CPRecycle: Recycling Cyclic Prefix for Versatile Interference Mitigation in OFDM based Wireless Systems

Saravana Rathinakumar*, Bozidar Radunovic†, Mahesh K. Marina*
* The University of Edinburgh † Microsoft Research UK

ABSTRACT

OFDM is currently the most popular PHY-layer carrier modulation technique, used in the latest generations of cellular, Wi-Fi and TV standards. OFDM systems use cycle prefix to mitigate inter-symbol interference. However, most of the existing systems over-provision the size of the cycle prefix considering the worst case scenarios which rarely occur. We propose a novel OFDM PHY receiver design, called CPRecycle , that exploits the redundant cycle prefix to reduce the effects of interference from neighboring nodes. CPRecycle is based on the key observation that the starting position of the FFT window within the cyclic prefix at the OFDM receiver does not affect the received signal but can substantially reduce interference from concurrent transmissions. We further develop an algorithm that is able to find the optimal starting position of the FFT window for each subcarrier using a Gaussian kernel density function and a fixed sphere maximum likelihood detector. Through implementation and extensive evaluations using USRP and off-the-shelf IEEE 802.11g transmitters/interferers, we show the effectiveness of CPRecycle in significantly mitigating interference. CPRecycle only requires local modifications at the receiver and does not require changes in standards, making it incrementally deployable.

1. INTRODUCTION

Orthogonal Frequency Division Multiplexing (OFDM) is a spectrally efficient digital modulation method that is at the heart of almost all modern wireless systems. In OFDM, the stream of symbols (that represent the digitally modulated form of user data) are multiplexed over closely spaced sub-carriers and transmitted as parallel sub-streams. Orthogonality of sub-carriers makes them non-interfering with each other and in turn leads to other benefits including robustness to frequency-selective fading, flexible/dynamic channel-aware allocation of sub-carriers to users and ease of spectrum aggregation. For these reasons, Wi-Fi (WLANs based on IEEE 802.11) standards since 802.11a/g have adopted OFDM as the physical layer underlying a CSMA/CA multiple access scheme. 4G LTE mobile networks take this further by incorporating a multiple access scheme called OFDMA that allocates different users to different subsets of subcarriers[1]. The most recent digital audio/video broadcasting standards are also based on OFDM.

In order to maintain orthogonality between consecutive OFDM symbols, an OFDM transmitter adds a cyclic prefix in front of each symbol. This prefix is a copy of the end of each OFDM symbol whose purpose is to maintain orthogonality. The length of the prefix is adjusted to match the worst case delay spread that can occur in any deployment. This value is typically over-provisioned. The first OFDM-based Wi-Fi standard, 802.11a/g, specified 0.8μs long cyclic prefixes which corresponded to a signal path of 240m. Newer versions allowed the cyclic prefix to be halved, which is still hugely over-provisioned, give that the range of most of the Wi-Fi links is only few tens of meters. Similarly, standard LTE cyclic prefix lasts about 5μs and covers a signal path of 1.5 km.

In this paper, we present a novel receiver design called CPRecycle that leverages the over-provisioned cyclic prefix to mitigate the interference from concurrent wireless transmissions. The key observation underlying CPRecycle design is that when the receiver performs FFT with different starting points in the redundant portion of the cyclic prefix, the resulting signal component remains the same across the different FFTs but interference can vary by as much as 40dB, as we demonstrate in our measurements.

The main design challenge is how to find the optimal starting point for the FFT as it depends on the content of the interfering packet and it varies across subcarriers. This is

CoNEXT '16, December 12-15, 2016, Irvine, CA, USA
 ISBN 978-1-4503-4292-6/16/12…$15.00
DOI: http://dx.doi.org/10.1145/2999572.2999577

[1]More precisely, LTE uses OFDMA in the downlink direction. A variant called SC-FDMA is used for the uplink to suit lower cost and battery operated mobile transmitters with non-linear amplifiers.

very difficult as we cannot observe the interference signal in isolation. Instead, we create an empirical model of the interference as a function of the starting position of the FFT transformation. We then use this model to perform a maximum likelihood detection using the decoding outputs of all starting positions.

We implement CPRecycle on USRPs. An attractive aspect of CPRecycle is that it is local to the receiver and does not require any modification of the existing protocols nor changes at the transmitter, thus it can work with legacy devices. It is applicable to any OFDM/OFDMA based PHY with overprovisioned cyclic prefix. The computation complexity of CPRecycle can be tuned and it gracefully degrades to a standard OFDM receiver in the worst case.

In our evaluation we show that CPRecycle is useful in two important scenarios. The first scenario, *co-channel interference*, is a common case in real-world Wi-Fi deployments where multiple nodes access the same Wi-Fi channel at the same time. This can cause interference and packet losses, in particular in hidden-node scenarios. In the co-channel interference case we observe up to 15dB reduction in interference through the use of CPRecycle , even in case of the highest modulation rate (64QAM) and lowest coding rate (3/4).

The second important scenario is the adjacent channel interference scenario. All wireless transmitters experience RF leakage and cause interference even outside of their own channel. OFDM is able to maintain orthogonality between carriers only in perfectly synchronized systems, which rarely occurs [46]. In practice, there is a non negligible out-of-band interference and a guard-band is reserved to prevent interference between adjacent channels. We study the performance of CPRecycle interference mitigation in the adjacent channel interference scenario where we remove the guardband and tightly pack channels together. We observe that CPRecycle can remove up to 25dB of interference.

Through extensive simulation and experimental evaluations using USRP and commodity Wi-Fi hardware, we demonstrate the effectiveness of CPRecycle in significantly improving receiver side decoding in presence of interference, thereby also enabling efficient spectrum use. The network level benefits are significant due to the sharp drop in the average number of interfering neighbors in the network. In summary, the key contributions of this paper are:

- We propose CPRecycle , a novel receiver design that improves performance of existing OFDM-based wireless systems through an improved signal processing at the receiver, leveraging commonly overprovisioned OFDM cycle prefixes.
- As a part of CPRecycle , we propose a novel decoding algorithm that improves decoding performance by jointly decoding received signal over multiple FFT window positions.
- In our evaluation we show that we can reduce the effects of co-channel interference on a Wi-Fi receiver by up to 15dB and the effects of adjacent channel interference by up to 25dB by implementing only local modifications at the receiver.

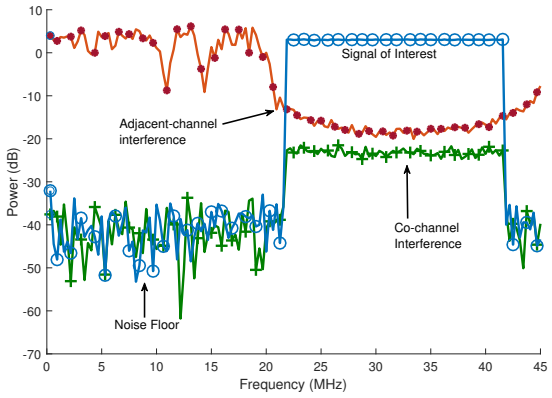

Figure 1: Illustration of Adjacent Channel Interference and Co-channel Interference.

2. BACKGROUND

This sections gives a brief overview of interference in OFDM based systems and the use of cyclic prefix for inter-symbol interference avoidance.

2.1 Interference in OFDM

Adjacent channel interference [9,10,34,49,58,59] occurs when an interferer while transmitting in its own channels, leaks part of its power into the adjacent channels, corrupting the signal received by a receiver in those adjacent channels. This can be due to imperfect filters at the transmitters or due to intermodulation of signals. Zubow et al. [59] analyze the effects of adjacent channel interference on 802.11 WLANs and observe that adjacent channel interference causes severe problems with the carrier sensing mechanism in 802.11. It was found that the carrier sensing mechanism can be too restrictive in some cases, leading the node to mistakingly defer its transmission, and too optimistic in some cases resulting in packet losses. An illustration of adjacent channel interference is shown in Fig.1. In this example, the sender is assigned a 20MHz channel (from 24 to 44MHz) in which it transmits the signal of interest. The interferer although assigned an adjacent non-overlapping 20MHz channel (1 to 21MHz in Fig.1) leaks energy into the adjacent band interfering with the signal of interest leading to a drop in SINR by about 15dB.

Another scenario where adjacent channel interference might occur is when two neighboring transmitters use partially overlapping channels, a very common scenario in IEEE 802.11 networks due the limited number of non-overlapping channels. In this scenario, there are three main problems caused due to adjacent channel interference. (i) Incorrect determination of a busy medium: when a transmitter performs carrier sensing before transmitting a packet, it may detect a high energy level due to an interferer leaking energy into its adjacent bands. This leads the transmitter to incorrectly assume that the medium is in use and defer its transmission. (ii) Signal corruption due to power heterogeneity: a weak signal

received by a receiver can be corrupted a high power interferer located close by leaking energy into the adjacent bands. (iii) Hidden terminals and exposed terminals that cause signal corruption due to adjacent channel interference cannot be handled through RTS/CTS, since the nodes are operating on a different channel, even though they are overlapping channels. One of the defining features of adjacent channel interference is the effect of interference power heterogeneity. The subcarriers closer to the channels occupied by the interferer are affected by a stronger interfering signal in relative to the other subcarriers, leading to a varying effect in different subcarriers.

Co-channel interference [13,28] occurs when multiple transmitters use the same subset of frequencies for communication. In IEEE 802.11 standards, co-channel interference is mitigated with the use of CSMA/CA, where transmitters would scan for an idle medium before transmissions. However, in dense IEEE 802.11 WLAN deployments, this situation cannot be avoided due to the limited number of non-overlapping channels in the 2.4GHz ISM band and overcrowding [12]. Gummadi et al [17], in their study of the effects of co-channel interference on 802.11 networks show how an interfering signal that is about 1000 times weaker can cause significant packet losses in a WLAN.

The presence of co-channel interference can have other indirect effects on the network performance as well. Using CSMA/CA, 802.11 nodes must scan the medium (for 4μs for 20MHz channel) and perform a clear channel assessment to determine if the channel is busy before transmission. The clear channel assessment can result in a busy medium when one of two following conditions are satisfied: (i) Carrier Sense – it is able to detect and decode an 802.11 preamble; (ii) Energy Detection – the energy detected in the channel is atleast 20dB greater than the minimum modulation and coding rate sensitivity. In the presence of co-channel interference, the transmitter would perform an exponential back-off which reduces the achievable throughput. Significant improvements in throughput [48] can be achieved by reducing this energy detection threshold.

In cellular networks, the use of femto cells can cause co-channel interference when deployed in a co-channel or hybrid configuration. In these configurations, a macrocell is overlaid with OFDM based femto cells assigned an overlapping set of channels. This can cause co-channel interference between neighboring femto cells sharing the same set of channels (co-tier interference) or between a femto cell and a macro cell (cross-tier interference) [39]. While co-tier interference can be managed through an efficient allocation of subcarriers, it is far more difficult to manage cross-tier interference due to limited availability of the wireless spectrum.

2.2 Cyclic Prefix

In OFDM based systems, the cyclic prefix (CP) or guard interval, illustrated in Fig. 2, is used primarily to prevent inter-symbol interference (ISI). ISI is a type of signal distortion that is caused when consecutively transmitted symbols interfere with each other at the receiver. This is due to the multi-path propagation characteristics of the wireless channel, where a transmitted signal may take multiple paths from the transmitter to the receiver with different propagation delays and the multiple copies of the signal may interfere with itself. The cyclic prefix acts as a guard period between successive OFDM symbols, thereby completely eliminating the ISI. The duration of the cyclic prefix is chosen to be greater than the largest delay spread expected by any user in the target environment. It is usually defined in the communication standards and cannot be changed adaptively based on the environment due to interoperability issues.

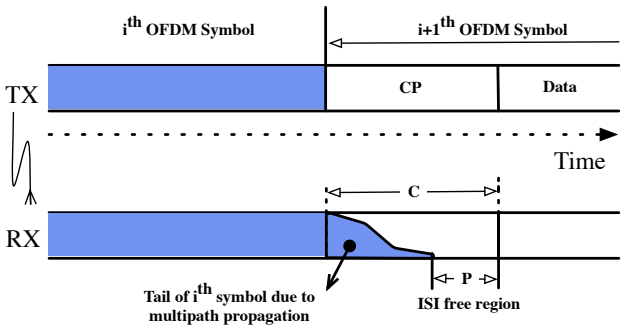

Figure 2: Illustration of Cyclic Prefix (CP) or guard interval.

In practice, cyclic prefix is a copy of a portion of the OFDM symbol towards its end, and it is inserted before the actual OFDM symbol. Because of the way cyclic prefix is constructed, only one symbol from the intended transmitter is received at any point in time during the whole course of duration spanning the cyclic prefix and actual OFDM symbol.

Standard	Bandwidth	FFT Size	CP Size	Duration
802.11a/g	20 MHz	64	16	0.8 μs
802.11n/ac	40 MHz	128	32 (16)	1.6 (0.8) μs
802.11n/ac	80 MHz	256	64 (32)	3.2 (1.6) μs
802.11n/ac	160 MHz	512	128 (64)	6.4 (3.2) μs

Table 1: Cyclic Prefix in 802.11 standards

The downside of using cyclic prefix is that it lowers the spectral efficiency since no additional information is transferred during the cyclic prefix period. Note that cyclic prefix duration is chosen based on the maximum delay spread which can result in substantial portion of the overall symbol period being consumed by the cyclic prefix. For example, in 802.11 systems about 20% of the symbol duration is allocated for the cyclic prefix.

Table 2.2 lists size and duration of cyclic prefix specified in different 802.11 standards with the default long guard interval as well as the short guard interval (in parentheses). In LTE, the normal cyclic prefix length is 4.7μs, an overhead about 7% in a OFDM symbol with actual data portion of about 66.7μs. There is also an extended cyclic prefix of length 16.7μs specified in LTE for broadcast services and environments with long delay spreads, increasing the cyclic prefix related overhead to 25% in this case.

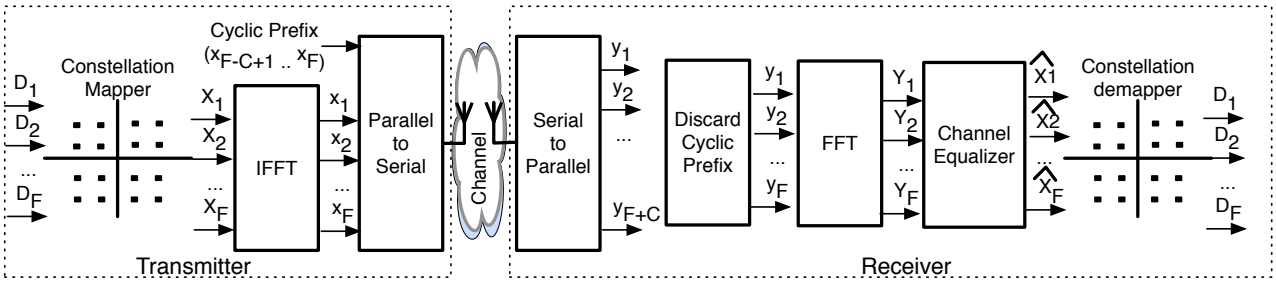

Figure 3: Schematic of a standard OFDM system.

Studies that model the indoor propagation characteristics [18,29,55] of wireless signals, however, show that in most of the cases the multi-path delay spread is in the order of nano-seconds, suggesting that cyclic prefix in practice is usually over-provisioned by a significant amount. In these measurement based studies, the power delay profile, which is the strength of the received signal plotted against time, is used to characterize the multipath channel. The time delay between the multipath arrivals is used to determine the maximum delay spread in the environment, which is in the order of nano-seconds for various environments [18, 29, 55]. Since the inter-symbol interference from an OFDM symbol on the following OFDM symbol is limited to the maximum delay spread, this suggests that the cyclic prefix is over-provisioned significantly in several environments.

Furthermore, the latest standards such as IEEE 802.11n/ac, support wider channel widths of up to 160MHz. With wider channels, as shown in Table 2.2, the duration of the cyclic prefix increases due to the increase in number of subcarriers. However, since the multipath delay spread is independent of the channel width, the number of samples that are unaffected by ISI (which is the portion of over-provisioned cyclic prefix) only increases with channel width.

3. OPPORTUNITIES IN CYCLIC PREFIX

In a standard OFDM system (illustrated in Fig. 3), the receiver discards the cyclic prefix before decoding the OFDM symbol. In this section we discuss the opportunities in retaining the cyclic prefix and using it to improve symbol decoding. We start by analyzing the effect of choosing different FFT windows on an OFDM symbol.

3.1 Sliding FFT Windows

Let us consider a discrete-time OFDM system, illustrated in Fig. 3. The system consists of F subcarriers onto which complex data symbols D_s are modulated using an inverse discrete Fourier transform (IDFT). Let vector

$$X_s = (X_s[0], \cdots, X_s[F-1])$$
$$\text{where,} \quad X_s[f] \in \mathcal{L} = \{l_1, l_2, \cdots, l_k\}$$

denote, in frequency domain, a complex vector representing the s^{th} OFDM symbol transmitted by the n^{th} user, and $\mathcal{L}$ denotes the finite set of alphabet from the transmitter's codebook, each corresponding to a lattice point. The time-domain representation of the OFDM symbol s transmitted by the n^{th} user is given by

$$x_s = (x_s[0], \cdots, x_s[F-1])$$

where,

$$x_s[t] = \frac{1}{F} \sum_{f=0}^{F-1} X_s[f] e^{i2\pi ft/F}, \quad 0 \le t < F$$

To eliminate the effects of dispersed channel distortion a cyclic prefix, which is a copy of a portion of the symbol, is prepended to each OFDM symbol. The time-domain signal with a cyclic prefix of size C transmitted by node n can be written as follows,

$$x'_s[t] = x_s[t \bmod F], \qquad -C \le t \le F-1$$

The received signal y_s for OFDM symbol s, contains $F+C$ samples, including the cyclic prefix of C samples. To perform DFT on the received signal, a segment of size F must be chosen with the rest of the C samples disregarded from y_s. Since there are P samples in the cyclic prefix that are not affected by ISI, as shown in Fig.2, there are P valid sampling windows which can be used to decode the data transmitted in symbol s. We refer to each of these P sampling windows as *segments*. After channel equalization, since these P segments are not affected by ISI, the signal received from the j^{th} segment at subcarrier f in OFDM symbol s can be written as,

$$\hat{X}_s^j[f] = \frac{1}{\hat{H}} \sum_{t=0}^{F-1} y_s^i[f] e^{-i2\pi(C-P+j)f/F} + \mathcal{E}_s^i[f] \qquad (1)$$

where, $\hat{H}$ is the estimated channel matrix and $\mathcal{E}_s^i[f]$ is the cumulative noise on that subcarrier from the environment and other interferers.

In the time domain, these different segments correspond to different cyclic shifts of the data transmitted in the OFDM symbol. However, this translates to a frequency dependent phase rotation in the frequency domain which can be com-

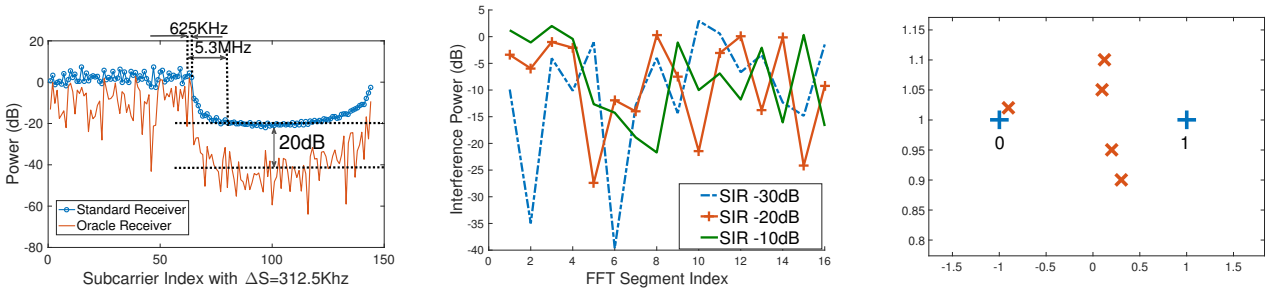

(a) (b) (c)

Figure 4: (a) Almost 20dB reduction in interference by choosing best FFT segment for each subcarrier; (b) Interference power in a subcarrier at different FFT segments showing significant variation; (c) Constellation plot showing two lattice points and received signal in 5 FFT segments illustrating problems with a simple metric.

puted (and easily corrected) for the segment j and subcarrier f as,

$$\theta^j[f] = e^{-i2\pi(C-P+j)f/F} \quad (2)$$

Hence, this predictable phase shift can be easily corrected to obtain P copies of the transmitted symbol.

PROPOSITION 3.1. *Choosing different FFT segments of an OFDM symbol does not affect the symbol except for a multiplicative phase shift due to the rotation in the time domain.*

3.2 Opportunities for Interference Mitigation

To understand the effects of interference in different FFT segments, we conduct real life experiments with USRPs and implement the OFDM system illustrated in Fig. 3. We consider the communication between an 802.11g access point and client in the presence of (adjacent/co-channel) interference. The transmitter is assigned a total of 64 subcarriers of 312.5KHz width and the duration of the cyclic prefix is fixed at 0.8 μs with 16 samples. To create a scenario with adjacent channel interference, contiguous subcarriers are assigned to the sender and interferer with 4 subcarriers as guardband in between. The interferer transmits the signal with a temporal offset that is greater than 0.8 μs, the duration of the cyclic prefix to create adjacent channel interference. To create co-channel interference, the interferer is assigned the same set of subcarriers used by the sender.

The key insight from analyzing the interference at the receiver is that the effect of interference varies significantly across the different FFT segments of the same OFDM symbol. For instance, an OFDM symbol received with -20dB SIR, is shown in Fig. 4a. In this scenario, the interferer occupies the subcarriers (68-132) adjacent to the sender (1 to 64) and due to a temporal offset greater than the duration of cyclic prefix, leaks energy into the adjacent bands distorting the sender's signal. The normalized interference power (obtained by muting the sender) at subcarrier 63 as seen by the receiver for different levels of SIR, over all 16 possible FFT segments of an OFDM symbol is shown in Fig.4b. It can be seen that the interference power varies significantly across the FFT segments. For instance, in the presence of adjacent channel interference with -30dB SIR, the interference power varies by almost 40dB, with the lowest at FFT segment 6.

Combining this insight with the fact that these P values for each subcarrier f have the same signal component as stated in Proposition 3.1, but are affected by a different interference component as shown in Eq. 1, it is clear that identifying the best FFT segments for each subcarrier can have significant benefits over discarding the cyclic prefix as done in existing OFDM based wireless systems.

First, minimizing interference power in each subcarrier would reduce the overall effects of interference, enabling signal decoding even in the presence of interference and it can be effective for different types of interference. In the example discussed above in Fig. 4b with SIR -30dB, a standard OFDM receiver would have discarded the cyclic prefix and selected the 16th FFT segment where the interference is almost 35dB stronger than in FFT segment 6. We refer to a scheme identifying the FFT segment yielding the lowest interference power as the Oracle scheme, which assumes perfect knowledge of the interference at the receiver. The interference power in different subcarriers with a standard OFDM receiver and an Oracle receiver are shown in Fig. 4a. The oracle scheme is able to reduce the effects of interference in the channel used by the sender by about 20dB as illustrated.

Second, the sharp spectrum mask realized by choosing the best FFT windows can reduce the number of subcarriers used as guard-band between contiguous bands assigned to neighboring transmitters. This means cognitive users can be allocated frequencies that are much closer to incumbents, improving efficiency of spectrum use. For instance, from Fig. 4a, the spectrum mask realized using the oracle scheme (shown in red) is very sharp compared to the vanilla case of not using any adjacent channel interference mitigation mechanism (shown in blue). And the required guard band is significantly reduced from 5.3MHz to just 625KHz for an adjacent channel interference threshold of about -20dB, enabling efficient use of the spectrum.

However to exploit these opportunities, we need to be able to decode the received symbol in the presence of P redun-

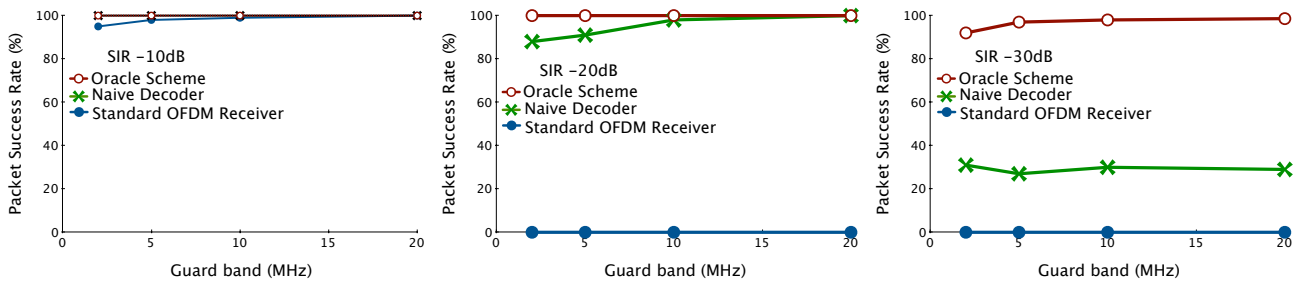

(a) SIR -10dB (b) SIR -20dB (c) SIR -30dB

Figure 5: Packet success rate using Oracle scheme and the naive decoder showing deteriorating performance when interference increases; experiment settings: single adjacent channel interferer, QPSK with 3/4 coding rate, varying guard band and SIR values.

dant copies of the signal and there are several challenges that must be overcome first.

3.3 Challenges with Decoding

It is not practical to possess perfect knowledge of interference at the receiver, without which the FFT segments with the minimum interference power cannot be identified. The Oracle scheme while it provides a clear picture of the opportunities for mitigating interference, is thus impractical.

Using simple statistical metrics to decode OFDM symbols using P redundant copies is not effective and as a result underlying opportunities for mitigating interference may be squandered using them. To understand this, we define a naive decoder to identify the closest lattice point around which the signal received in different FFT segments is scattered. For each subcarrier we compute the average deviation of the received complex vector from the various possible lattice points for the modulation scheme used, over all the FFT segments. Then the lattice point with the minimum average deviation is assumed to be the correct one [38]:

$$l* = \arg\min_{l \in \mathcal{L}} \sum_{i=1}^{P} |\hat{X}_s^i[f] - l| \quad (3)$$

To evaluate the naive decoder, we use USRPs and the same WiFi settings described above for the experiments. We vary the SIR for different modulation schemes and the packet error rates for QPSK modulation is shown in Fig. 5 for different guardband sizes. As expected, the metric performs well at lower interference power. When the SIR is about -10dB, both the Oracle scheme and the naive decoder are able to eliminate the packet errors. However, at SIR -20dB, while the Oracle scheme is able to decode all the packets successfully, using the naive decoder only results in marginal improvements. In the presence of strong interference (with SIR less than -10dB), the shortcomings of the naive decoder are apparent. The performance of the oracle scheme with strong interference shows that there are FFT segments where the received signal can be successfully decoded, however, the naive decoder is unable to find the right lattice points. In analyzing the scenarios where the naive decoder fails, we identify three main sources for these errors.

To illustrate this, we use an example scenario shown in Fig. 4c, with the set of possible lattice points of the transmitted signal (blue plus marker) and the signal received in different FFT segments (red cross marker). For simplicity, we consider only two lattice points (BPSK) and $P = 5$ (five FFT segments are used for decoding). In this instance, the transmitted signal corresponds to lattice point 1, and due to varying interference in different FFT segments, the received signal is scattered around the transmitted lattice point. To illustrate different scenarios where errors occur, we consider that one of the FFT segments suffers from strong interference and the received signal is close to lattice point 0 even though the transmitted signal corresponds to lattice point 1.

The first source of the error is the use of arithmetic mean in the naive decoder to determine the central tendency of the signal received in different FFT segments. It is well known that arithmetic mean is susceptible to outliers making it ineffective either due to a small sample size or if the underlying distribution is skewed. In the example discussed above, the received signal in four of the five FFT segments are closer to lattice point 1. However, due to a single outlier, on an average the five points are closer to lattice point 0 and hence are incorrectly identified. The small number of ISI free FFT segments further increase the proportion of these outliers.

Second the naive decoder assumes that the received signal from different FFT segments are on the correct lattice point. At the receiver, the signal corresponding to a lattice point would have been affected by fading and AWGN noise due to the wireless medium. The constellation decoders work under the assumption that of the received signal with the effects due to fading and noise perfectly removed would be exactly one of the lattice codes. However, this is not true in the presence of interference. With interference affecting each of the FFT segments apart from fading effects and channel noise, the received signal would be at a certain distance from the correct lattice point. In the example discussed above, four of the five points are at a similar distance away from the lattice point. Instead, if the received signal is expected to be at a certain distance from lattice point 1 then the decoder has a better chance of identifying the outlier.

Finally, the naive decoder completely ignores the phase errors due to interference. It only takes into account the am-

plitude effects of interference by computing the Euclidean distance between the lattice points. Phase noise can be introduced due to the fluctuation of the oscillators in the transmitters and the performance degradation [36, 45, 54] of OFDM systems in the presence of phase noise has been well studied. The example discussed above shows such a case, where the same phase error on the outlier would have a much larger change in the euclidean distance between the outlier and the lattice points.

4. CPRECYCLE

Considering the aforementioned issues with using simple statistical metrics for decoding, we design CPRecycle , a novel OFDM receiver that creates an interference model from the preambles to effectively utilize the opportunities provided by the redundant samples in the cyclic prefix.

4.1 Modeling the effect of interference

In OFDM based systems, the symbols with data is usually preceded by one or more training symbols of known data called preambles for channel estimation and synchronization. These preambles typically use a robust modulation scheme that can be decoded even at low SNR values. In CPRecycle receiver, using the P ISI free segments of each of the preambles, P complex values are generated for each subcarrier with every preamble symbol. These P complex values can be used to create a model of the interference effects. We now discuss the various issues that needs to be addressed in generating such a model.

The first hurdle with using the preambles to generate a model is that the modulation schemes in preambles and the data symbols could be different. Lattice codes are generated by selecting a finite number of points from a two dimensional Euclidean space $\mathcal{R}^n$ depending on the modulation scheme. Hence the received signal in the preambles cannot be directly used to create a model for the data symbols to use. To facilitate this, we compute the variations of the received signal in different FFT segments relative to the lattice point being considered. It can now be applied to a signal corresponding to any lattice point and hence used across different modulation schemes.

Another issue is the limited number of samples that are available to create and use the interference model. Most of the OFDM standards use at most two preambles for channel estimation and in each preamble the maximum number of samples for each subcarrier is the number of samples in the cyclic prefix. Furthermore, since the receiver does not possess any information about the interference, it is not accurate [28] to assume a standard probability distribution (e.g., Gaussian). Hence care has to be taken to design a non-standard probability distribution that works well with a small sample size.

Finally, there is the need to decouple the amplitude and phase effects of interference in different FFT segments, mainly because there is no correlation between them. In scenarios with strong interference, it is reasonable to expect that the interfering signal is carrying data that either amplitude or phase modulated. In such cases too it is beneficial to consider phase errors independently. Also, with amplitude and phase errors decoupled, a weighted function can be used to tune the impact of these errors to improve the accuracy of the interference model.

Based on the issues discussed above, to effectively utilize the opportunities provided by the redundant samples in the cyclic prefix, we need a non-parametric density estimation from the amplitude and phase changes in the different FFT segments that works well with a small sample set.

The simplest method to estimate the probability density of the interference is to use bins of constant or variable width in phase and amplitude and construct a bivariate histogram. However, there are two main problems with using bivariate histograms to model the effect of interference in our context: (i) with a small sample set there are discontinuities in the estimated density due to empty bins; (ii) it assumes that there is no relation between the data in adjacent bins.

So we instead employ a more effective alternative called kernel density functions [40, 42, 47] to generate a non parametric density. Unlike histograms, kernel density functions does not have discontinuities and can produce a smooth distribution with a small sample set. Furthermore, the amplitude and phase changes can be integrated using a bivariate product kernel density function where the weight for amplitude and phase variations can be tuned.

In order to generate a probability density function with the preamble data in each subcarrier, we use a bivariate gaussian product kernel density estimation function with a variable bandwidth. Let $R_A^j[f]$ and $R_\phi^j[f]$ denote the set of amplitude and phase variation values observed on a subcarrier f, $1 \leq j \leq P$, from the preambles, which can be computed as:

$$R_A^j[f] = \mathcal{A}(\hat{X}_s^j[f] - X_s[f]), \quad 1 \leq s \leq N_p, \quad 1 \leq j \leq P$$
$$R_\phi^j[f] = \Phi(\hat{X}_s^j[f] - X_s[f]), \quad 1 \leq s \leq N_p, \quad 1 \leq j \leq P$$

where N_p is the number of preambles used in modeling the interference. Then the probability density function can be written as:

$$f_m(a_{obs}, \phi_{obs}) = \frac{1}{P * N_p} \sum_{j=1}^{P*N_p} \Big[K_a(\frac{a_{obs} - R_A^j[f]}{B_a}) \times K_\phi(\frac{p_{obs} - R_\phi^j[f]}{B_\phi}) \Big]$$

where,

$$K_a(a) = \frac{1}{2\pi} e^{-a^2/2} \text{ and } K_\phi(p) = \frac{1}{2\pi} e^{-p^2/2} \tag{4}$$

B_a and B_ϕ are the kernel-bandwidths which are smoothing parameters that determine the range of amplitude and phase over which the sample points are averaged to generate the probability density.

It is well known that the choice of the kernel-bandwidths has a significant impact [21] on the accuracy of density estimation and it is crucial to identify right value. To illus-

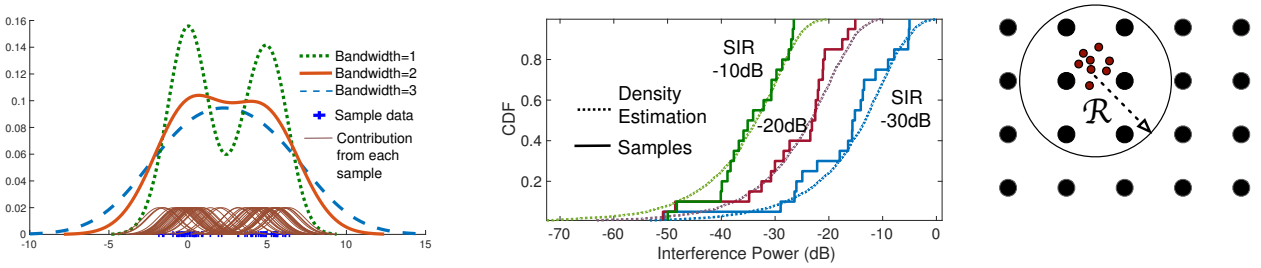

(a) (b) (c)

Figure 6: (a) Kernel density estimation with varying bandwidth; (b) Density estimates and samples of amplitude variations showing accurate modeling for different SIR scenarios; (c) Illustration of lattice points with a sphere of radius $\mathcal{R}$ centered at the centroid of signal received in 7 FFT segments.

trate this consider an example of a set of amplitude variations along with the kernel density function with three different bandwidths shown in Fig. 6(a). Larger bandwidths result in over smoothing of the density estimate and smaller bandwidths introduce large errors between the data points. In general, it is beneficial to have a larger bandwidth at low densities and a smaller bandwidth at high densities of data.

In CPRecycle , we use the data driven approach to determine the best bandwidth which is possible in the presence of at least two preambles. The Gaussian kernel density function shown above generates a smooth bivariate density function and the probability density function is recomputed each time a new set of preambles are received.

The density estimation of amplitude variations and the variations observed in the data symbols, for different SIR values are shown in Fig. 6(b). The kernel density functions accurately predict a density that is applicable for the amplitude variations in the data symbol.

4.2 Maximum likelihood decoding

To decode the received symbols $\hat{X}_s[f]$ to the correct lattice point, we use a maximum likelihood decoder that identifies the lattice point with the maximum probability of the received symbol corresponding to that point. When only one FFT segment is used, the maximum likelihood decoder reduces to a minimum Euclidean distance decoder that identifies the codeword that is closest to the received symbol. However, with CPRecycle receiver, each symbol transmitted on a subcarrier results in P received symbols, one per FFT segment.

Let the transmitted symbols $X_s[f]$ be drawn from a known finite alphabet $\mathbb{L} = \{l_1, l_2, \cdots, l_k\}$ each corresponding to a point in the lattice. The maximum likelihood decoder can be defined as:

$$l^* \triangleq \underset{X_s[f] \in \mathbb{L}}{\operatorname{argmax}} \quad \mathcal{P}(X_s[f]|\hat{X}_s[f]) \tag{5}$$

where

$$\mathcal{P}(X_s[f]|\hat{X}_s[f]) = \prod_{j=1}^{P} \frac{\mathcal{P}(\hat{X}_s^j[f]|X_s[f])}{\mathcal{P}(X_s[f])}$$

where $\mathcal{P}(X_s[f])$ sent is constant and $\mathcal{P}(\hat{X}_s^j[f]|X_s[f])$ can be computed from the probability density function defined in Eq. 4 as follows:

$$\mathcal{P}(\hat{X}_s^j[f]|X_s[f]) = f_{X_s[f]}(\mathcal{A}(\hat{X}_s^j[f] - X_s[f]), \Phi(\hat{X}_s^j[f] - X_s[f]))$$

With higher modulation schemes the search space for the decoder increases exponentially with the number of lattice points (as 2, 4, 16, 64, 256 for BPSK, QPSK, 16QAM, 64QAM, and 256QAM respectively). Hence it is essential to reduce the number of possible lattice points for comparision. In CPRecycle , to select a subset of possible lattice points we use the concept of a fixed sphere decoder.

The concept of a fixed sphere has been shown to be effective [5, 8, 50] to reduce the search space in identifying the closest lattice point. For single antenna receivers, the decoder searches through the lattice points that are located within a sphere of radius $\mathcal{R}$ centered at the received signal. However, a slight variation is required in our case since the decoder receives P signal values from which the lattice points need to be identified, instead of one.

In CPRecycle , to identify the point around which the sphere is centered, we compute the centroid of the cluster of P complex signal values. The centroid is simply the average of the real and imaginary values of all the P values. Only the subset of lattice points that fall in the sphere of radius $\mathcal{R}$ from the centroid of the P samples constitute the search space for the decoder. The choice of lattice points for a sphere decoder is illustrated in Fig. 6(c). In this instance, only the six lattice points that fall within the sphere are considered as possible transmitted codes by the decoder. This significantly reduces the number of operations required in decoding the received symbol.

4.3 Putting It All Together

Algorithm 1 shows the overall procedure followed by CPRecycle receiver from putting the above components together. When CPRecycle receiver receives a preamble, it computes the number of ISI free samples in the CP to determine P. The P segments in the preamble are used to generate a unique probability density function for each subcarrier. These prob-

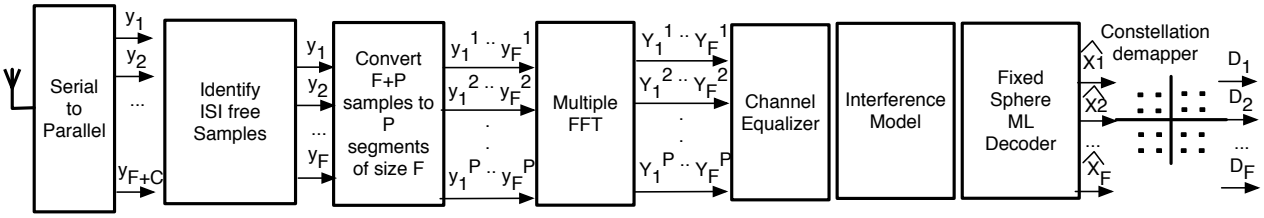

Figure 7: Block diagram of CPRecycle receiver as implemented.

ability density functions are constantly updated when subsequent preambles are received. Once the interference is modeled, the subsequent OFDM symbols are decoded using the maximum likelihood decoder. The set of lattice points over which the maximum likelihood detector searches for the transmitted symbol is computed using the radius R which is an input parameter to the CPRecycle receiver.

Algorithm 1 CPRecycle Receiver

procedure CPRECYCLE(B_a,B_ϕ,$\mathcal{R}$,$\hat{X}_s$,P)
  if OFDM Symbol s is a preamble **then**
    for each segment $j \in \{1, 2, \cdots, P\}$ **do**
      $Y_s^j = FFT(y_s^j)$
      $R_{\mathcal{A}}^j[f] = \mathcal{A}(Y_s^j[f] - X_s[f])\ , \forall f \in \mathcal{F}$
      $R_{\phi}^j[f] = \Phi(Y_s^j[f] - X_s[f])\ , \forall f \in \mathcal{F}$
    end for
  else
    for each subcarrier $f \in \mathcal{F}$ **do**
      $L^c \subset L$ s.t. $\forall l \in L$
      $l \in L^c$ if $\mathcal{A}(l - \text{Centroid}(\hat{X}_s[f])) < \mathcal{R}$
      $X_s[f] = \underset{l \in \mathbb{L}^c}{\text{argmax}} \quad \mathcal{P}(l|\hat{X}_s[f])$
    end for
  end if
end procedure

Note that from the above description it is clear that CPRecycle receiver does not need to explicitly know the precise nature of interference (e.g., adjacent channel interference, co-channel interference). It can leverage the preambles used for channel estimation. The effectiveness of CPRecycle relies on the extent to which channel and interference characteristics seen from a preamble apply to the subsequent OFDM symbols with data. For rapidly varying or sporadic interference, more frequent preambles are needed to accurately model the interference.

5. EXPERIMENTAL EVALUATION

In this section, we experimentally evaluate the effectiveness of CPRecycle in mitigating different types of interference that can be mapped to practical scenarios.

5.1 Implementation

We have implemented a prototype of the CPRecycle receiver using the USRP radio platform [11], Ziria [44] an SDR programming environment, and the GNU Radio software package [3].

CPRecycle Receiver: We implement two variants of the CPRecycle receiver to run on the USRP: (i) IEEE 802.11g receiver (ii) A generic configurable OFDM baseband receiver. For the IEEE 802.11g receiver, we modify the GNU Radio based receiver [3] as shown in Fig. 7. Instead of discarding the CP, the ISI free portion of the CP is used to generate P segments that are then passed on to the FFT block to generate P values for each subcarrier corresponding to the signal transmitted on an individual subcarrier. The maximum likelihood decoder then detects the signal transmitted on the subcarrier using the interference model generated with the preambles for each subcarrier.

IEEE 802.11g Setting : The CPRecycle receiver is applicable to IEEE 802.11a/g/n radios which are based on an OFDM PHY. For our experiments we use an off-the-shelf 802.11g Linksys access point running tomato firmware. Each 20MHz channel is composed of 64 subcarriers, spaced 312.5 KHz apart, of which 52 subcarriers are used for data and 4 subcarriers for pilots. Each OFDM symbol has a duration of 4μsec. and each data payload is preceded by a long training field that contains two OFDM symbols for a duration of 8μsec, to enable synchronization and channel estimation. The variation of the signal in different segments in this long training field is used to create the interference model. For our experiments, we choose three MCS modes, QPSK 1/2 (9 Mbps), 16-QAM 1/2 (24 Mbps), and 64-QAM 2/3 (36 Mbps).

5.2 Results

We now evaluate the performance of CPRecycle in the presence of adjacent channel interference and co-channel interference.

5.2.1 Adjacent Channel Interference

Single Interferer. For the adjacent channel interference case, we use an off-the-shelf 802.11g access point (Linksys) that continuously transmits 400 byte packets, in channel 11 (2462MHz). To generate interference, we use a USRP (B210) to continuously transmit 802.11 traffic in an overlapping channel, in this case channel 8 (2447MHz). A CPRecycle receiver running on another USRP B210, that is capable of decoding 802.11g packets is placed in a fixed location. To choose the appropriate SNR for each MCS, the Linksys router is re-positioned from the receiver until that MCS mode has

the highest throughput. Once the SNR for the MCS mode is fixed, the SIR is varied by moving the interferer that generates 802.11 packets in the adjacent channel. We transmit a total of 2000 packets for each scenario and the average values for packet success rate is shown in Fig. 8.

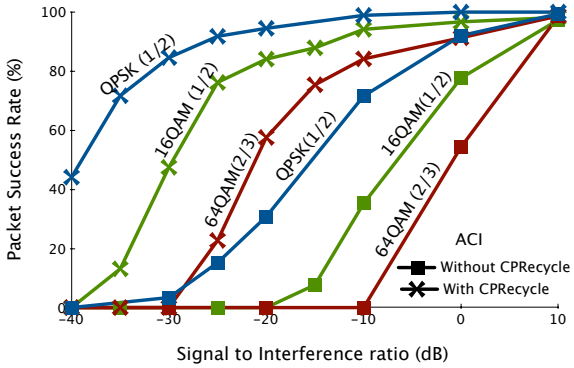

Figure 8: Packet success rates for different modulation and coding schemes with one adjacent channel interferer

The severity of the effect of adjacent channel interference on the packet success rates can be seen from the figures. At an SIR value of 0dB, where the power of the signal and the interference is the same, the success rates of packet delivery drops significantly for all MCS modes. Being the highest rate, 64QAM suffers almost 50% packet loss and is unable to transmit a packet when SIR is -10dB. This effect is slightly less pronounced for the lower modulation schemes such as QPSK, however, the increase in packet loss rate with SIR is still steep, and becomes unusable when SIR decreases to 10dB.

With the CPRecycle receiver, the packet success rates are significantly improved for all the MCS schemes with similar packet success rates achieved with atleast 15dB of adjacent channel interference and in several cases upto 25dB of adjacent channel interference for lower modulation schemes. Considering the packet success rates at -10dB, for example, it can be observed that for all MCS modes the improvement in packet success rates is significant and with higher modulation schemes communication is made possible (with almost 80% packet delivery rate) that would not have otherwise been possible without CPRecycle receiver.

Multiple Interferers. The effect of two interferers creating adjacent channel interference on either side of the channel allocated to a transmitter is shown in Fig. 9. For this experiment, the Linksys access point is allocated channel 10 (2457MHz) and the interferers are allocated channels 7 (2442MHz) and 13 (2472MHz) respectively. This is a common scenario in dense deployments of WLANs where overlapping channels has to be allocated to neighboring access points. The packet success rates is noticeably lower for all the modulation schemes, since the number of subcarriers that are affected by adjacent channel interference is almost doubled. However since the interference model is maintained independently per subcarrier, it does not have a significant impact on the performance of the CPRecycle receiver. For example, when the SIR is -10dB, CPRecycle is able to decode more than 80% of the packets successfully in most of the cases.

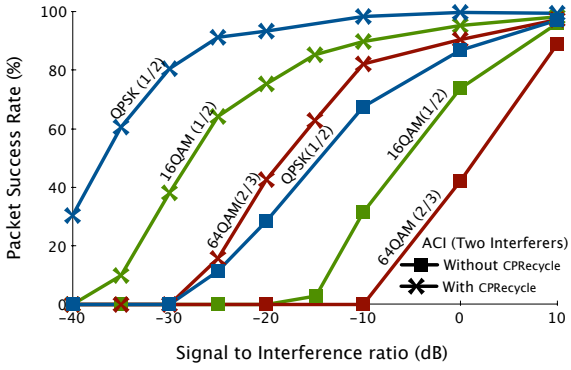

Figure 9: Packet success rates with two adjacent channel interferers

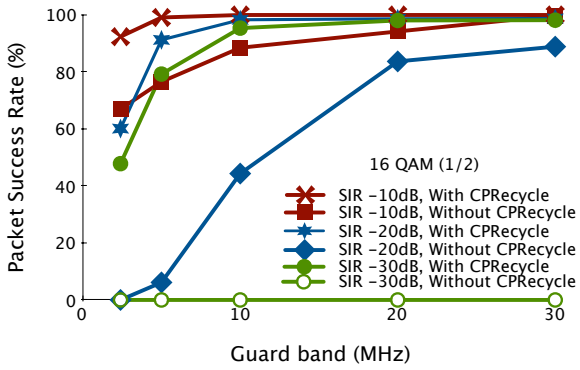

Figure 10: Packet success rates with varying guardband sizes with adjacent legacy transmitter

Guard band needed with adjacent legacy OFDM transmitter. The effect of adjacent channel interference with different sizes of guard-bands for 16QAM is shown in Fig. 10 respectively. For this experiment, the set of subcarriers assigned for the first transmitter is fixed and the set of contiguous subcarriers assigned to the second transmitter is varied to generate settings with different guard-bands between the two transmitters. It can be observed that with CPRecycle the amount of guard-bands required to achieve the same packet success rates is significantly lower for both the modulation schemes. This shows that with CPRecycle , a cognitive radio can be allocated frequencies much closer to a licensed band achieving a significantly more efficient use of the wireless spectrum. For example, considering the case with 16QAM, if a cognitive user is allocated a cluster of subcarriers adjacent to a licensed TV transmitter, whose signal is 10 times stronger, then the required guard-band would be reduced from about 15MHz to less than 5MHz to achieve a similar packet success rate.

5.2.2 *Co-Channel Interference*

Single Interferer. To generate co-channel interference, we use a setup that is similar to the adjacent channel interference scenario, except, we use a USRP 802.11 transmitter. This is so clear channel assessment can be turned off to enable simultaneous use of the same channel by both the transmitter and the interferer. Similar to the adjacent channel interference case, the SNR for each MCS mode is chosen such that any higher modulation scheme would result in a lower throughput. In total, 2000 packets of size 400 bytes, are transmitted for each scenario for each MCS mode and a given SIR setting, and the average packet success rates are computed. The results are shown in Fig. 11.

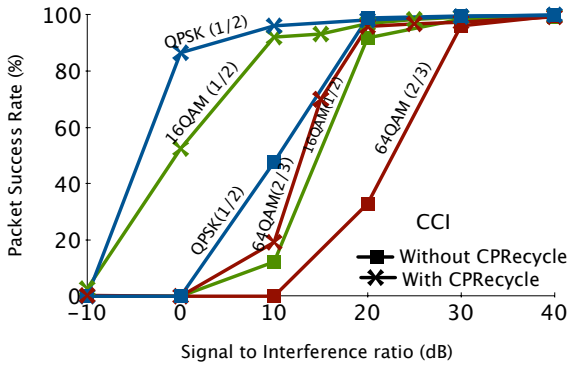

Figure 11: Packet success rates for different modulation and coding schemes with single co-channel interferer

As expected, the effect of co-channel interference on 802.11 WLANs is far more severe than adjacent channel interference, which is evident from the figures. Even with SIR 10dB, when the signal of interest is three times stronger than the interference, the packet reception rate drops steeply for all the MCS schemes. This is mainly due to two reasons. (i) Unlike adjacent channel interference, the co-channel interference is in-band. (ii) The number of subcarriers affected by interference is much higher in the co-channel interference scenario. In most cases all the subcarriers used by the transmitter is affected by strong interference.

Another observation is the steepness of the drop in packet reception rates with increasing co-channel interference. The range of co-channel interference tolerated by both with and without CPRecycle receiver is about 15dB in most cases, where as it was about 30dB of adjacent channel interference for most MCS modes. This is mainly due to the significantly higher number of subcarrier affected by interference when compared to adjacent channel interference. However, CPRecycle is able to recover most of these errors since it maintains a separate interference model for each subcarrier from the preamble data.

Multiple Interferers. The effect of multiple co-channel interferers is shown in Fig.12. For this experiment, we setup an 802.11 transmitter with carrier sensing disabled, and two interferers in the same channel, placed at the same distance from the transmitter. The SNR is chosen for each MCS mode

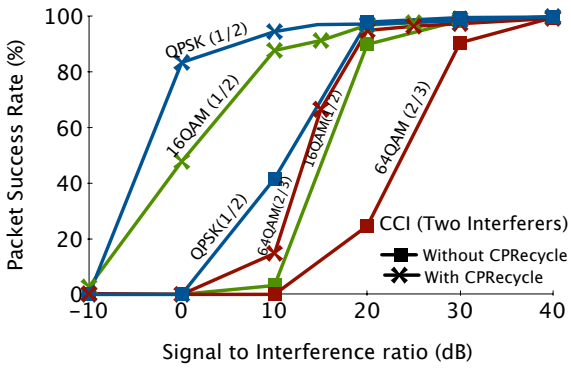

Figure 12: Packet success rates with two co-channel interferers

similar to the other experiments. The SIR is varied by increasing the transmit power in both the interferers. It can be observed that unlike in the case of adjacent channel interference, co-channel interference does not have a significant impact on packet reception. This can be attributed to the fact that the number of subcarriers affected by the higher number of interferers does not change where as it almost doubles in the case of adjacent channel interference. The improvement in packet success rate with CPRecycle is again significant even though the variance of interference is presumably higher with more interferers, while the total power of the interference remains the same. This is primarily due to the nature of the interference model that considers both amplitude and phase changes in the interference to generate the probabilistic model for each subcarrier.

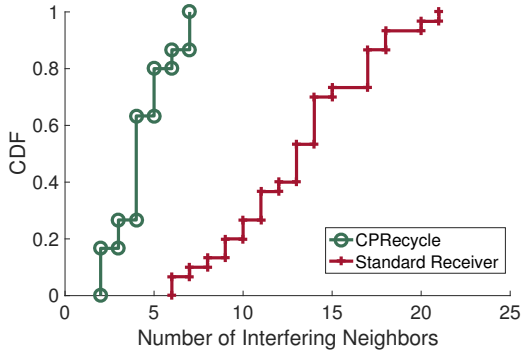

Figure 13: CDF of number of interfering neighbors for access points in a real office environment with and without CPRecycle receiver.

Network Level Improvements. While it is clear that CPRecycle can decode signals even in the presence of strong interference, the network level benefits of this are not obvious. To highlight this, we plot the CDF of number of interfering neighbors for access points in a real indoor office environment shown in Fig.13. From Fig.11, it is evident that with the CPRecycle receiver, the level of co-channel interference

that can be tolerated is atleast 15dB for all the MCS modes. This is a direct measure of the increase in energy detection threshold that the CPRecycle receiver would be able to tolerate without additional packet errors.

We consider our office building [32] which has five floors with a large atirum and most of the walls are made of glass. There are 40 access points deployed in the building with mostly the same place for access points in each floor. We measure the signal strength of access points that can be detected at each of these locations and determine the number of neighbors for the access points by reducing the threshold by 15dB derived from Fig.11. It can be seen that the number of neighbors with CPRecycle is significantly reduced. For instance, with a standard receiver, more than 80% of access points have atleast 12 interfering neighbors where as with CPRecycle more than 80% of the access points have utmost 6 neighbors. This shows how CPRecycle can significantly improve the network capacity of a dense WLAN by reducing the potential interferers in the network.

6. DISCUSSION

Detecting ISI free portion of CP : Several methods [4, 37, 43, 57] have been proposed in the literature for the detection of ISI-free region in the cyclic prefix. In each of these schemes a correlation coefficient is computed between samples in a given window and a threshold is used to estimate the range of ISI free samples in the CP.

The effect of the duration of the ISI free region over the performance of CPRecycle is shown in Fig. 14, where the number of FFT segments represents the duration of the ISI free region. A key observation here is that even when a significant portion (about 60%) of the cyclic prefix is affected by ISI, CPRecycle is able recover a significant percentage of the erroneous packets. This suggests that CPRecycle can even be used in multipath environments with a significant delay spread.

Computational Complexity and Oversampling: The computational complexity of CPRecycle is O(PN_p^2f), where P is the number of ISI free samples in the CP, N_p is the number of preambles and f is the number of subcarriers. Since the number of preambles is not a configurable parameter, we study the effect of P, the number of samples. We conduct experiments for the ACI scenario with varying number of FFT samples, with five preambles to observe the behavior of CPRecycle .

The packet success rate for three different SIR conditions with 16QAM modulation is shown in Fig. 14. An interesting behavior we observe with the number of FFT segments is that, the benefits of the increasing the number of FFT segments for interference modeling saturates when P reaches about 60% of the samples even at very high interference (SIR -30dB). With lower levels of interference, even 20% of the CP is enough to reduce the packet error rates significantly. There are two advantages to this behavior with CPRecycle : (i) scenarios with high multi-path delays where the number of ISI free samples in the CP is limited, can still make use of CPRecycle to improve the performance of the receivers. (ii) on devices with limited computational capability the number of FFT segments can be tuned to the capabilities of the device, which gracefully degrades to a standard OFDM receiver with one FFT segment, in the worst case. Hence it can be used in a wide variety of hardware configurations with varying computational capabilities.

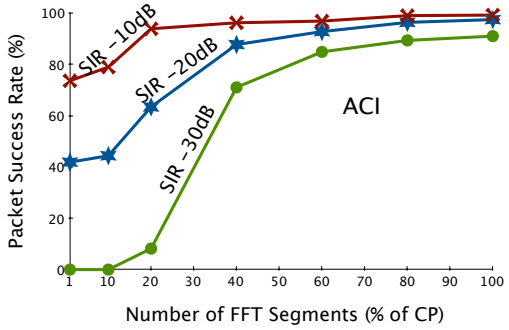

Figure 14: Packet success rates with varying number of FFT segments

When unconstrained by computational capability, it is also beneficial to increase P beyond the number of ISI free samples available in the CP. This is possible through oversampling with new devices that support higher sampling rates.

7. RELATED WORK

Adjacent Channel Interference: OFDM systems are known to suffer from high levels of out-of-band emissions. Several techniques have been proposed to reduce this out-of-band radiation. Windowing is a time domain technique [51], where the signal is multiplied with a windowing function before transmission to reduce the energy in the side lobes. Techniques such as Subcarrier weighing [6], Multiple-choice sequences [7], Cancellation carriers [41], constellation Expansion [35], and Adaptive symbol transition [31] are some of the techniques that manipulate the frequency domain signal at the transmitter to enable out of band reduction. A comprehensive comparison of these side lobe reduction has been presented in [24]. One of the defining features of adjacent channel interference is that only the subcarriers the band assigned to a transmitter is affected. The schemes that suppress ACI are designed to mitigate the interference in the edge subcarriers and hence are not suitable to suppress other types of interference such as co-channel interference.

In LTE, fractional frequency reuse and adaptive power management techniques are used to reduce the level of interference in the network. Active interference mitigation schemes such as interference rejection combining [27], coordinated multipoint transmission (COMP) [26], and channel coding are being used to mitigate co-channel interference.

Co-Channel Interference: Co-channel interference management techniques can be grouped into two categories (i) schemes that mitigate interference by modifying the trans-

mitted signal (ii) schemes that decode the signal of interest in the presence of interference.

Interference mitigation schemes such as [20, 52], adapt the transmissions to be more resilient to interference. Their application is limited to niche scenarios and moreover they require changes to the existing standards and are not backward compatible. Interference alignment [1, 2, 15] is a recently proposed technique that in this category. However, they require communication over the wired backbone and are not backward compatible. Similarly, Swarun et al, propose OpenRF [25] a cross-layer architecture for interference management that enables access points to cancel their interference at the clients significantly improving the network capacity and is applicable only to multiple antenna systems.

Several schemes have been proposed to decode the signal of interest in the presence of co-channel interference. Kong et al [23] propose MZig, a physical layer technique to decode simultaneous transmissions from multiple ZigBee devices to provide an m-fold increase in throughput. Gollakota et al, propose TIMO [14], an IEEE 802.11n receiver that can decode the packets in the presence of cross-technology interference. While TIMO can work even when the interference is persistent and lasts over a few seconds, unlike, CPRecycle , it can only be applied to receivers with multiple antennas. Yan et al, propose WizBee, [56], a ZigBee receiver that can decode ZigBee packets in the presence of strong interference from 802.11 nodes limiting its application.

In contrast to the interference management schemes discussed above, CPRecycle can mitigate different types of interference on single antenna systems and is also backward compatible with legacy OFDM systems.

Partial Packet Recovery: Partial packet recovery is a class of techniques that attempt to recover corrupt packets instead of retransmitting them. Several approaches [16, 19, 20, 22, 33, 53] have been proposed to address this inefficiency in retransmitting an entire packet due to a few bit errors. They can broadly be categorized into (i) co-operative packet recovery and (ii) cross-layer packet recovery.

In co-operative packet recovery schemes such as [22, 30, 33, 53] multiple access points coordinate with each other to recover partially corrupted packets by exploiting receiver diversity. SOFT and MRD use PHY layer information to identify corrupt blocks of bits that needs to be transmitted. ZipTx uses adaptive FEC codes to improve probability of repairing bit errors in co-operation with other APs in the vicinity. However, co-operative packet recovery techniques (including MRD and SOFT) demand additional constraints such as multiple coordinating APs, hardware changes, incompatible with IEEE 802.11, and hence are not useful in scenarios where CPRecycle is applicable. These techniques are useful in 802.11 mesh networks, where only the correct bits of a packet are forwarded on and the receiver combines multiple such copies to recover the entire packet, and in scenarios where these multiple APs coordinate through a wired backbone to share partial packets with the receiver.

Cross-layer partial packet recovery techniques such as [19, 20] attempt to recover partially corrupt retransmissions of the same packet and are in a way extensions of the chase combining decoder (where multiple noisy copies of a packet are combined to recover the packet). These techniques however require modification at both the transmitter and receiver to use additional parity bits to identify corrupt blocks for retransmission.

Furthermore, both categories of partial packet recovery techniques are complementary to CPRecycle and can be used in combination to improve the packet reception rate further. For example, SOFT and PPR use a confidence measure on decoding a bit as '0' or '1' cooperatively with multiple APs to improve decoding accuracy. When used in combination with CPRecycle, it would receive higher confidence measures on the decoding decision since CPRecycle exploits multiple copies of the signal in the cyclic prefix to select the FFT window with the signal that is closest to the correct lattice point.

8. CONCLUSIONS

In this paper, we have considered the problem of mitigating different types of interference experienced by OFDM based wireless systems. Exploiting the fact that OFDM based wireless standards over-provision the cyclic prefix (CP) that is meant for preventing inter-symbol interference, we presented a novel OFDM receiver design called CPRecycle that takes advantage of the redundant portion of the CP towards interference mitigation. Specifically, CPRecycle models the effect of interference in each subcarrier using a Gaussian kernel density function using the preamble symbols and uses a fixed sphere maximum likelihood detector to decode the following data carrying OFDM symbols subject to interference. Using off-the-shelf IEEE 802.11g transmitters and interferers, we experimentally show the effectiveness of CPRecycle for mitigating adjacent-channel interference and co-channel interference. We also show that two preambles and small portion of CP are sufficient to realize significant benefits in terms of packet success rate with CPRecycle . The application of CPRecycle in multiple antenna systems and evaluation of CPRecycle in the presence of different types of interference is left for future work.

9. ACKNOWLEDGMENTS

We thank our shepherd Michael Mitzenmacher and the anonymous reviewers for many helpful suggestions on improving the paper. We also thank Rik Sarkar for many helpful discussions on the statistical methods used in error correction mechanisms.

10. REFERENCES

[1] ADIB, F., KUMAR, S., ARYAN, O., GOLLAKOTA, S., AND KATABI, D. Interference alignment by motion. In *Proceedings of the ACM MobiCom* (2013), pp. 279–290.

[2] BANSAL, T., ZHOU, W., SRINIVASAN, K., AND SINHA, P. Robinhood: sharing the happiness in a wireless jungle. In *Proceedings of ACM HotMobile* (2014), p. 22.

[3] BLOESSL, B., SEGATA, M., SOMMER, C., AND DRESSLER, F. An IEEE 802.11 a/g/p OFDM receiver for GNU radio. In *Proceedings of Software radio implementation forum* (2013).

[4] CHEN, X., ZHANG, C., AND LUO, Y. Low-complexity ISI-free region detection for OFDM systems in high-mobility fading channels. In *Proceedings of IEEE HWMC* (2014).

[5] CONWAY, J. H., AND SLOANE, N. J. A. *Sphere packings, lattices and groups*, vol. 290. Springer Science & Business Media, 2013.

[6] COSOVIC, I., BRANDES, S., AND SCHNELL, M. Subcarrier weighting: a method for sidelobe suppression in OFDM systems. *Communications Letters, IEEE 10*, 6 (2006), 444–446.

[7] COSOVIC, I., AND MAZZONI, T. Suppression of sidelobes in OFDM systems by multiple-choice sequences. *European transactions on telecommunications 17*, 6 (2006), 623–630.

[8] DAMEN, M. O., EL GAMAL, H., AND CAIRE, G. On maximum-likelihood detection and the search for the closest lattice point. *Information Theory, IEEE Transactions on 49*, 10 (2003), 2389–2402.

[9] DING, Y., HUANG, Y., ZENG, G., AND XIAO, L. Channel assignment with partially overlapping channels in wireless mesh networks. In *Proceedings of ICWI* (2008), p. 38.

[10] DRAVES, R., PADHYE, J., AND ZILL, B. Routing in multi-radio, multi-hop wireless mesh networks. In *Proceedings of MobiCom* (2004), pp. 114–128.

[11] ETTUS, M. Usrp users and developers guide. *Ettus Research LLC* (2005).

[12] FARSHAD, A., MARINA, M. K., AND GARCIA, F. Urban WiFi characterization via mobile crowdsensing. In *Proceedings of NOMS IEEE* (2014), pp. 1–9.

[13] FENG, C., CUI, H., MA, M., AND JIAO, B. On statistical properties of co-channel interference in OFDM systems. *Communications Letters, IEEE 17*, 12 (2013), 2328–2331.

[14] GOLLAKOTA, S., ADIB, F., KATABI, D., AND SESHAN, S. Clearing the RF smog: making 802.11 n robust to cross-technology interference. *ACM SIGCOMM CCR 41*, 4, 170–181.

[15] GOLLAKOTA, S., PERLI, S. D., AND KATABI, D. Interference alignment and cancellation. In *ACM SIGCOMM Computer Communication Review* (2009), vol. 39, ACM, pp. 159–170.

[16] GOWDA, M., SEN, S., CHOUDHURY, R. R., AND LEE, S.-J. . J. Cooperative packet recovery in enterprise WLANs. In *INFOCOM, 2013 Proceedings IEEE* (2013), pp. 1348–1356.

[17] GUMMADI, R., WETHERALL, D., GREENSTEIN, B., AND SESHAN, S. Understanding and mitigating the impact of RF interference on 802.11 networks. *ACM SIGCOMM CCR 37*, 4 (2007), 385–396.

[18] HOLLOWAY, C. L., COTTON, M. G., AND MCKENNA, P. A model for predicting the power delay profile characteristics inside a room. *Vehicular Technology, IEEE Transactions on 48*, 4 (1999), 1110–1120.

[19] HUANG, J., XING, G., NIU, J., AND LIN, S. Coderepair: Phy-layer partial packet recovery without the pain. In *Proceedings of (INFOCOM)* (2015), pp. 1463–1471.

[20] JAMIESON, K., AND BALAKRISHNAN, H. PPR: Partial packet recovery for wireless networks. *ACM SIGCOMM Computer Communication Review 37*, 4 (2007), 409–420.

[21] JONES, M. C., MARRON, J. S., AND SHEATHER, S. J. A brief survey of bandwidth selection for density estimation. *Journal of the American Statistical Association 91*, 433 (1996), 401–407.

[22] KATTI, S., KATABI, D., BALAKRISHNAN, H., AND MEDARD, M. Symbol-level network coding for wireless mesh networks. In *ACM SIGCOMM Computer Communication Review* (2008), vol. 38, pp. 401–412.

[23] KONG, L., AND LIU, X. mzig: Enabling multi-packet reception in zigbee. 552–565.

[24] KRYSZKIEWICZ, P., BOGUCKA, H., AND WYGLINSKI, A. M. Protection of primary users in dynamically varying radio environment: practical solutions and challenges. *EURASIP Journal on Wireless Communications and Networking*, 1 (2012), 1–20.

[25] KUMAR, S., CIFUENTES, D., GOLLAKOTA, S., AND KATABI, D. Bringing cross-layer MIMO to today's wireless LANs. In *ACM SIGCOMM CCR* (2013), vol. 43, pp. 387–398.

[26] LEE, D., SEO, H., CLERCKX, B., HARDOUIN, E., MAZZARESE, D., NAGATA, S., AND SAYANA, K. Coordinated multipoint transmission and reception in LTE-advanced: deployment scenarios and operational challenges. *Communications Magazine, IEEE 50*, 2 (2012), 148–155.

[27] LÉOST, Y., ABDI, M., RICHTER, R., AND JESCHKE, M. Interference rejection combining in LTE networks. *Bell Labs Technical Journal 17*, 1 (2012), 25–49.

[28] LI, Y., WANG, X., AND MUJTABA, S. A. Cochannel interference avoidance algorithm in 802.11 wireless LANs. In *Proceedings of VTC* (2003), pp. 2610–2614.

[29] LIM, C.-P. . P., VOLAKIS, J. L., SERTEL, K., KINDT, R. W., AND ANASTASOPOULOS, A. Indoor propagation models based on rigorous methods for site-specific multipath environments. *Antennas and Propagation, IEEE Transactions on 54*, 6 (2006), 1718–1725.

[30] LIN, K. C.-J. . J., KUSHMAN, N., AND KATABI, D. Ziptx: Harnessing partial packets in 802.11 networks. In *Proceedings of the MobiCom* (2008), pp. 351–362.

[31] MAHMOUD, H., AND ARSLAN, H. Sidelobe suppression in OFDM-based spectrum sharing systems using adaptive symbol transition. *Communications Letters, IEEE 12*, 2 (2008), 133–135.

[32] MAJECKA, B. Statistical models of pedestrian behaviour in the forum. *Master's thesis, School of Informatics, University of Edinburgh* (2009).

[33] MIU, A., BALAKRISHNAN, H., AND KOKSAL, C. E. Improving loss resilience with multi-radio diversity in

wireless networks. In *Proceedings of ACM MobiCom* (2005), pp. 16–30.

[34] NACHTIGALL, J., ZUBOW, A., AND REDLICH, J.-P. . P. The impact of adjacent channel interference in multi-radio systems using IEEE 802.11. In *Proceedings of IWCMC* (2008), pp. 874–881.

[35] PAGADARAI, S., RAJBANSHI, R., WYGLINSKI, A. M., AND MINDEN, G. J. Sidelobe suppression for OFDM-based cognitive radios using constellation expansion. In *IEEE WCNC* (2008), IEEE, pp. 888–893.

[36] POLLET, T., VAN BLADEL, M., AND MOENECLAEY, M. Ber sensitivity of OFDM systems to carrier frequency offset and wiener phase noise. *Communications, IEEE Transactions on 43*, 2/3/4 (1995), 191–193.

[37] RAMASUBRAMANIAN, K., AND BAUM, K. An OFDM timing recovery scheme with inherent delay-spread estimation. In *Proceedings of IEEE GLOBECOM* (2001), vol. 5, pp. 3111–3115.

[38] RATHINAKUMAR, S. M., RADUNOVIC, B., AND MARINA, M. K. ShiftFFT: An efficient approach to mitigate adjacent channel interference in OFDM systems. In *Proceedings of the HotWireless in Conjunction with MobiCom* (2015), pp. 11–15.

[39] SAQUIB, N., HOSSAIN, E., LE, L. B., AND KIM, D. I. Interference management in OFDMA femtocell networks: issues and approaches. *Wireless Communications, IEEE 19*, 3 (2012).

[40] SCOTT, D. W. *Multivariate density estimation: theory, practice, and visualization*. John Wiley & Sons, 2015.

[41] SELIM, A., MACALUSO, I., AND DOYLE, L. Efficient sidelobe suppression for OFDM systems using advanced cancellation carriers. In *Proceedings of IEEE ICC* (2013), pp. 4687–4692.

[42] SHEATHER, S. J. Density estimation. *Statistical Science 19*, 4 (2004), 588–597.

[43] SHEU, C.-R. . R., AND HUANG, C.-C. . C. A novel guard interval based ISI-free sampling region detection method for OFDM systems. In *Proceedings of IEEE VTC 2004*, vol. 1, pp. 515–519.

[44] STEWART, G., GOWDA, M., MAINLAND, G., RADUNOVIC, B., VYTINIOTIS, D., AND PATTERSON, D. Ziria: language for rapid prototyping of wireless PHY. In *Proceedings of the 2014 ACM conference on SIGCOMM* (2014).

[45] STOTT, J. The effects of phase noise in COFDM. *EBU technical Review* (1998), 12–25.

[46] TAN, K., FANG, J., ZHANG, Y., CHEN, S., SHI, L., ZHANG, J., AND ZHANG, Y. Fine-grained channel access in wireless lan. *ACM SIGCOMM CCR 41*, 4 (2011), 147–158.

[47] TERRELL, G. R., AND SCOTT, D. W. Variable kernel density estimation. *The Annals of Statistics* (1992), 1236–1265.

[48] THORPE, C., AND MURPHY, L. A survey of adaptive carrier sensing mechanisms for IEEE 802.11 wireless networks. *Communications Surveys & Tutorials, IEEE 16*, 3 (2014), 1266–1293.

[49] VILLEGAS, E. G., LOPEZ-AGUILERA, E., VIDAL, R., AND PARADELLS, J. Effect of adjacent-channel interference in IEEE 802.11 WLANs. In *CROWNCOM* (2007), pp. 118–125.

[50] VITERBO, E., AND BOUTROS, J. A universal lattice code decoder for fading channels. *Information Theory, IEEE Transactions on 45*, 5 (1999), 1639–1642.

[51] WEISS, T., HILLENBRAND, J., KROHN, A., AND JONDRAL, F. K. Mutual interference in OFDM-based spectrum pooling systems. *IEEE Xplore 4* (5 2004), 1873–1877 Vol.4.

[52] WOO, G. R., KHERADPOUR, P., SHEN, D., AND KATABI, D. Beyond the bits: cooperative packet recovery using physical layer information. In *Proceedings ACM MobiCom* (2007), ACM, pp. 147–158.

[53] WOO, G. R., KHERADPOUR, P., SHEN, D., AND KATABI, D. Beyond the bits: cooperative packet recovery using physical layer information. In *Proceedings of ACM MobiCom* (2007), pp. 147–158.

[54] WU, S., AND BAR-NESS, Y. OFDM systems in the presence of phase noise: consequences and solutions. *Communications, IEEE Transactions on 52*, 11 (2004), 1988–1996.

[55] WYSOCKI, T. A., AND ZEPERNICK, H.-J. . J. Characterization of the indoor radio propagation channel at 2.4 GHz. *Journal of telecommunications and information technology* (2000), 84–90.

[56] YAN, Y., YANG, P., LI, X.-Y. . Y., ZHANG, Y., LU, J., YOU, L., WANG, J., HAN, J., AND XIONG, Y. Wizbee: Wise ZigBee coexistence via interference cancellation with single antenna. *Mobile Computing, IEEE Transactions on 14*, 12 (2015), 2590–2603.

[57] ZHANG, C., CHEN, X., AND LUO, Y. Threshold optimization for ISI-free region detection in high-mobility fading channels. In *Proceedings of IEEE HMWC* (2015).

[58] ZHOU, K., JIA, X., XIE, L., CHANG, Y., AND TANG, X. Channel assignment for WLAN by considering overlapping channels in SINR interference model. In *Proceedings of ICNC* (2012), IEEE, pp. 1005–1009.

[59] ZUBOW, A., AND SOMBRUTZKI, R. Adjacent channel interference in IEEE 802.11n. *WCNC 2012* (4 2012).