

CSpy: Finding the Best Quality Channel without Probing

Souvik Sen
HP Labs

Božidar Radunović
Microsoft Research

Jeongkeun Lee
HP Labs

Kyu-Han Kim
HP Labs

ABSTRACT

Wireless performance depends directly on the quality of the channel. A wireless transmitter can improve its performance by estimating and transmitting on only the strongest channel, which can be of significantly higher quality than a weak channel (yielding up to 100% rate improvement). It is considered impossible to predict the quality of the unseen channels. Thus, the only way to identify the strongest channel is by probing each channel individually, incurring large overheads. The key contribution of this paper is a discovery of previously unobserved properties of the wireless channel that makes it possible to predict the the strongest of a set of channels from the measurements collected *only on a single channel*. We confirm the properties through measurements and present a theoretical analysis that explains their nature. Our proposed system, *CSpy*, utilizes these observations to predict the strongest channel. *CSpy* is the first to reliably estimate the strongest channel by utilizing channel responses extracted from off-the-shelf wireless chipsets, without probing any additional channels. By tracking the strongest channel, *CSpy* improves performance by up to 100% in comparison to channel agnostic schemes.

Categories and Subject Descriptors

C.2.1 [Network Architecture and Design]: Wireless communication

General Terms

Algorithm, Design, Experimentation, Performance

Keywords

Wireless, Cross-Layer, Channel Estimation

1. INTRODUCTION

There has been extensive research towards improving wireless performance using both link layer (scheduling, coordina-

tion, rate-control) and physical layer (MIMO, coding, beamforming) innovations. Most of these techniques try to achieve the highest performance that can be obtained under a given channel condition. However, wireless performance is ultimately limited by the quality of the channel that the link is operating on. The channel qualities vary significantly across the frequency band [1–3]. E.g., a few of the 21 available channels in the 5GHz band may be strong, and the others weak. TV whitespaces can provide more opportunities, with up to 40 available channels [4]. Due to environmental changes and user mobility, the qualities of these channels vary over time and *across channels*. We find that the quality of the strongest channel can be significantly better than the weakest, or even an average channel. Thus, it is possible to improve the wireless performance if we can successfully track and communicate on the strongest channel at all times.

Wireless signals traverse multiple paths before arriving at the receiver. The signals traversing different paths undergo different attenuations, delays and phase shifts. The phase shifts are further affected by the carrier frequency. The channel qualities at different frequencies thus depend on how different complex multipath signal components combine at the receiver. Due to the phase shift induced by the carrier frequency, signals from some paths that add constructively at one frequency may combine destructively at another frequency. Because of a large number of paths in a multi-path channel and a limited resolution of the receivers, it is considered impossible to predict the quality of any channels other than the ones being observed. This is why all of the existing channel estimation algorithms (e.g. [1, 5, 6]) use channel probing to identify the *Strongest Channel Index (SCI)*¹. But probing is expensive. Moreover, identifying the strongest channel is not a one time operation and tracking the strongest channel will require frequent channel probes, incurring a very high overhead.

In this paper we present a novel technique that allows us to *determine the strongest channel from a set of channels while probing only a single channel from the set*. We first present extensive channel measurements showing that links at different locations with the same strongest channel index exhibit similarities in their multi-path structures. This is surprising, given that channels at different locations are typically assumed independent [7, 8]. We then analyse the observed correlation under a well-known Saleh-Valenzuela (SV) channel model [9]

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
MobiCom'13, September 30–October 4, Miami, FL, USA.
Copyright 2013 ACM 978-1-4503-1999-7/13/09 ...\$15.00.

¹We define strongest channel index as the channel which yields the highest expected throughput across all available channels.

and show how to construct a simple estimator that can provably determine on average the strongest of the two adjacent channels. This analysis also helps us identify the intrinsic characteristics of the channel that make the prediction possible. However, even constructing such a simple estimator of SCI was not obvious. We thus turn to machine learning techniques to help us automatically learn the best estimator using a training set of observed channel. We find that the performance of the machine learning estimator improves as it is trained with diverse set of environments, e.g., indoors vs. outdoors. However, it is possible to infer the strongest channel with up to 80% accuracy using our approach, without probing any additional channel. Further, by using the strongest channel, wireless performance can be improved by up to 100%.

Translating the above key ideas to a real system entails a host of technical challenges. First, how are various channel profiles in practice, in different real-world environments, amenable for prediction? Second, is it possible to determine the strongest channel even if it is far away from the current operating channel? Third, how do we deal with mobile scenarios where the channel can change frequently? In such mobile scenarios, even if it is possible to find the SCI on a per-packet basis, tracking the strongest channel may incur prohibitive channel switching costs.

To address the above questions, we present a system called *CSpy*. To the best of our knowledge, *CSpy* is the first system that is able to predict the best channel without probing. We show that *CSpy* performs well under realistic channel conditions, providing significant performance gains over channel agnostic approaches and no channel probing overhead. Moreover, we leverage physical layer information to estimate the coherence time of the channel, and trigger channel switching only if estimated SCI can be used for reasonable duration. Also, note that inferring the strongest channel index is orthogonal to detecting the least used channel [10–12]. Ideally, both should be combined (as in e.g. [13]), but this is out of scope of our current work. The main contributions of this paper are as follows:

- **A measurement study to understand the channel quality across the 5GHz frequency band in a typical office environment.** We find that the achievable throughput gain by tracking the strongest channel, compared to fixed channel strategies, can be substantial (up to 100%). The channels with the same SCI resemble to each other while they differ from channels with different SCIs.
- **A theoretical analysis on the effect of multi-path on channel inferring capability.** We prove that even with limited multipath information from a single 20MHz channel, it is possible to determine the strongest amongst the adjacent channels in a typical indoor environment.
- **A learning based approach to determine the strongest channel by utilizing the channel responses extracted from off-the-shelf WiFi chipsets.** We extract the multipath information from commodity chipsets, and utilize it to determine the strongest channel index.
- **A demonstration of performance gains.** We implement *CSpy* algorithms in our test-bed, and quantify our gains for various scenarios. We show that we can predict the SCI

among 5 adjacent channels on each side achieving 85% throughput of the optimal “oracle” algorithm, and yielding around 50% gain in rates over an average channel, and up to 2.5x gain over the worst channel.

The subsequent sections expand on each of these contributions, including controlled measurements, followed by analysis, design, and evaluation.

2. MULTIPATH CHANNEL RESPONSE

We start with a quick primer on channel response, how it affects the link quality and why it is difficult to predict it.

Channel impulse response (CIR): The wireless signal from the transmitter to the receiver traverses through multiple paths, undergoing reflections, diffractions, and scattering. Essentially, the received signal contains multiple time-delayed, attenuated, and phase-shifted copies of the original signal. If $x(t)$ is the transmitted signal at time t , and $h(t, \tau)$ captures the *Channel Impulse Response (CIR)* at time t to an impulse transmitted at time $t - \tau$, the received signal is [7]:

$$y(t) = \int_{-\infty}^{\infty} h(\tau)x(t - \tau)d\tau + w(t), \quad (1)$$

where $w(t)$ is additive white noise. The channel impulse response h can be considered time-invariant during the packet duration (thus we drop the dependency on t), and is defined as:

$$h(\tau) = \sum_{p=0}^{P-1} A(p)\delta(\tau - \tau(p)), \quad (2)$$

where $A(p) = a(p)e^{i\phi(p)}$ is a complex response of path p , P is the (very large) number of paths between the transmitter and the receiver and $a(p)$, $\phi(p)$, $\tau(p)$ are the attenuation, phase and delay of the signal traversing on path p . The Fourier transform $H(f) = \mathcal{F}(h(t))$ of CIR is often called the *Channel Frequency Response (CFR)*. An equivalent of (1) in frequency domain is

$$Y(f) = X(f)H(f), \quad (3)$$

where $Y(f) = \mathcal{F}(y(t))$ and $X(f) = \mathcal{F}(x(t))$ are Fourier transform of the received and transmitted signal, respectively.

Extracting channel information from commodity chipsets: Most modern digital radios use OFDM communication, and transmit signals across orthogonal subcarriers at different frequencies. An OFDM receiver implementation includes a channel estimation logic in the hardware as a part of basic operations. Some WiFi chipsets, such as Atheros 9390, can export the estimated discrete CFR to the driver. Let f_c be the carrier frequency of channel c , W be the bandwidth of the radio and n be the number of subcarriers. In most 802.11 a/g/n receivers, we have $W = 20$ MHz and $n = 64$. The reported discrete CFR $\hat{\mathbf{H}}^c = [H(f_c - W/2), H(f_c - W/2 + 1), \dots, H(f_c - W/2 + n - 1)]$ is a complex vector that describes the quality of channel c at each subcarrier. Figure 1 shows examples of a few discrete CFR vectors.

The channel impulse response $\hat{\mathbf{h}}^c = \mathcal{D}^{-1}\hat{\mathbf{H}}^c$ can be obtained by applying an inverse (fast) discrete Fourier transform on the CFR. By applying the aforementioned transforms we get

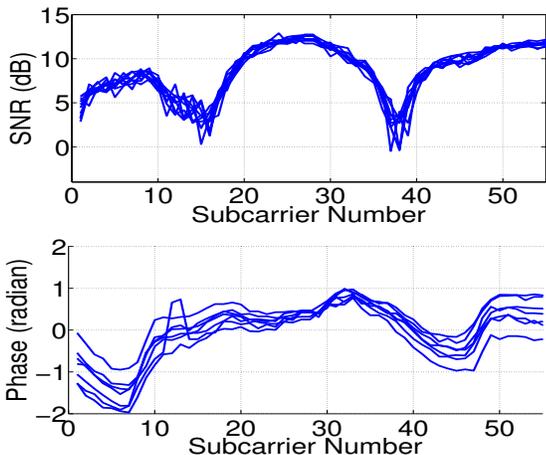


Figure 1: Magnitude (top) and phase (below) of CFR reported for 10 consecutive packets.

a discrete CIR $\hat{\mathbf{h}}^c$ with n number of taps:

$$\hat{\mathbf{h}}^c = [\hat{h}^c[0], \dots, \hat{h}^c[n]] \text{ where}$$

$$\hat{h}^c[l] = \frac{1}{P} \sum_{p=0}^{P-1} A(p) e^{-2\pi(f_c - W/2)\tau(p)} \frac{1 - e^{-i2\pi W\tau(p)}}{1 - e^{i2\pi(l/n - W\tau(p)/n)}. \quad (4)$$

Intuitively, $\hat{h}^c[l]$ approximately aggregates the effects of paths p with $\tau(p) \in [l/W, (l+1)/W]$. Figure 2 shows an example of a discrete CIR vector. Note from (4) that $\hat{\mathbf{h}}^c$ changes with the carrier frequency f_c , since it is a sum of complex vectors whose phase shift depends on the carrier frequency.

Difficulty of predicting the quality of unseen channels:

Our goal is to predict a performance of an unseen channel c' using an observed discrete CIR $\hat{\mathbf{h}}^c$ from a different channel $c \neq c'$. We will now argue that it is impossible to do so without assuming some structure on the amplitude $a(p)$ and phase $\phi(p)$ of each individual path. To see that, consider the system of equations (4) jointly for channels c and c' . This system has $2n$ equations and P unknown complex variables $A(p)$. Because the number of paths P is very large (c.f. [9]), much larger than the number of subcarriers n , this system is underdetermined and has infinitely many solutions. Thus, there exists a set of values for $\{a(p), \phi(p)\}_{p=0, \dots, P-1}$ that would achieve any discrete CIRs $\hat{\mathbf{h}}^c$ and $\hat{\mathbf{h}}^{c'}$ in channels c and c' . Without some assumption on the channel we are not able to infer anything about channel c' only by observing channel c . In reality, the channel $\{a(p), \phi(p)\}_{p=0, \dots, P-1}$ has some structure which we will exploit. In the subsequent section we will first show that we can indeed make some prediction about unseen channels and in Section 4 we will explain why it works.

Metric for determining channel quality: Channel quality can be measured by the per-packet SNR metric that 802.11 drivers usually provide in the form of Received (or Relative) Signal Strength Indicator (RSSI). However RSSI is known to be an unreliable indicator of channel quality [14–16]. In search of a better metric, we find that the effective SNR (eSNR) value, as described in [14] can correctly capture the quality of the channel from the CFR. eSNR first uses the CFR to compute the SNR and bit-error rate (BER) of each subcarrier. From the subcarrier-specific BER, a channel-wide BER is computed and translated into eSNR. eSNR can further be used to determine the highest bit-rate that the channel can support. As shown

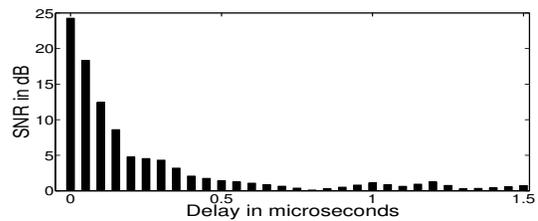


Figure 2: Amplitude of a part of a CIR vector obtained from a CFR after applying IFFT.

in [14], eSNR can estimate the correct bit-rate of the channel with great accuracy, and performs significantly better than the per-packet SNR metric. Out of all the available channels in the 5GHz frequency band, we define the strongest channel index (SCI) as the channel that supports the highest bit-rate. We also adjust for link layer overheads (using [17]) to calculate the *Expected Throughput* of a channel from the bit-rate prescribed by the eSNR metric. In the rest of our paper we use the eSNR-based expected throughput metric to compare the throughput of the different channels from their CFR vectors.

3. MEASUREMENT AND HYPOTHESES

This section aims to show that PHY layer CIR information can be an indicator of the strongest channel. We also demonstrate that tracking the strongest channel can indeed improve the network performance. We present three main hypotheses if CIRs from a single channel can be used to infer the strongest channel:

1. *The quality of the channel varies across frequency.* Wireless performance can be significantly improved by identifying and tracking the strongest channel, in comparison to a random fixed channel.
2. *Links with the same strongest channel shares some statistical similarity in their multipath structure.* It is possible to identify some features from the CIR information that can indicate the SCI.
3. *Links with different strongest channel index exhibits different multipath characteristics.*

To verify our hypotheses, we first describe our experimental methodology, followed by the key findings.

3.1 Measurement Setup

Our initial experiments were conducted at 10 locations in an office building. We use laptops equipped with Atheros 9390 WiFi card operating on the 5GHz frequency spectrum as transmitters and receivers. We instrument the Atheros 9390 driver to extract the CFR information on a per-packet basis. The Atheros firmware exports a subset of the CFR (56 out of 64 subcarriers), and we only use these to compute the CIR. Although there are 21 available channels in the 5GHz frequency band, the Atheros 9390 driver is not Dynamic Frequency Selection (DFS) enabled [18] and hence can operate on only 9 of the 21 available channels. Hence, we perform experiments on channel indices 36 – 48 and 149 – 165 only.

We design experiments to measure the quality of all the 9 available channels at 5GHz. We make necessary driver modifications at both the transmitter and the receiver to cycle

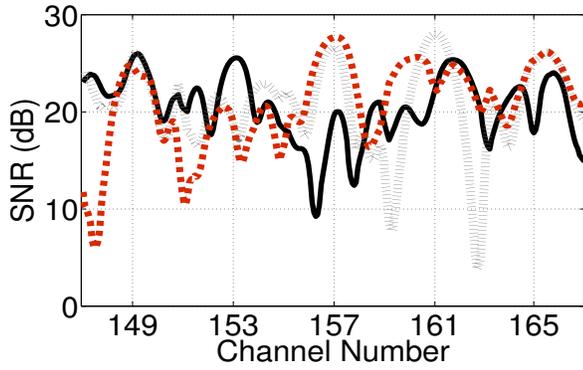


Figure 3: Magnitude of CFR vectors collected from channels 149 – 165.

through the channels. The transmitter starts from channel 36 at 5GHz and transmits 10 short packets, before switching onto the next channel. To maintain synchronization, the receiver also follows the same predefined sequence. Due to software restrictions, a single round takes considerable amount of time – approximately 70ms. To ascertain no arbitrary multipath changes within a single round, we consider only those measurements where the channel characteristics were highly correlated across two consecutive rounds. We use the measurements to verify our key hypotheses in the following subsection.

3.2 Measurement and Verification

Hypothesis 1: *The quality of the channel varies across frequency.*

The CFR vectors, collected from a single round of transmissions, provide a snapshot of the characteristics of the channels across all the available frequencies in the 5GHz band. Figure 3 shows 3 such snapshots of 5 different channels. We observe that the CFRs across different frequencies are quite dissimilar. To understand how the wireless performance may vary across different channels, we compute the expected throughput from the CFRs. Figure 4 plots the difference between the expected throughput of the strongest channel over the weakest and a random fixed channel. Evidently, the performance gap is considerably large. On average, the expected throughput of the strongest channel is almost 2 – 3 times than that of the weakest channel. Thus, we conclude that *the channel quality varies significantly across frequency. If we can identify the strongest channel, wireless performance will improve significantly.*

Impact of temporal variations: Channel quality can change over time due to environmental changes. Figure 5(a) illustrates the performance of a single link across 5 different channels over a duration of 50 seconds. A single channel demonstrates gradual but significant throughput variations due to the slow fading caused by environmental variations. We find that the performance also varies widely across channels. The quality of a channel depends on how the different multipath components combine at its operating frequency. As the environment changes, the multipath structure of the signal also changes. The multipath components may combine constructively at one frequency and destructively on another, while over time, the multipath structure can change in such a way that the previously stronger channel becomes weaker and vice versa. Thus, we find that the strongest channel index also

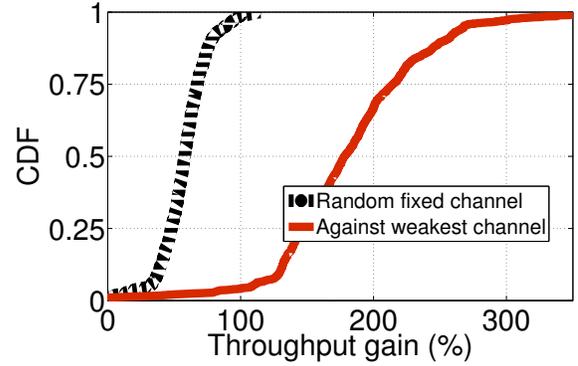


Figure 4: Distribution of expected throughput gain by using the strongest channel.

changes over time (figure 5(b)). *It is not enough to identify the strongest channel only once, we need to track it to maximize performance.*

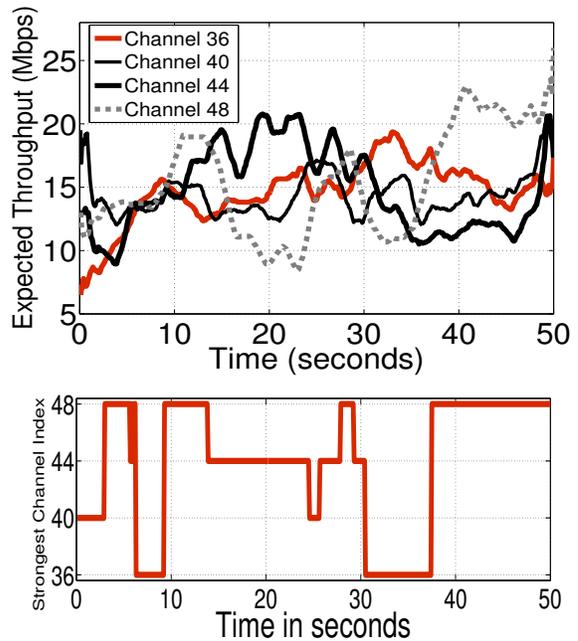


Figure 5: (a) The expected throughput of a particular link for different channels. (b) SCI changes over time.

Hypotheses 2: *Links with the same strongest channel shares some statistical similarity in their multipath structure.*

The channel quality depends on the multipath characteristics observed at a location. If at some instance, two links have similar multipath characteristics, they will also probably have the same strongest channel. To verify this hypotheses, let group $G_{i,j}$ contain all the CIRs collected from channel number i when the SCI was j . In other words, each group contains CIRs from the same channel, but collected at different times and locations, such that the SCI was the same at those instances. According to our hypotheses the CIRs belonging to the same group should be similar to each other. We find the similarity between two CIRs, a and b , by using the the cross correlation metric:

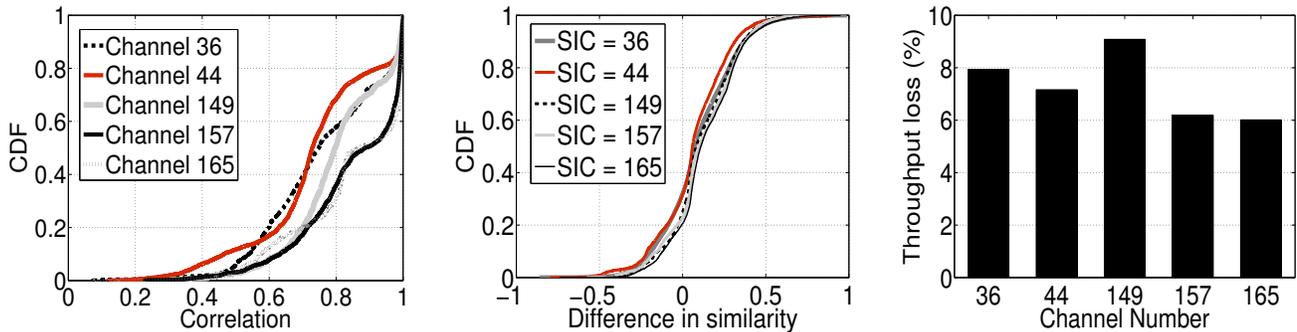


Figure 6: (a) Similarity between the CIRs gathered from channel 36 at different locations when the SCI was the same. (b) CDF of difference in similarities $C_{own} - C_{others}$, observed at different locations, for different SCI values. (c) Throughput loss due to mispredicting the SCI.

$$c(\mathbf{a}, \mathbf{b}) = \frac{\sum_s a_s b_s}{\sqrt{\sum_s a_s^2} \sqrt{\sum_s b_s^2}}. \quad (5)$$

Figure 6(a) shows that CIRs from the same group have high correlation, suggesting that the same SCI may imply similar multipath structure. If we can a priori learn some unique features about the multipath, we may be able to infer the SCI by observing the CIR from a single channel. Our observation in figure 6(a) may appear coincidental. However we will explain in section 4 why it may be possible to identify such differentiating features about the multipath at all.

Hypotheses 3: *Links with different strongest channel index exhibits different multipath characteristics.*

We evaluate the (dis)similarity between two groups of CIRs that were collected from the same channel but had a different SCI. Whenever a CIR from a group $G_{i,j}$, is more similar to the CIRs of a different group, $G_{i,k|k \neq j}$, than those from its own group, we will mispredict the strongest channel. To evaluate the probability of such mispredictions, we arbitrarily choose a test CIR from a group $G_{i,j}$, and use correlation to find the similarity between the test CIR with a randomly chosen CIR from the same group. The correlation value, denoted as C_{own} , indicates *similarity* of a sample test CIR with responses which have the same SCI. Now, for the CIRs from all other groups, $G_{i,k|k \neq j}$, we find the *similarity* of a randomly chosen CIR with the test CIR – we denote this similarity as C_{others} . If a test CIR is more similar to a different group than its own group, we will naturally misclassify its SCI.

Figure 6(b) plots the CDF of the difference in similarities, $C_{own} - C_{others}$, for 5 different strongest channels. If the difference is negative, then the SCI is likely to be misclassified. We find that the CIRs are often sufficient to correctly classify the SCI. However, in some cases more than 25% of the responses are misclassified. To understand these mispredictions, figure 6(c) plots the difference in expected throughput between the mispredicted strongest channel and the correct strongest channel. We find the performance gap is small, implying that there can be multiple strong channels with similar quality. As long as we can infer one of the stronger channels, we will be able to improve the wireless performance.

Impact of resolution on SCI detection: Ideally we want to find out all the multipath signal components that traverse be-

tween the transmitter and the receiver. However, due to limited (20MHz) bandwidth of WiFi, we can only determine a discrete channel impulse response (CIR). From equation (4) in section 2, it is clear that the discrete CIR depends on the operating frequency of the WiFi chipset. Thus, we expect that the structure of the CIR will vary across channels. Adjacent channels may have a relatively similar CIR structure and CIRs from widely separated channels can be quite dissimilar (figure 7). Thus, it is not possible to use the CIRs collected from a single channel in the 5GHz frequency band, to infer the strongest channel in the 2.4GHz band. However, we will show in our evaluation that it is possible to predict the strongest channel, amongst atleast a few adjacent channels with reasonably high accuracy.

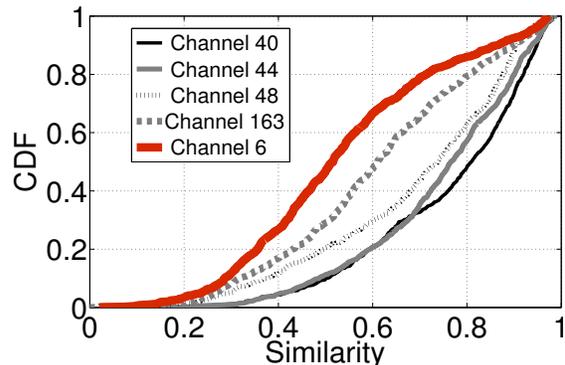


Figure 7: Similarity of aggregate CIRs collected from channel 36 with other channels. The similarity is high for nearby channels (40, 44, 48), but low for channel 163 and even lower for channel 6 which is in 2.4GHz.

4. WHY DOES PREDICTION WORK?

Measurement results from the previous section clearly show that we can predict the best among the adjacent channels based on the observations from the current channel. However, as discussed in Section 2, this is surprising, and has to follow from a certain structure of the channel response. The goal of this section is to explain the intuition of why are we able to make any prediction at all. To this end, we construct a simplified channel model and a simple predictor based on the observations in the current channel and prove that it is able to predict the best among the two adjacent channels. Note that the actual prediction algorithm we use in our implementation is different, and described in the following section. The actual

algorithm is more general but is not amenable to a simple theoretical analysis. In this section we are deliberately focusing on a simple but analytically tractable scheme, to illustrate the intuition of why the prediction is possible at all.

Firstly, we define a predictor. Let

$$\Delta^c = \tan(\angle \hat{h}^c[1]) - \tan(\angle \hat{h}^c[0])$$

be the phase difference between the first and the second tap of the CIR of channel c (where $\angle x$ denotes the phase of a complex number x). Note that this is purely based on the observations we make in channel c . We claim that when $\Delta^c < 0$ in channel c we expect that channel $c+1$ is better than channel $c-1$. Secondly, we defined a simplified channel model. We start with a well known Saleh-Valenzuela (SV) channel model [9]. The key observation from [9] we use is that the path phases $\phi(p)$ are random (i.e. $\phi(p)$ and $\phi(q)$ are independent when $p \neq q$) and uniformly distributed on $[0, 2\pi]$. Also, let T_d be the delay spread of the channel such that $a(p) = 0$ for $\tau(p) > T_d$ (and $a(p)$ can be an arbitrary value for p such that $\tau(p) \leq T_d$). This is also in accordance with the SV model. We then have the following proposition

PROPOSITION 1. *Let us consider a $W = 20\text{MHz}$ bandwidth radio and a channel with the delay spread $T_d = 20\text{ns}$. Under the assumptions on the channel model above, we have*

$$\mathbb{E} \left[\Delta^c \left(\left| \hat{h}^{c+1}(0) \right|^2 - \left| \hat{h}^{c-1}(0) \right|^2 \right) \right] < 0 \quad (6)$$

PROOF. Due to a lack of space we only give a sketch of the proof and ignore little- o terms. First, let $\hat{A}(p) = a(p)e^{i\hat{\phi}(p)} = A(p)e^{-2\pi(jc - W/2)\tau(p)}$. It is easy to see that the phase $\hat{\phi}(p)$ is also random, uniformly distributed in $[0, 2\pi]$ and independent across paths. Using Taylor expansions (and omitting little- o terms) we then obtain the following approximation

$$\begin{aligned} \Delta^c &= \tan(\angle \hat{h}^c[1]) - \tan(\angle \hat{h}^c[0]) \approx \left. \frac{\partial \tan(\angle \hat{h}^c[l])}{\partial l} \right|_{l=0} \\ &\approx \sum_{p,q} a(p)a(q) \sin \hat{\phi}(p) \cos \hat{\phi}(q) \left(\frac{n}{W\tau(p)} - \frac{n}{W\tau(q)} \right). \end{aligned}$$

Also, we have

$$\begin{aligned} \left| \hat{h}^{c+1}(0) \right|^2 &= \\ &= \frac{1}{n^2} \sum_{p,q} a(p)a(q) \cos(\hat{\phi}(p) - \hat{\phi}(q) - 2\pi W(\tau(p) - \tau(q))). \end{aligned}$$

We are interested in $\mathbb{E}[\Delta^c \left| \hat{h}^{c+1}(0) \right|^2]$. Note that $\mathbb{E}[\sin(\hat{\phi}(p))] = \mathbb{E}[\cos(\hat{\phi}(p))] = 0$. Since different paths have independent phases, the expectation of all cross-products containing different phases will be zero (e.g. $\mathbb{E}[\sin(\hat{\phi}(p)) \cos(\hat{\phi}(q))] = 0$ for $p \neq q$). Taking this into account when expanding the terms in $\mathbb{E}[\Delta^c \left| \hat{h}^{c+1}(0) \right|^2]$, and with some elementary manipulations, we get

$$\begin{aligned} \mathbb{E}[\Delta^c \left| \hat{h}^{c+1}(0) \right|^2] &\approx \frac{1}{n^2} \sum_{p=0}^{P-1} \sum_{q=0}^{p-1} 2a(p)a(q) \\ &\times \sin(2\pi W(\tau(p) - \tau(q))) \left(\frac{n}{W\tau(p)} - \frac{n}{W\tau(q)} \right). \end{aligned}$$

Let us order the paths such that $\tau(p) > \tau(q)$ for $p > q$. We then have in every term that $\tau(p) - \tau(q) > 0$ and $W(\tau(p) - \tau(q)) \leq WT_d < \frac{1}{2}$, thus each sine term is positive and each term in the sum is negative. Consequently, we have

$\mathbb{E}[\Delta^c \left| \hat{h}^{c+1}(0) \right|^2] < 0$. Similarly we get $\mathbb{E}[\Delta^c \left| \hat{h}^{c-1}(0) \right|^2] > 0$ for channel $c-1$. \square

The interpretation of the proposition is as follows. Measurements show that most of the channel energy is concentrated in the first tap $\hat{h}^c(0)$. Term $\left| \hat{h}^{c+1}(0) \right|^2 - \left| \hat{h}^{c-1}(0) \right|^2$ is thus proportional to the difference in qualities between channels $c+1$ and $c-1$. If this term and Δ^c were uncorrelated, then the expectation of the product would be zero. Since the expectation is negative, this implies that negative Δ^c yields positive difference between channel qualities of $c+1$ and $c-1$, on average. Thus if one observes negative Δ^c , it can conclude that channel $c+1$ is better (on average) than channel $c-1$ and conversely for positive Δ^c channel $c-1$ is most likely better than channel $c+1$.

We further verify the finding of Proposition 1 on the measured data. In Figure 8(a) we plot the empirical PDF functions of the simple estimator (6) over all measured channel responses. We verify that the observation from Proposition 1 holds, and that the mean is negative. Moreover, we see that the simple estimator has 70% accuracy (figure 8(b)) in practice when predicting the strongest of the adjacent channels.

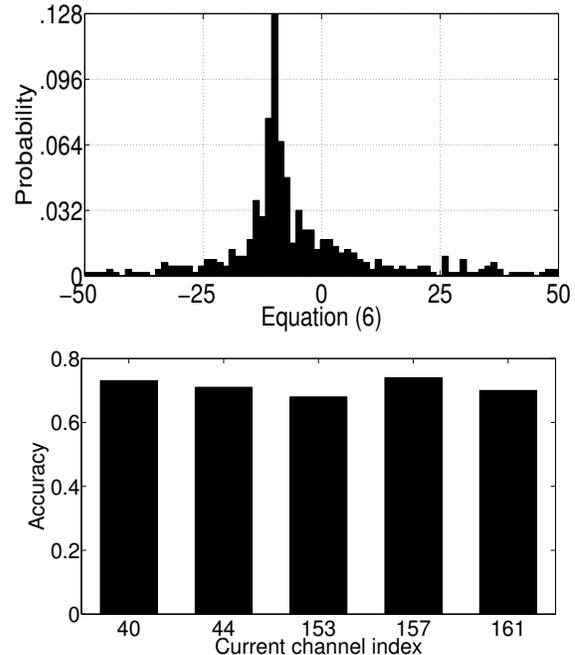


Figure 8: (a) The empirical PDF of the estimator (6) calculated over the entire set of measured CIR samples. (b) Accuracy of strongest adjacent channel prediction for different current channels.

The main high-level intuition behind the result is that the prediction works because the phases of different paths are uniformly random and independent from each other. Thus if we carefully construct an estimator based on relative phases, as

it was shown in this example, we can recover a bias that is correlated with the channel quality. But constructing such an estimator is a difficult task. In the subsequent section we will use machine learning techniques to systematically explore the space of estimators and construct the best one for our problem, which improves the accuracy and expands the prediction range.

5. DESIGN AND IMPLEMENTATION

In this section we present the design and implementation details of CSpy. We have seen in Section 4 a simple predictor that is able to predict the strongest of two adjacent channels. However, this simple predictor captures only one specific of the underlying channel structure. There might be others that can further improve the accuracy and expand the prediction power to more distant channels, but as we have seen, they are not easy to discover. Instead, we use machine learning techniques to train a predictor on a large training set and automatically discover various hidden correlations. We first describe the feature representation and selection process, and then define the classifier. We further describe another important building block that estimates the coherence time of the channel, and finally we present an overview of the entire system design.

5.1 Feature representation and selection

The analysis from Section 4 tells us that the relative phase differences between the first two taps can provide some information on SCI. Thus we decided to generalize this observation and use a feature vector that encodes relative phases of all taps. To obtain such a vector we take a complex discrete CIR vector observed by the NIC and we normalize its amplitudes and phases: we divide the discrete CIR vector by its amplitude and we subtract the phase of the first tap from the phases of all taps. Such a normalized feature vector then effectively encodes the descriptions of phases of different taps relative to the phase of the first one.

We obtain the discrete CIR vector by applying the inverse FFT on a discrete CFR vector obtained from a NIC, as described in Section 2. We extract the discrete CFR estimated for each received packets from Atheros 9390 chipsets with minor driver modifications. We note that the 802.11n standard allows higher resolution CFR reports, than the Atheros 9390 chipsets, that can further improve the performance of our system [19]. The extracted CFR measurements contains arbitrary phase offsets due to the lack of synchronization between the transmitter and the receiver.² We use the techniques described in [8, 20] to compute the phase offset and adjust the phase of each CFR before utilizing them for further calculations.

Finally, we note that the l -th tap of the discrete CIR vector $\bar{\mathbf{h}}$ predominantly contains the aggregate of paths with delays $\tau(p) \in [l/W, (l+1)/W]$, as discussed in Section 2. We observe that most of the significant multipath components appear in the first few taps of $\bar{\mathbf{h}}$. This is because the major multipath signals from the transmitter arrive at the receiver within a short time period called the *delay spread* [21], which is around 50ns in indoor environments. In our evaluation, we will show that

²This is not a problem for a conventional OFDM receiver that only needs to remove, but not learn the channel response.

the first 10 taps, which carries most of the energy, of the CIR vector $\bar{\mathbf{h}}$ can be a reliable indicator of the SCI. Using more taps will only increase the classification complexity, without any noticeable improvement in performance.

5.2 Training and Classification

Another takeaway from Section 4 is that the predictor (6) works well under a fairly general channel model assumption. This suggests that a simple classifier should be effective on a large class of channel profiles. Therefore, we decided to use a multi-class support-vector machine (SVM) as a simple linear classifier. In the training phase we collect CIR samples from multiple channels at various locations. For each channel we train our classifier on the training set comprising $\{\bar{\mathbf{h}}^c(x), \text{SCI}(x)\}_x$ over all locations x . Once a new CIR is obtained in the wild, we apply the classifier on it, derive the SCI and switch to the best channel. Since SVM is a linear classifier, it is simple and fast to calculate.

It is important to note that our classifier does not need to be trained on samples from each individual location. It is not built to learn all possible channel responses. It is designed to learn correlations that exist under very general assumptions, like the one illustrated in Section 4. In practical terms this means that we can train the classifier in one building and may be able to use it in another, as we demonstrate in our evaluation in Section 6. We also emphasize that the performance of the classifier does not reflect directly in the classification accuracy. Often, even if the classifier misestimates the SCI of the current location, the estimated channel can still have good performance. We evaluate this aspect in detail our evaluation.

5.3 Tracking the Strongest Channel

Even if it is possible to identify the strongest channel, channel switching may not always be beneficial due to the switching cost. To understand how quickly the SCI changes due to changes in the environment, we collected a few CFR measurements by tuning both the transmitter and the receiver at channel 36 using a 40MHz bandwidth. We disabled channel switching as it introduces an unacceptable overhead (3ms in our implementation) in our measurement. We postprocess each CFR vector over 40MHz to create CFRs of 4 consecutive 10MHz channels, and compute the strongest amongst only these 4 channels. Figure 9 shows that for a static link, we may need switch the channel only every 1–2s and thus the switching cost can be amortized. We observe that although environmental variations affect the channel characteristics, the strongest channel may not always be more affected than other weak channels. However, if the link is mobile, the strongest channel remains consistent only for less than 50ms. Thus, to track the SCI of a mobile link, we may need to switch channels frequently, which may incur some overhead.

Although it is difficult to estimate the exact duration for which the SCI can be used, the same duration is proportional to the coherence time of the channel (T_c). We compute the channel coherence time as the time-lag beyond which the correlation of the CFR vectors is less than 0.6. To amortize the cost of the channel switching time, CSpy ascertains that the strongest channel can be used for a duration significantly more than the channel switching duration (T_s). It prescribes a channel change only when $T_c > kT_s$ where k is a constant. We use $k = 10$ in our implementation to prevent CSpy from hurting

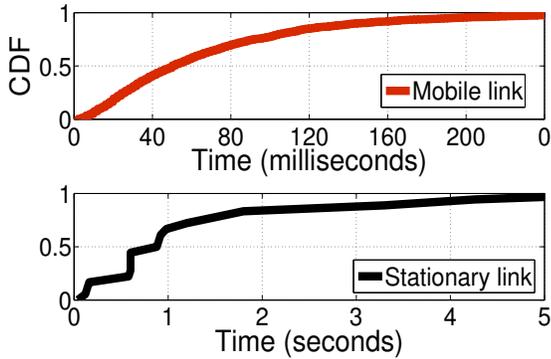


Figure 9: Time duration between two consecutive channel changes while tracking the strongest channel.

performance by triggering frequent channel switching. Observe that the SCI coherence time (50ms under mobility) is almost 5x longer than the channel coherence time (typically 10ms for walking speeds). Moderately fast switching time in recent chipsets (3ms in our implementation) and a reasonable SCI coherence time allows us to tightly track the SCI in most scenarios.

5.4 Overall System Design

Figure 10 presents the overall design of CSpy. CSpy trains a SVM classifier with the CIR feature vectors, obtained from the CFR readings, after appropriate phase and amplitude sanitization. It groups the CIR feature vectors according to the empirically determined SCI. Once the classifier is trained with sufficiently diverse set of CIR readings, it can be used to infer the strongest channel. During run-time, CSpy estimates the CIR feature vector as well as the channel coherence time from the CFR readings obtained from a single channel. It classifies the CIR vector using the previously trained classifier, to obtain the SCI. Finally, CSpy switches to the strongest channel only if the coherence time of the channel is significantly more than the channel switching time.

6. EVALUATION

We evaluate CSpy across 70 different links in two different office buildings, and a parking lot. We conduct experiments to answer 2 key performance questions. (1) What is CSpy’s strongest channel estimation accuracy? (2) How much performance gain can CSpy achieve by tracking the strongest channel? We begin by explaining our evaluation methodology.

6.1 Experiment Design

We employ a trace-driven simulation approach to evaluate CSpy. We use the same measurement setup as described in section 3.1 and collect traces during daytime. At 70 fixed locations, we cycle through 9 different channels and gather the CFRs. For each round, we compute the strongest amongst all the channels by using the expected throughput metric. We also collect 8 walking traces from a single 40MHz channel (channels 36 bonded with channel 40), and thereafter divide the CFR into 4 contiguous 10MHz channels. We use only these 4 channels to evaluate CSpy’s performance under mobility.³ To evaluate CSpy’s performance for a particular link l_i , we

³It is not possible to evaluate all possible channels because the multipath structure changes frequently under mobility.

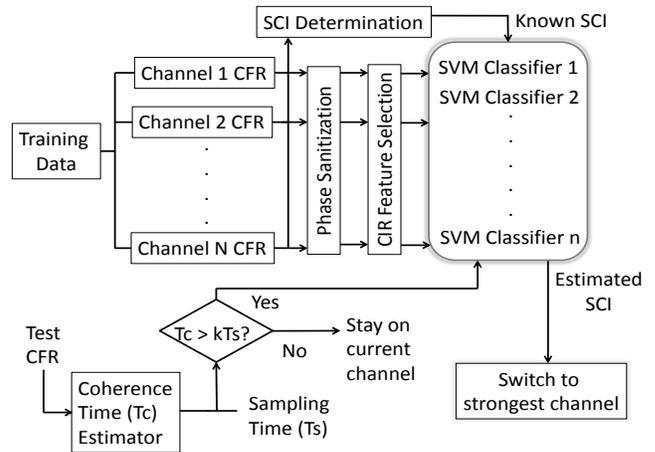


Figure 10: CSpy design.

train the SVM classifier using the CIR features collected from all other links $l_j |_{i \neq j}$. Unless specified, we use the 10 taps specified in Section 5.1 as our feature vector.

Metrics and comparison: We use the following three metrics to evaluate CSpy’s performance: (1) Accuracy – the fraction of cases when the expected throughput of the estimated SCI is within 90% of the actual strongest channel. (2) False positives (FP) – the fraction of cases in which the classifier estimates a weak channel, whose expected throughput is less than the average throughput across all available channels. In other words, false positives account for scenarios when CSpy hurts performance by switching to a weak channel. (3) Throughput – We compute CSpy’s throughput as the expected throughput of the estimated SCI, aggregated over time. We also compute the optimal throughput from the known SCI determined empirically. Apart from the optimal scheme, we compare CSpy with a scheme which always uses a fixed channel and a probing scheme which periodically scans all possible channels to determine the SCI.

6.2 SCI Estimation Performance

Accurate estimation of the SCI will help in improving the wireless throughput by switching to a stronger channel. Otherwise, CSpy may hurt performance by switching to a weak channel. Figure 11(a) shows CSpy’s SCI estimation performance across 50 locations in a single building. Evidently, CSpy’s classification technique can determine the SCI from the observed CIR with a median accuracy of 83%. In rare occasions (mean 2.5%) CSpy underperforms and estimates a weak channel as the SCI. However, we observe poor accuracy (less than 50%) at a few locations. Careful investigation showed that these locations received packets at weak signal strengths, that results into low fidelity of the estimated CIRs [8].

CSpy collects the CIR from the current operating channel and utilizes it to estimate the strongest channel. Figure 11(b) further analyzes CSpy’s performance and shows the accuracy/false-positives on a per operating channel basis. The accuracy varies across different channels because the richness of the CIRs may vary across channels. We find that the channels at the edge of the frequency band, indices 36, 161, 165,

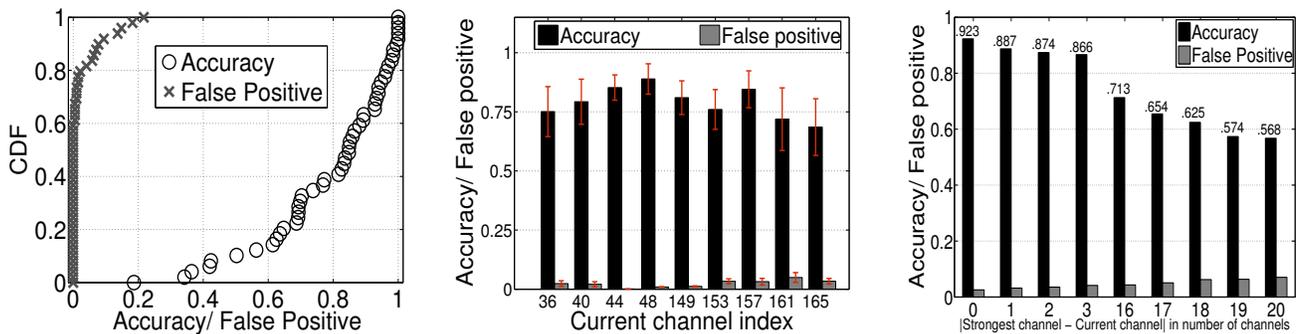


Figure 11: Mean accuracy/false-positive of SCI determination (a) at 50 different locations, (b) for CIRs collected from different operating channels, (c) with increasing frequency gap between the operating channel and the actual SCI.

performs poorly in comparison to other channels. Due to the dependency of the CIR on the operating frequency (section 3.2), CSpy’s SCI determination accuracy is low if the SCI is far away (in terms of frequency) from the current operating channel (figure 11(c)). Nonetheless, it is still possible to use the CIRs from a single channel to reliably estimate the strongest channel within a few (4 – 5) adjacent channels.

Effect of training data: CSpy trains an SVM classifier with CIRs collected from different locations. The diversity of the training data affects CSpy’s SCI estimation accuracy. The accuracy increases with the amount of training data (figure 12(a)). By using the training data from 30 locations, it is possible to estimate the SCI with 79% accuracy.

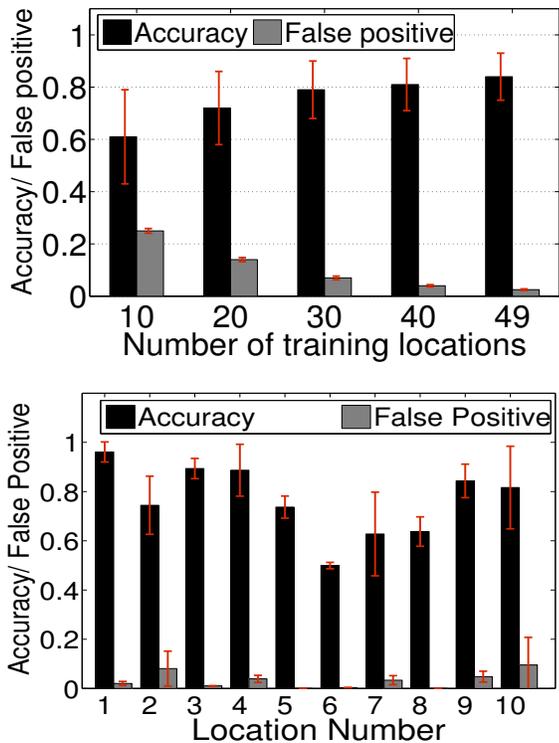


Figure 12: SCI estimation performance (a) at 50 locations with different number training locations in the same building, (b) in the second building by using the training data collected from only the first building.

The CIR features used in CSpy depend solely on the properties of the multipath components, and does not specifically depend on the location they are collected from. Thus, we find that we can train the classifier in the first building and use it to correctly classify the CIRs observed at 10 different locations in a second building (figure 12(b)). Figure 13 presents the accuracy and false-positive graphs for the parking lot experiment. In this scenario, we placed the transmitter right outside the first building, and collected traces at 10 random locations in the adjacent parking lot. We achieved only a modest accuracy in estimating the SCI at the outdoor locations by using the training data collected *only* from the indoor measurements, implying that outdoor and indoor channel characteristics in our measurements are quite different (figure 13(a)). Of course, the accuracy improves if the classifier is also trained with outdoor measurements (figure 13(b)). We further observe that although training separate classifiers for indoor and outdoor measurements may be possible, it improves CSpy’s performance only marginally. We conclude that if a single classifier is sufficiently trained with a variety of CIRs, it may be able to determine the strongest channel at any location with reasonable accuracy.

Effect of feature length: The SCI estimation performance depends on the length of the chosen CIR feature vector. If the length is small, the estimation accuracy may not be satisfactory, whereas, if we use all the 56 complex taps of the CIR vector as the distinguishing feature, it may complicate the classification operation. Figure 13(c) plots the performance of SCI estimation for different feature lengths. Accuracy plateaus after a feature length of 10 because most of the multipath signals arrive at the receiver within a short duration. Figure 13(c) also specifies the average runtime of the classification operation for different feature lengths using a 2.8GHz CPU. As expected, the runtime as well as accuracy increases with the feature length. CSpy uses a feature length of 10 for training and classification to strike a balance between runtime and performance.

6.3 Performance Gain from SCI Detection

In this subsection we demonstrate how CSpy improves the wireless performance by tracking the strongest channel. In a typical WLAN scenario, a wireless link operates on a single fixed channel and does not exploit diversity across channels. CSpy tracks channel quality variations, and switches to the strongest channel whenever necessary. Figure 14(a) demonstrates that CSpy provides similar throughput gain over any

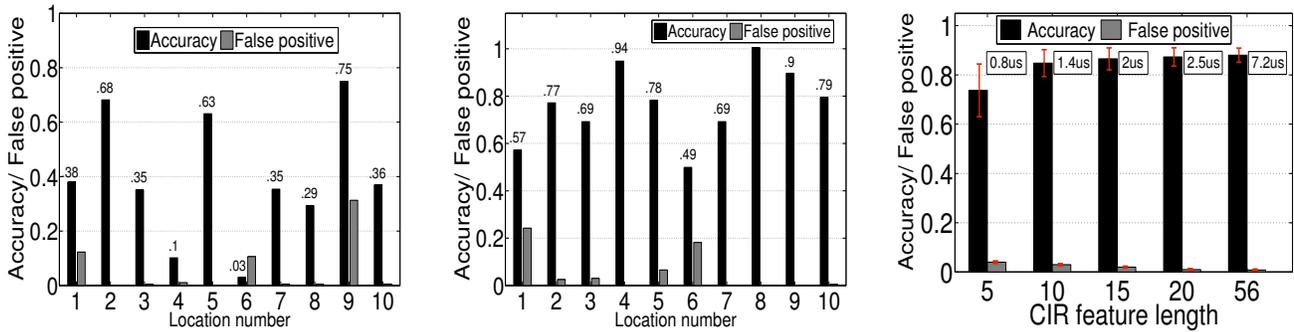


Figure 13: SCI estimation performance at 10 locations in the parking lot by using the training data collected from (a) the first, and the second building only, (b) the parking lot as well as from the two buildings. (c) SCI estimation performance for different CIR feature lengths.

fixed channel scheme. Because channel conditions change quite randomly at different locations, we also find that CSpy's throughput gain varies across different locations, with gain of upto 100% (figure 14(b)) over an average channel. When the current operating channel is the weakest of all the possibilities, CSpy can provide upto 250% improvement in throughput. On average, CSpy also achieves 85% of the throughput achieved by an optimal algorithm which always picks the correct SCI (figure 14(b)).

the channel qualities. Probing-based schemes are inapplicable in mobile scenarios because the overhead is often more than the channel coherence time of a mobile link. Even for static links, it is difficult to avoid the trade-off between probing frequency and SCI estimation accuracy. Frequent probing can track the SCI but incurs large overhead, whereas, the estimated SCI becomes stale if the channels are probed only occasionally. CSpy outperforms probing-based mechanisms by eliminating the need of periodic channel scanning (figure 15).

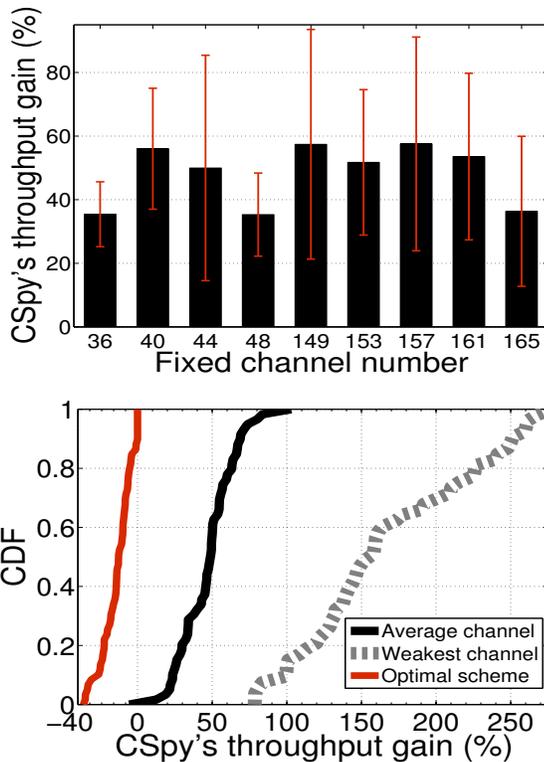


Figure 14: CSpy's throughput gain (a) against fixed channel scheme, (b) against the weakest, an average channel, and the optimal scheme, across 60 indoor locations.

Comparison with probing schemes: Probing all the available channels incurs an overhead (approximately 40ms in our implementation) because it requires scanning across all the available channels, along with packet exchanges to evaluate

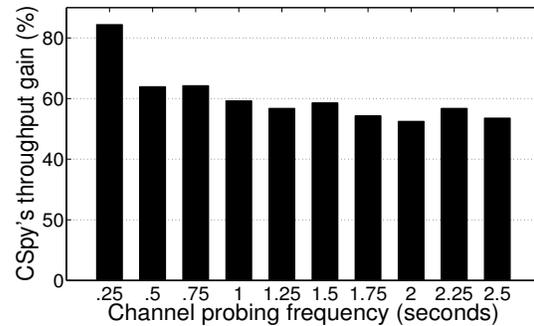


Figure 15: Throughput gain of CSpy over a probing-based scheme.

Effect of mobility: The SCI may change frequently for mobile links, and thus it may not be always beneficial to switch channels due to the costs involved. CSpy switches to the strongest channel if the estimated coherence time is significantly more the channel switching overhead. To evaluate CSpy's performance under mobility, we include the channel switching time (3ms in our implementation) as an overhead in our simulation. Figure 16 shows reduced throughput gain when the link is mobile because CSpy pays a higher overall switching cost. However, thanks to the light-weight estimation scheme, CSpy can still achieve a moderate performance improvement (median 24%) in mobile scenarios, when the channel is changing frequently.

7. RELATED WORK

Opportunistic channel selection is a well studied area in the context of cognitive radio systems [10–12, 22–26]. The goal of most of these studies is to reduce the effect of interference by communicating on channels that are either free or sparsely occupied. WiFi based multi-channel solutions also aim at avoiding channels which are heavily occupied by other

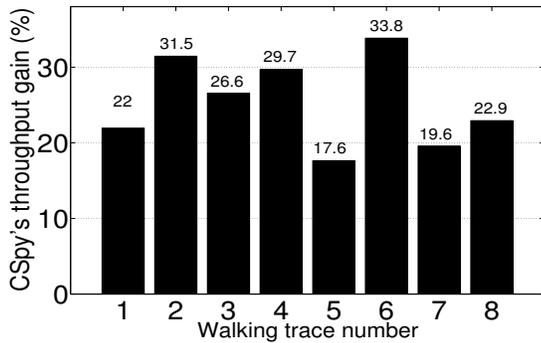


Figure 16: Throughput gain of CSpy for 8 different walking traces

devices. Early proposals required multiple radios to monitor all the available channels [27], which was later resolved using temporal synchronization [28], and channel hopping [29]. Interference-aware channel assignment was also studied in the context of mesh networks [30] and WLANs [31]. Authors in [32, 33] demonstrate that wireless throughput can be improved by allocating spectrum based on bandwidth demand. Recently PHY layer information has been exploited to select channels based on occupancy. Authors in [34] use a subset of OFDM subcarriers to detect narrowband interference, while Jello [35] identifies idle spectrum fragments using power spectrum density maps. While channel selection solutions optimize for a free or lightly-used channel, CSpy's aim is to determine the strongest channel that can deliver the highest PHY bit-rate. CSI-SF [36] combines CFRs from two adjacent channels to reconstruct the CFR of one bonded channel and use it for channel quality prediction; this still requires probing individual sub-channels.

Cellular networks [7] uses opportunistic scheduling by assessing the channel quality of all the registered users and serving clients which have a better channel quality. The concept of opportunistic scheduling has also been applied in multi-channel systems mostly by using probes [1, 5, 6, 37]. The main idea is to evaluate the quality of all the channels to a client by using short probe packets, and thereafter estimate and utilize the strongest channel. The issue is that probing techniques require frequent channel switching which incurs a large overhead. Author in [1] attempts to reduce the probing overhead in the context of whitespaces by utilizing opportunistic sampling, but still incurs a prohibitive overhead in finding the strongest channel. CSpy can be used in the cellular systems to remove or reduce the probing overhead. Cellular systems also employ coherence time based training to predict throughput and bitrate based on SNR [38]. While CSpy can benefit from this framework, it relies on features of multipath to estimate the strongest channel.

8. DISCUSSION AND FUTURE WORK

Dependency on training: As we discuss in Section 5.2, our training algorithm learns the generic properties of CIR. As our evaluation demonstrates, we need to retrain the classifier only for different types of environment: an indoor-trained classifier works in various indoor environments but not outdoors, and vice versa. This is expected, given that the proof of Proposition 1 depends on the channel's delay spread. We leave the

full classification of different classes of environments and how to deal with them for future work.

CSpy in WLAN scenarios: This paper demonstrates the feasibility of improving the performance of a single link by estimating and communicating on the strongest channel. Our findings are directly applicable to scenarios where the transmitter has only one receiver to cater to, e.g., wifi-direct, wireless tethering, wireless displays, etc. If the transmitter (AP) has to serve multiple receivers (clients), it can either choose a fixed channel that is the strongest for majority of the clients, or periodically switch between the strongest channel of each client. Of course, CSpy can perform significantly better if the AP can switch to the strongest channel on a per-packet basis, for each individual client, without incurring large switching costs. This may be possible by using sophisticated physical layer techniques such as RODIN [39] and WiFi-NC [40]. IEEE 802.11ac also defines an enhanced RTS/CTS mechanism that enables the AP and a client to find and use the best sub-channel within a given 40 or 80 MHz bonded channel on a per-packet basis, but it incurs a large RTS/CTS probing overhead. CSpy can also be incrementally deployed across multiple APs, although APs employing CSpy may experience a higher performance than those without CSpy.

Combining CSpy with interference-aware channel selection: Wireless throughput depends on both channel quality and the degree of interference. In some cases, a highly loaded channel may offer better performance over a less loaded channel, if the quality of the highly loaded channel is better. Balancing the channel quality and medium access in an optimal and distributed way is still an open research problem. CSpy optimizes for the former, without considering the effect of interference. Since channels at different locations are independent [7, 8, 41], it is reasonable to assume that SCI will be random and independent at different locations. Thus we speculate that using the channel quality as the channel selection metric across the entire network will yield a balanced distribution of links across channels. Alternatively, one can combine the existing works on interference-aware dynamic channel selection schemes [13, 42, 43] with CSpy. We leave both these directions for future work.

SCI estimation and MIMO: It is straightforward to extend CSpy to a MIMO system. However, MIMO already exploits channel diversity through spatial diversity, hence the expected gains of CSpy over MIMO will be smaller. Indeed, we measure a smaller but still substantial gain of 29% over an average channel and 62% over the weakest channel in the same experimental setting with 2x2 MIMO. This will further decrease with a larger number of antennas. But CSpy could potentially be beneficial in a different MIMO setting, similar to WLAN scenario discussed above, where an AP with a large number of antennas serves multiple clients concurrently. Since in this case the AP effectively assigns one antenna to each client, CSpy can guide the client and channel selection to improve the overall performance.

CSpy for bandwidth selection: In this paper, we assumed channels having homogeneous bandwidth. Recent wireless technologies allow various bandwidth options (20, 40, or 80 MHz in IEEE 802.11ac). In order to maximize the performance gain, not only the center frequency (channel index)

but also the width of channel should be optimally decided. Thus, training and classifying heterogeneous bandwidth options, along with the channel index, is a natural next step.

9. CONCLUSION

Finding a quality channel has been a long-sought but still critical problem in wireless communications, particularly with the systems using unlicensed bands. While existing solutions rely on either active or passive measurements (probing), we propose a mechanism that can find the best quality channel among the candidate channels, purely based on the measurement done in one single channel. We leverage the multipath structure embedded in channel frequency response (CFR), which is a byproduct of OFDM/MIMO processing, to infer the best quality channel. Although the limited granularity of CFR doesn't allow us to reconstruct the CFRs of unseen channels, the aggregate multipath structure extracted from CFR has enough information for us to employ a machine learning classifier to find the best channel. To our best knowledge, this work is the first to present a viable solution to find the best quality channel without probing, and with reasonable accuracy (80%). Our evaluations show reasonable classifier training overhead and realtime classification overhead, hinting an immediate deployment in production networks.

10. ACKNOWLEDGMENT

We sincerely thank Prof. Kyle Jamieon for shepherding our paper, as well as the anonymous reviewers for their valuable comments and feedback. We would also like to extend our gratitude to Sungro Yoon.

11. REFERENCES

- [1] B. Radunovic et al. Dynamic channel, rate selection and scheduling for white spaces. In *ACM CONEXT*, 2011.
- [2] Jeongkeun Lee et al. Understanding the effectiveness of a co-located wireless channel monitoring system. In *IEEE ICC*, 2010.
- [3] Anand Subramanian et al. Understanding channel and interface heterogeneity in multi-channel multi-radio wireless mesh networks. *PAM*, 2009.
- [4] M. Mishra et al. How much white space is there? *University of California, Berkeley, Tech. Rep. UCB/EECS-2009-3*, 2009.
- [5] P. Chaporkar et al. Scheduling with limited information in wireless systems. In *ACM MobiHoc*, 2009.
- [6] A. Sabharwal et al. Opportunistic spectral usage: Bounds and a multi-band csma/ca protocol. *IEEE ToN*, 2007.
- [7] D. Tse and P. Viswanath. *Fundamentals of Wireless Communication*. Cambridge University Press, 2005.
- [8] S. Sen et al. Spot localization using phy layer information. In *MobiSys*, 2012.
- [9] A. Saleh et al. A statistical model for indoor multipath propagation. *IEEE Journal on Selected Areas in Communications*, 1987.
- [10] R. Murty et al. Senseless: A database-driven white spaces network. *IEEE Transactions on Mobile Computing*, 2012.
- [11] H. Zheng et al. Device-centric spectrum management. In *IEEE DySPAN*, 2005.
- [12] Q. Zhao et al. Decentralized cognitive mac for opportunistic spectrum access in ad hoc networks: A pomdp framework. *IEEE Journal on Selected Areas in Communications*, 2007.
- [13] N. Ahmed and S. Keshav. SMARTA: a self-managing architecture for thin access points. In *CoNext*, 2006.
- [14] D. Halperin et al. Predictable 802.11 packet delivery from wireless channel measurements. In *ACM SIGCOMM*, 2010.
- [15] M. Vutukuru et al. Cross-layer wireless bit rate adaptation. In *ACM SIGCOMM*, 2009.
- [16] S. Sen et al. Accurate: Constellation based rate estimation in wireless networks. In *USENIX NSDI*, 2010.
- [17] R. Chandra. A case for adapting channel width in wireless networks. In *ACM SIGCOMM*, 2008.
- [18] WIFI Alliance. Spectrum sharing in the 5 ghz band-dfs best practices. *Spectrum and Regulatory Committee*, 2007.
- [19] E.H. Ong et al. Ieee 802.11 ac: Enhancements for very high throughput wlans. In *IEEE PIMRC 2011*.
- [20] D. Halperin et al. Tool release: gathering 802.11 n traces with channel state information. *ACM SIGCOMM CCR*, 2011.
- [21] S. Ghassemzadeh et al. Measurement and modeling of an ultra-wide bandwidth indoor channel. *IEEE Transactions on Communications*, 2004.
- [22] I. Akyildiz et al. Next generation/dynamic spectrum access/cognitive radio wireless networks: a survey. *Computer Networks*, 2006.
- [23] S. Huang et al. Optimal sensing-transmission structure for dynamic spectrum access. In *IEEE INFOCOM 2009*.
- [24] H. Kim et al. In-band spectrum sensing in cognitive radio networks: energy detection or feature detection? In *ACM MobiCom*, 2008.
- [25] N. Chang et al. Optimal channel probing and transmission scheduling for opportunistic access. In *MobiCom'07*.
- [26] S. Yoon et al. Quicksense: Fast and energy-efficient channel sensing for dynamic spectrum access networks.
- [27] N. Jain et al. A multichannel csma mac protocol with receiver-based channel selection for multihop wireless networks. In *IEEE ICCN*, 2001.
- [28] Jungmin So et al. Multi-channel mac for ad hoc networks: handling multi-channel hidden terminals using a single transceiver. In *ACM MobiHoc*, 2004.
- [29] P. Bahl et al. Ssch: slotted seeded channel hopping for capacity improvement in ieee 802.11 ad-hoc wireless networks. In *ACM MobiCom*, 2004.
- [30] K. Ramachandran et al. Interference-aware channel assignment in multi-radio wireless mesh networks. In *IEEE INFOCOM*, 2006.
- [31] E. Rozner et al. Traffic-aware channel assignment in enterprise wireless lans. In *IEEE ICNP*, 2007.
- [32] T. Moscibroda et al. Load-aware spectrum distribution in wireless lans. In *IEEE ICNP*, 2008.
- [33] S. Rayanchu et al. Fluid: Improving throughputs in enterprise wireless lans through flexible channelization. *ACM MobiCom*, 2011.
- [34] H. Rahul et al. Learning to share: narrowband-friendly wideband networks. *ACM SIGCOMM*, 2008.
- [35] Lei Yang et al. Supporting demanding wireless applications with frequency-agile radios. In *USENIX NSDI*, 2010.
- [36] Riccardo Crepaldei et al. CSI-SF: Estimating wireless channel state using csi sampling & fusion. In *Infocom*, 2012.
- [37] P. Chaporkar et al. Optimal joint probing and transmission strategy for maximizing throughput in wireless systems. *IEEE Journal on Selected Areas in Communications*, 2008.
- [38] A. Schulman et al. Bartendr: a practical approach to energy-aware cellular data scheduling. In *ACM Mobicom*, 2010.
- [39] E. Chai et al. Building efficient spectrum-agile devices for dummies. In *ACM MobiCom*, 2012.
- [40] K. Chintalapudi et al. Wifi-nc: Wifi over narrow channels. In *USENIX NSDI*, 2012.
- [41] Junxing Zhang et al. Advancing wireless link signatures for location distinction. In *MobiCom*. ACM, 2008.
- [42] A. Anandkumar et al. Opportunistic spectrum access with multiple users: learning under competition. In *IEEE INFOCOM*, 2010.
- [43] G. Kasbekar et al. Opportunistic medium access in multi-channel wireless systems: A learning approach. In *Allerton*. IEEE, 2010.