where we refer to the links with SNR larger than 20dB as *Strong*, between 10dB and 20dB as *Moderate*, and below 10dB as *Weak*. We see that detection is highly reliable, *i.e.* the error rate is always less than 10%. The largest errors occur when both S_s and $I_{i \to s}$ are weak, and their distortion distributions display large randomness.

7.3 Overall Efficiency

While the above experiments use representative topologies, we also evaluate Ganache using 20 randomly generated 4-link topologies. To emulate heterogeneity in transmit power settings, we randomly select two links and set their transmit power to be 10dB lower than the rest.

Figure 13 shows the CDF of the performance of C-planning, Model, Uni-Aggr and Uni-Cons. We see that even the conservative Uni-Cons scheme with 22-subcarrier guardbands fails to protect links in half of the topologies. Uni-Aggr improves link throughput slightly by using more frequency, but creates large packet losses. Ganache's dynamic guardband configuration ("Model") effectively controls the impact of cross-band interference, reducing packet losses to a minimum. Adding centralized planning on top further reduces the guardband overhead, decreasing the total impact to 10% in average and always below that of Uni-Cons (40+%).

Local Adaptations. The above results excluded local adaptation in order to understand Ganache's static planning components. From the same experiments, we also examined the portion of all links that experienced >5% packet loss, and thus needed local adaptation. With a random frequency layout, 20% of all links required local adaptation. This number dropped to 11.25% with central planning. This shows that Ganache's central planning effectively reduces power heterogeneity to a range where the empirical model is more reliable. While the gain is marginal in this experiment, local adaptation can still help with link dynamics over time.

8. RELATED WORK

We classify the related work into the following four categories:

Spectrum Sharing Systems. Recent work has explored the use of dynamic spectrum sharing systems to improve spectrum utilization and network throughput [1, 3, 9, 22, 27, 34]. Most of these designs do not consider cross-band interference. Those that do use fixed-size guardband configurations [22, 34]. As our work has shown, fixed guardband configurations do not perform well in high density networks. Some recent projects set variable guardband sizes between devices. In [27], frequency-agile wireless devices gradually adjust their frequency separation from a legacy narrow-band device to prevent interference to the latter. The adjustment is triggered by poking the legacy device and observing its reactions. In [9], the guardband between TV white-space devices and DTV channels is determined by the devices' transmission power and spectrum mask. Ganache differs from these efforts by using central frequency planning to reduce power heterogeneity and by applying local guardband adjustments based on self cross-band interference detection.

Cross-band Interference in WiFi. Previous studies measured the impact of cross-band interference among WiFi devices [2, 6] or between WiFi and other devices [35], but did not provide any systematic solutions to suppress the interference. Recent work [28] verified the severity of cross-band interference among densely deployed WiFi devices (with carrier sensing), and proposed placing interfering links on well-separated channels. The authors apply a greedy algorithm to assign neighbor links to the farthest channels possible without estimating the guardband based on link conditions. Ganache, in contrast, can better manage spectrum and reduce guardband overhead by estimating guardband usage and planning frequency usage to reduce power heterogeneity.

Adaptive Guard Interval. In OFDM systems, Guard Interval (GI) is used between consecutive symbols to overcome inter-symbol interference. GI is typically set as a fix time that is higher than the maximum delay spread in a network. This fixed GI configuration could lead to unnecessary overhead, because different devices experience difference delay spreads. In [7, 8, 19], the authors propose to adapt GI according to each station's channel condition. While adding GI effectively reduces inter-symbol interference due to large delay spread in the time domain, Ganache addresses a different problem of cross-band interference due to out-of-band emission in the frequency domain.

Using Physical Layer Hints. Prior work has exploited physical layer information to assist protocol designs [18, 21, 32]. They use bit-level confidence measures to estimate the fraction of packets to retransmit or to configure transmission rates. Ganache's interference detection is inspired by these designs, but addresses an orthogonal problem of distinguishing packet loss caused by cross-band interference from loss caused by conventional channel impairments.

9. CONCLUSION

We study the impact of cross-band interference on high density networks. We find that it can have a drastic impact on wireless throughput, and that the current practice of using fixed-size guardbands is ineffective. We propose and prototype Ganache, a new guardband configuration system that sets and adapts guardbands to regain frequency orthogonality at a minimum overhead. Detailed experiments show that Ganache can produce throughput gains over current fixed-size solutions of up to 150%. To the best of our knowledge, Ganache is the first system to effectively control cross-band interference via dynamic guardband configuration.

Acknowledgments

We thank our shepherd Ranveer Chandra and the anonymous reviewers for their helpful suggestions. This work is supported in part by NSF Grants CNS-0916307, IIS-0847925, CNS-0832090, CNS-0546216, and CNS-0905667. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.

10. REFERENCES

[1] AKYILDIZ, I. F., LEE, W. Y., VURAN, M., AND MOHANTY, S. NeXt generation/dynamic spectrum access/cognitive radio wireless networks: A survey. *Computer Networks Journal (Elsevier)* (2006).

[2] ANGELAKIS, V., PAPADAKIS, S., SIRIS, V., AND TRAGANITIS, A. Adjacent channel interference in 802.11a: Modeling and testbed validation. In *Proc. of IEEE Radio and Wireless Symposium* (2008).

[3] BAHL, P., ET AL. White space networking with Wi-Fi like connectivity. In *Proc. of SIGCOMM* (2009).

[4] BRANDES, S., COSOVIC, I., AND SCHNELL, M. Reduction of out-of-band radiation in OFDM systems by insertion of cancellation carriers. In *IEEE Commu. Letter* (2006).

[5] CHANDRA, R., ET AL. A case for adapting channel width in wireless networks. In *Proc. of SIGCOMM* (2008).

[6] CHENG, C.-M., HSIAO, P.-H., KUNG, H., AND VLAH, D. Adjacent channel interference in dual-radio 802.11a nodes and its impact on multi-hop networking. In *Proc. of Globecom* (2006).

[7] DAS, S., FITZEK, F., CARVALHO, E., AND PRASAD, R. Variable guard interval orthogonal frequency division multiplexing in presence of carrier frequency offset. In *Proc. of Globecom* (2005).

[8] DAS, S., RAHMAN, M., FITZEK, F., AND PRASAD, R. Variable guard interval for OFDM based WLANs. In *Proc. of IEEE PIMRC* (2005).

[9] DEB, S., SRINIVASAN, V., AND MAHESHWARI, R. Dynamic spectrum access in DTV whitespaces: design rules, architecture and algorithms. In *Proc. of MobiCom* (2009).

[10] GE, F., ET AL. Software defined radio execution latency. In *Proc. of SDR* (2008).

[11] GOLDBERG, D. E. *Genetic Algorithms in Search, Optimization and Machine Learning*. Addison-Wesley, 1989.

[12] GOLLAKOTA, S., AND KATABI, D. Zigzag decoding: combating hidden terminals in wireless networks. In *Proc. of SIGCOMM* (2008).

[13] GUMMADI, R., NG, M. C., FLEMING, K., AND BALAKRISHNAN, H. Airblue: A system for cross-layer wireless protocol development and experimentation. In *MIT Tech. Report* (2008).

[14] GUTIN, G., YEO, A., AND ZVEROVICH, A. Traveling salesman should not be greedy: domination analysis of greedy-type heuristics for the TSP. *Discrete Appl. Math 117* (2002), 81–86.

[15] HALPERIN, D., ANDERSON, T., AND WETHERALL, D. Taking the sting out of carrier sense: interference cancellation for wireless LANs. In *Proc. of MobiCom* (2008).

[16] HOU, W., YANG, L., ZHANG, L., SHAN, X., AND ZHENG, H. Understanding the impact of cross-band interference. In *Proc. of ACM Coronet* (2009).

[17] IEEE 802.11-2007. `http://standards.ieee.org/getieee802/802.11.html`.

[18] JAMIESON, K., AND BALAKRISHNAN, H. PPR: partial packet recovery for wireless networks. In *Proc. of SIGCOMM* (2007).

[19] LAI, L., ET AL. An adaptive OFDM system with variable guard interval. In *Proc. of SPIE, Vol. 5284, 243* (2004).

[20] LEE, J., ET AL. An experimental study on the capture effect in 802.11a networks. In *Proc. of WinTECH* (2007).

[21] LIN, K. C.-J., KUSHMAN, N., AND KATABI, D. ZipTx: Harnessing partial packets in 802.11 networks. In *Proc. of MobiCom* (2008).

[22] MAHESHWARI, R., ET AL. Adaptive channelization for high data rate wireless networks. *SUNY Tech. Report* (2009).

[23] MAHMOUD, H. A., AND ARSLANU, H. Sidelobe suppression in OFDM-based spectrum sharing systems using adaptive symbol transition. In *IEEE Commu. Letter* (2008).

[24] MISHRA, A., ROZNER, E., BANERJEE, S., AND ARBAUGH, W. Exploiting partially overlapping channels in wireless networks: Turning a peril into an advantage. In *Proc. of IMC* (2005).

[25] NYCHIS, G., ET AL. Enabling MAC protocol implementations on software-defined radios. In *Proc. of NSDI* (2009).

[26] RAHUL, H., EDALAT, F., KATABI, D., AND SODINI, C. Frequency-aware rate adaptation and MAC protocols. In *Proc. of MobiCom* (2009).

[27] RAHUL, H., ET AL. Learning to share: narrowband-friendly wideband networks. In *Proc. of SIGCOMM* (2008).

[28] RAMAN, V., AND VAIDYA, N. H. Adjacent channel interference reduction in multichannel wireless networks using intelligent channel allocation. *UIUC T.R.* (2009).

[29] SHRIVASTAVA, V., ET AL. CENTAUR: realizing the full potential of centralized WLANs through a hybrid data path. In *Proc. of MobiCom* (2009).

[30] TAN, K., ET AL. SORA: High performance software radio using general purpose multicore processors. In *Proc. of NSDI* (2009).

[31] Universal Software Radio Peripheral, http://www.ettus.com.

[32] VUTUKURU, M., BALAKRISHNAN, H., AND JAMIESON, K. Cross-layer wireless bit rate adaptation. In *Proc. of SIGCOMM* (2009).

[33] Wireless open-access research platform, http://warp.rice.edu.

[34] YANG, L., HOU, W., CAO, L., ZHAO, B. Y., AND ZHENG, H. Supporting demanding wireless applications with frequency-agile radios. In *Proc. of NSDI* (2010).

[35] ZHU, J., WALTHO, A., YANG, X., AND GUO, X. Multi-radio coexistence: Challenges and opportunities. In *Proc. of ICCCN* (2007).