Supporting Real-time Wireless Traffic through a High-Throughput Side Channel

Haoyang Lu and Wei Gao

Department of Electrical Engineering and Computer Science
University of Tennessee at Knoxville
{hlu9,weigao}@utk.edu

ABSTRACT

Performance of modern cognitive and interactive mobile applications highly depends on the data transmission delay in the wireless link that is vital to supporting real-time wireless traffic. To eliminate wireless network congestion caused by large amounts of concurrent network traffic and support such real-time traffic, traditional schemes adopt various flow control and QoS-aware traffic scheduling techniques, but fail when the amount of network traffic further increases. In this paper, we present a novel design of high-throughput wireless side channel, which operates concurrently with the existing wireless network channel over the same spectrum but dedicates to real-time traffic. Our key idea of realizing such a side channel is to exploit the excessive SNR margin in the wireless network to encode data as patterned interference. We design such patterned interference in form of energy erasure over specific subcarriers in an OFDM-based wireless network, and achieve a data rate of 1.25 Mbps in the side channel without affecting the existing wireless network links. Experimental results over software-defined radio platforms demonstrate the effectiveness of our side channel design in reducing the latency of real-time wireless traffic, while providing sufficient data throughput for such traffic.

CCS Concepts

• **Networks** → *Network protocol design; Network experimentation;*

Keywords

wireless networks, wireless side channel, real-time wireless traffic

1. INTRODUCTION

Mobile computing has been an indispensable part of every aspect of modern life, by enabling highly cognitive and interactive mobile applications over heterogeneous types of mobile devices. Representative examples of these mobile

applications include mobile cloud computing [14], resource sharing among mobile systems [2], wearable computing [9], and multi-party mobile gaming [13], which fundamentally transform the way people access information and interact with each other. Performance of these applications depends on the transmission latency of wireless link connecting mobile devices with each other and the remote cloud, so as to support real-time wireless traffic that is vital to prompt application response.

In practice, the real-time wireless traffic of these applications may be seriously delayed when competing with other data traffic being transmitted concurrently over the same wireless channel. Traditional designs of wireless networks eliminate such network congestion via Quality of Service (QoS)-aware traffic scheduling [3] or flow control [4], but fail when the amount of network traffic further increases. Another option to alleviate such network congestion is to allocate additional wireless spectrum that is exclusively used for real-time wireless traffic. For example, a dedicated spectrum is designated as the control plane in cellular networks [16]. However, such exploitation of additional spectrum is infeasible due to the severe scarcity of wireless spectrum resources nowadays.

Instead, another viable solution to removing this fundamental limitation on supporting real-time wireless traffic is to explore a *wireless communication side channel*, which operates concurrently with the existing wireless channel¹ over the same spectrum but dedicates to transmitting real-time traffic. When the main channel is congested, real-time traffic is transmitted through the side channel. Hence, communication of delay-sensitive applications between mobile devices will never be delayed by concurrent wireless traffic, and its latency only depends on the link propagation delay. The key insight of such a side channel is that the Signal-to-Noise Ratio (SNR) of a wireless channel is usually higher than the SNR required to support the data rate being used, due to inaccurate SNR estimation and conservative rate adaptation in wireless networks. This in-band *SNR margin* can be exploited to encode data as *patterned interference* over the main channel. The impact of such a side channel on packet decoding over the main channel, on the other hand, could be efficiently eliminated by limiting the amount of additional patterned interference within the scope of the main channel's SNR margin.

In order to support real-time wireless traffic, the wireless

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.

MobiHoc'16, July 04-08, 2016, Paderborn, Germany

© 2016 ACM. ISBN 978-1-4503-4184-4/16/07...\$15.00

DOI: <http://dx.doi.org/10.1145/2942358.2942386>

¹The existing channel being used in current wireless networks is denoted as the "main channel" throughout the rest of this paper.

side channel has to provide sufficient data throughput that is required by modern interactive mobile applications. Existing schemes have designed the side channel over commodity wireless networks such as WiFi [5] and ZigBee [27], but are limited to achieving a data throughput of several hundreds of kbps. This side channel, hence, can only be used to transmit small network control messages (e.g., RTS/CTS frames) but is far from sufficient for mobile application traffic. Further improving the side channel throughput, in theory, needs to encode more bits into the patterned interference. Such throughput improvement, however, is difficult because the amount of interference that can be exploited for data encoding is limited by the channel SNR margin.

In this paper, we present a practical high-throughput design of the wireless side channel that efficiently supports real-time wireless traffic, by exploiting the unique properties of modern digital modulation methods, particularly OFDM which will be the technical foundation of next-generation high-speed wireless networks (e.g., LTE-A and 5G) [25]. Our basic idea of the side channel design is to encode data as patterned interference by erasing the energy of specific subcarriers in the main channel's OFDM symbols. Since such energy erasure does not increase the RF transmit power, it can be used to encode data into every OFDM symbol in the main channel, hence dramatically increasing the side channel throughput. On the other hand, since OFDM modulates data into separate subcarriers in both time and frequency domains, the amount of patterned interference could be efficiently controlled by interfering only a small portion of OFDM subcarriers, without affecting the packet decoding in the main channel and its resistance to channel contention.

The major challenges of such high-throughput side channel design, however, are two-fold. First, it is hard to ensure precise detection of patterned interference in the form of energy erasure over OFDM subcarriers, which can be easily corrupted by channel noise. A straightforward solution is to detect energy erasure as the subcarrier with the minimum level of energy. However, when the channel noise is sufficiently high, it may increase the energy level of the erased subcarrier and hence results in detection errors. Second, our proposed design also raises unique challenge to channel operation and control. It is difficult to distinguish data transmitted in the side channel from background noise, without knowing the existence of such a side channel in advance. Lightweight techniques for connection establishment over the side channel are hence needed without continuous detection of patterned interference.

To overcome the aforementioned challenges, we first propose a probabilistic approach to detecting patterned interference, which takes the channel noise model into account and achieves a high detection rate over the side channel. This approach is also resistant to the multipath channel fading and ambient noise. We then indicate the existence and the destination of a wireless side channel by exploiting the idle time periods between data frames in the main channel and inserting a specialized preamble between these frames. This approach, therefore, does not require active cooperation of the main channel and will not be affected by the traffic status in the main channel. We have implemented the proposed system design over practical Software-Defined Radio (SDR) hardware platforms, and evaluated the effectiveness of our design over realistic VoIP applications. The experimental results show that our system can provide a side channel

throughput of up to 2.5 Mbps, which is 10 times higher than existing work with minimal impairment to the main channel performance. It also reduces the data transmission delay in commodity 802.11 networks by up to 90%, and significantly eliminates the chance of delay jitters in such networks.

The rest of this paper is organized as follows. Section 2 reviews the existing work. Section 3 gives a brief overview on the principle of our design. Sections 4 and 5 present the technical details of the designs of the side channel and connection establishment. Section 6 evaluates the performance of our system. Finally, Section 7 discusses and Section 8 concludes the paper.

2. RELATED WORK

To efficiently eliminate the network congestion, various QoS control schemes have been proposed to prioritize delay-sensitive network traffic [20], and TCP flow control also helps reduce the channel congestion [12]. However, these designs cannot scale to the increase of network traffic.

Instead, recent studies aim to reduce the transmission latency by exploiting extra communication opportunities beyond the existing wireless channel. Existing designs of the in-band side channel are mainly applied for delivering coordination and control information, so as to improve the network performance without introducing extra costs compared to the out-of-band approaches. The side channel conveying the control messages is operated concurrently with data transmissions in the main channel, without jeopardizing the latter's performance. Wu et al. [27] observes the under-utilization of interference-tolerance provided by most physical layer implementations, and encodes a small amount of control messages by intentionally interfering several chips in each ZigBee symbol. Flashback [5] builds a control plane with a 400 kbps data rate in OFDM-based WiFi packets, by interpolating high-energy spikes into OFDM subcarriers. These spikes transmit at a power level that is 64 times greater than the power used for regular data symbols, hence allowing easier detection. However, all these designs are limited to conveying small-size network control messages over low-throughput side channels. In contrast, the major focus of this paper is to develop a side channel with a throughput that is sufficiently high to support real-time wireless traffic of mobile applications when the main channel is congested. We build our side channel upon commodity WiFi standard and achieve a high throughput by interfering one or more subcarriers in *every* OFDM symbol.

Orthogonal Frequency-Division Multiple Access (OFDMA) [6], a multi-user version of the OFDM design, can be potentially used as an alternative of the side channel, by dedicating a certain amount of subcarriers exclusively for the delay-sensitive traffic. However, the channel throughput for other traffic will inevitably suffer from degradation. For example, if two subcarriers are dedicated to delay-sensitive traffic, the throughput degradation for other traffic will be degraded 4% (2/48). In addition, since our design is implemented at the PHY layer and requires no modification of MAC, it can also work in parallel with other advanced MAC-layer network standards such as 802.11e [10], which prioritizes the delay-sensitive traffic at the MAC. Applications which require immediate responses such as workload offloading in the mobile cloud [7, 8, 24], furthermore, can also benefit from our work.

Our design for connection establishment and control in

Figure 1: Channel BER under different SNR conditions

the side channel is partly inspired by another type of side channel design, which is achieved through appending preambles to the wireless channel. E-Mili [29] constructs the M-preamble to facilitate the sampling-rate invariant packet detection, so that the WiFi receivers could recover from the downclocked idle listening mode once a valid preamble is received. The M-preamble is comprised of duplicated versions of complex Gold sequences (CGS), whose length implicitly conveys the address information. Gap Sense [30] allows interaction between mobile devices with heterogeneous types of wireless networks, by prepending legacy packets with a preamble consisting of multiple energy pulses with quiet periods in between. This quiet period conveys control messages, and can be detected by neighboring nodes with incompatible physical layers. Being different from existing schemes which develop new preambles and hence require specialized hardware or software to operate these preambles, our approach uses a simple duplication of regular WiFi preambles. Hence, our proposed schemes could be easily deployed over commodity WiFi transceivers by being integrated into their firmware, without any modification to the hardware.

3. OVERVIEW

In this section, we first elaborate our motivation of developing a side channel for supporting real-time wireless traffic, by demonstrating the existence of SNR margin in practical wireless channels. Next, we briefly introduce the background of the OFDM-based WiFi physical layer standard. Finally, we present the high-level picture of our proposed side channel design.

3.1 Motivation

As described above, our major motivation of designing the wireless side channel is the existence of SNR margin in the wireless channel, i.e., the difference between the actual channel SNR and the minimum channel SNR that is required to support the channel data rate being used. In practical WiFi networks, such minimum channel SNR is specified in order to maintain a sufficiently low Bit Error Rate (BER). For example, 10.7 dB is the minimum channel SNR required by the 802.11a standard to ensure a channel BER lower than 1%, when a 9 Mbps data rate (BPSK with 3/4 code rate) is being used [23].

To further evaluate the amount of SNR margin in practical WiFi networks, we conduct a preliminary experiment over USRP N210 SDR boards with GNURadio toolkit². More

²<http://gnuradio.org/redmine/projects/gnuradio>

specifically, we emulate different levels of channel SNR at the sender side by adjusting the transmit gain of the sending USRP board. At the receiver side, we estimate the channel status at real-time and compute the BER by comparing the received data frames with expected ones. The experiment results in Figure 1 show the relationship between channel SNR and BER when different data rates are being used³. For example, it shows that 5.8 dB is the minimum required SNR for the 9 Mbps data rate, and the SNR margin is hence 4.9 dB (calculated as the difference between 5.8 dB and the required 10.7 dB by the 802.11a standard). Such SNR margin will further increase when the wireless channel condition improves and a higher data rate is adopted accordingly.

The existence of such SNR margin, therefore, validates our proposed side channel design. Particularly, note that traditional designs of wireless side channels only exploit a small portion of the SNR margin to deliver small-size control messages. In contrast, we envision that the SNR margin can be further exploited to increase the throughput of the side channel without impairing the performance of main channel. The key challenge, however, is how to appropriately limit the amount of patterned interference being applied.

3.2 OFDM Primes

Our proposed design of the wireless side channel builds on Orthogonal Frequency Division Multiplexing (OFDM), which is a multicarrier modulation technique and is widely used in existing WiFi and cellular networks because of its high efficiency of spectrum utilization and resistance to multipath channel fading and inter-carrier interference. Data in OFDM is transmitted concurrently in closely spaced subcarriers, each of which corresponds to a frequency band that is orthogonal to others. Specifically, each 20 MHz WiFi channel comprises of 64 OFDM subcarriers that are spaced 312.5 kHz apart. In an OFDM symbol in the IEEE 802.11a standard, 12 subcarriers (including DC) are null as the guard band, 4 subcarriers are occupied by pre-known pilots for frequency offset compensation, and 48 subcarriers are used for data transmission. The duration of each OFDM symbol is 4 μ s.

As a result, one OFDM packet could be viewed as a two-dimensional time-frequency grid, as shown in Figure 2. All data subcarriers within one packet use the same modulation symbol mapper (e.g., BPSK, QPSK, 16-QAM or 64-QAM) and the same code rate (e.g., 1/2, 2/3 or 3/4), according to the data rate being adopted. Different settings of the modulation symbol mapper approach and code rate yield different data rates and anti-interference capability. For example, BPSK with 1/2 code rate yields the lowest 6 Mbps data rate but is the most resistant to noise, and is hence adopted in poor SNR scenarios. On the contrary, 64-QAM with 3/4 code rate provides the maximum 54 Mbps throughput by using a denser constellation diagram, and is applied in good SNR scenarios.

In each OFDM symbol, inverse FFT is applied to transform the signal from the frequency domain into the time domain. The symbol mapped into each subcarrier thereby can be viewed as a combination of the magnitude and phase information, or frequency response of that particular subcarrier. Therefore, if we manually replace the original data in

³Different data rates are adopted by using different modulation methods and code rates. For example, 6 Mbps is realized by BPSK with the code rate 1/2.

Figure 2: Encoding data as patterned interference into an OFDM symbol

one subcarrier with “0”, it is equivalent to erasing the power in that subcarrier. This characteristic will be the foundation of our side channel design which applies patterned interference onto OFDM subcarriers via energy erasure.

3.3 Big Picture

As shown in Figure 2, our primary design of the side channel is to erase the energy of one specific subcarrier in each OFDM symbol. At the sender’s side, the patterned interference is only applied to one OFDM subcarrier in each symbol, and hence controlled within the scope of channel SNR margin. At the receiver’s side, an erased subcarrier can be detected as the subcarrier with minimum level of energy. Since we do not require any additional hardware or increase the RF transmit power, little energy consumption is incurred by the side channel. As a result, when we erase the energy of one of the 48 OFDM subcarriers⁴, $\lceil \log_2 48 \rceil = 5$ data bits can be encoded into each symbol and the maximum data rate of the side channel is $5b/4\mu s = 1.25$ Mbps. For example in Figure 2, erasing subcarrier 24 transmits data bits $(11000)_2$ over the side channel. In practice, this data rate depends on the accuracy of detecting patterned interference, and we propose a probabilistic approach to ensure precise detection of such interference in Section 4.

In our side channel design, the main channel information modulated in the subcarriers being erased is unrecoverable, which could potentially impact the performance of the main channel. However, such impact could be efficiently controlled and mitigated if prior knowledge of the corrupted (erased) subcarriers is available. More specifically, the wireless receiver first detects the OFDM subcarriers being corrupted by the side channel before feeding the OFDM symbols to the Viterbi decoder in the main channel. Then, the Viterbi decoder in the main channel could utilize such information to reduce the weight of those bits contained in the corrupted subcarriers when doing its best-path search, known as the erasure operation, so as to avoid decoding error and data packet corruption.

4. SIDE CHANNEL DESIGN

In this section, we first describe the overall system design of the wireless side channel over commodity WiFi networks. Then, we present the probabilistic approach to detecting patterned interference, which ensures the detection rate of patterned interference in a wireless channel with additive white Gaussian noise (AWGN).

⁴Four pilot subcarriers are not used.

Figure 3: Side channel design over WiFi networks

4.1 Overall System Design

The system architecture of the side channel over WiFi networks is shown in Figure 3. Specifically, the real-time traffic is denoted as D' and conveyed in the side channel, and other network traffic is denoted as D and transmitted in the main channel. Energy erasure of OFDM subcarriers is implemented through a simple masking and multiplication approach. More specifically, in the transmitter side, a Pulse Position Modulator (PPM) [22] is applied onto D' . Correspondingly, a sequence of binary numbers, denoted as the mask M , is generated with ‘1’ and ‘0’ referring to either reserve or erase the energy in a specific subcarrier. For example, in Figure 2, erasing the energy of the 24th OFDM subcarrier, which corresponds to encode data $(11000)_2$ in the side channel, maps to the following mask:

$$\underbrace{1}_{c_{31}} \underbrace{1 \dots 1}_{c_{24}} \underbrace{0}_{c_{24}} \underbrace{1 \dots 1}_{c_0} \underbrace{1}_{c_0}$$

This mask contains a single ‘0’ in the position of c_{24} , and is then applied on the main channel subcarriers, which is equivalent to erasing subcarrier c_{24} while maintaining all other ones unchanged.

Based on such scheme, the wireless network system consisting of both a main channel and a side channel works as follows. First, the modulation process in the main channel remains unchanged. As shown in Figure 3, the convolutional encoder, which uses one type of the Forward Error Correction (FEC) coding scheme, introduces redundancy into the modulation process for error correction. The symbol mapper transforms the binary sequence to a sequence of constellation points P , according to the modulation methods being used. After the masking procedure, the carrier mapper would assemble an OFDM packet by filling the data sequence into the time-frequency grids. The FFT module then transforms the data from frequency domain to time domain for transmission. Second, at the receiver side, the signal distortion introduced by multipath fading is compensated by the frequency domain equalizer, by measuring the misshaping of the pre-known long training sequences. The packet $\hat{P}$ after equalization is then fed to the symbol demapper and Viterbi decoder for demodulation, and to side channel detector for detection of the erased subcarriers. The side channel information D' is then recovered through pulse position demodulator. Since our design retains the OFDM-based PHY layer unchanged, it could be easily deployed onto commodity WiFi systems through firmware integration without any modification to existing hardware or requiring any extra hardware.

4.2 Improvement of Data Throughput

Based on this basic design of the side channel, the major approach to further improve the data throughput of the side

Figure 4: Constellation diagram of 16-QAM (3/4) under different channel SNR conditions

channel is to erase multiple subcarriers in an OFDM symbol. Let K be the number of subcarriers being erased in each OFDM symbol, the number of data bits in the side channel that can be transmitted per symbol is hence $\lfloor \log_2(C_{48}^K) \rfloor$, which will quickly increase along with the value of K . For example, a throughput of 2.5 Mbps can be achieved by erasing two subcarriers per OFDM symbol. Essentially, this scheme improves the data throughput of the side channel by increasing the amount of patterned interference applied to the main channel. Since such interference is applied as energy erasure, its amount can be arbitrarily increased, in theory, without breaching the radiation power limit.

However, increasing the amount of patterned interference improves the side channel throughput at the cost of increasing the decoding errors and BER in the main channel, as more information being transmitted is removed from the main channel. Therefore, to efficiently control the amount of patterned interference within the scope of channel SNR margin, we limit $K \leq 2$ (4.1% subcarriers) in our design. In practice, the value of K is declared in the SIGNAL symbol in a regular WiFi frame. According to the 802.11 standard [1], the SIGNAL symbol is the first symbol after preambles, followed by DATA symbols, determining the data rate and the length of the frame. To ensure reliable delivery, the SIGNAL symbol is always transferred in the lowest data rate, e.g. 6 Mbps using BPSK. Similarly, we apply the lowest data rate ($K = 1$) in the SIGNAL symbol, to encode the information about the side channel data rate being used over the subsequent DATA symbols.

4.3 Probabilistic Detection of Energy Erasure

Precise detection of patterned interference in form of energy erasure is the vital factor that determines the performance of both the main channel and the side channel. Detection errors not only reduce the throughput of the side channel, but also increase the main channel BER by misinforming the Viterbi decoder to erase the correct main channel bits while keeping the corrupted bits. The impact of such detection errors is even more severe when multiple subcarriers are being erased.

The major challenge hindering precise detection of patterned interference is channel noise. When sufficient noise exists in the erased OFDM subcarrier, this subcarrier may not be the one with the minimum level of energy and results in detection error, when the straightforward detection approach presented in Section 4.1 is being used. According to [21], the received signal $\mathbf{Y}$ in the frequency domain (after FFT) can be written in time-frequency matrix notation as

$$\mathbf{Y} = \text{diag}(\mathbf{X})\mathbf{H} + \mathbf{N} \quad (1)$$

where $\mathbf{N}$ is AWGN noise and $\mathbf{H}$ is the Fourier Transform of the system function of the multipath channel $h(\tau)$. Since OFDM can efficiently reduce the impact of multipath fading by adopting low-rate symbols and pilot-based frequency-domain equalization [28, 15], we consider AWGN as the major source of the distortion.

Based on such modeling of channel noise, our basic idea of improving the detection accuracy is to probabilistically evaluate the chance for each OFDM subcarrier to contain the patterned interference, and determine such probabilities based on subcarriers' energy levels and up-to-date characteristics of channel noise. We denote $\mathcal{S}$ as the set of points on the constellation diagram corresponding to the digital modulation method used in the main channel, and X and Y as the constellation symbols of the transmitted and received signals in a subcarrier, respectively. From the viewpoint of constellation diagram, the procedure of subcarrier erasure is equivalent to designate an extra constellation point, i.e., the origin $O = (0, 0)$, to the side channel signals. Therefore, $X \in \mathcal{S} \cup O$. Then, the posterior probability for this subcarrier to contain the patterned interference can be calculated by the receiver as

$$\mathbb{P}\{X = O|Y\} = \frac{\mathbb{P}\{Y|X = O\} \cdot \mathbb{P}\{X = O\}}{\sum_{s \in \mathcal{S} \cup O} \mathbb{P}\{Y|X = s\} \cdot \mathbb{P}\{X = s\}}. \quad (2)$$

When there is no channel noise, Y always equals to X and the subcarrier with the minimum energy always contains the patterned interference. Otherwise, the smaller Y is, the higher probability is computed in Eq. (2). The prior probability of $\mathbb{P}\{X = O\} = \frac{K}{N}$ where N is the number of OFDM subcarriers that are used to transmit data in the main channel. $\mathbb{P}(X = s) = 1/|\mathcal{S}| \cdot (1 - K/N)$ for all $s \in \mathcal{S}$. With AWGN, the conditional probability $\mathbb{P}\{Y|X = s\}$ can be computed from a two-dimensional Gaussian distribution $\mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$ with $\boldsymbol{\mu} = O$, based on the distance between X and Y in the complex plane. The noise level is reflected by the co-variance matrix $\boldsymbol{\Sigma}$ of Gaussian distribution, which correlates to channel SNR and is continuously estimated at real-time [18].

Since Eq. (2) computes $\mathbb{P}\{Y|X = s\}$ for every $s \in \mathcal{S} \cup O$, such probabilistic detection of patterned interference with a higher order modulation incurs additional computational overhead. For example, the above probability needs to be computed 65 times if 64-QAM modulation is applied. To reduce such overhead, we investigate the spacial distribution of the erased subcarriers in the constellation diagram. As shown in Figure 4(b) and Figure 4(c), with a moderate channel condition where the main channel information could be correctly decoded, almost all the received side channel signals in the constellation diagram are within the circle C that

centers on the origin, whose radius r is the distance between origin and its closest constellation point in the main channel. In other words, the side channel signals are very likely to be falsely classified as the main channel signals with the lowest power in poor SNR scenarios.

In this case, we only need to concern about the received signals with a magnitude less than r , and classify others as the main channel data. We hence replace the set $\mathcal{S}$ with $\mathcal{S}'$ in the denominator of Eq. (2), where $\mathcal{S}'$ only consists of the weakest main channel signals in the constellation diagram. $\mathcal{S}'$ contains at most 4 elements (2 elements for BPSK), which significantly reduces the computational workload of the probabilistic detection approach.

5. CONNECTION ESTABLISHMENT AND TRANSMISSION CONTROL

In our design, since data is transmitted over the side channel as patterned interference, it is difficult to distinguish the side channel from background noise, without knowing the existence of such a side channel in advance. Lightweight techniques for connection establishment over the side channel are hence needed without continuous detection of patterned interference. In this section, we indicate the existence and destination of a side channel by appending specialized preambles to the regular data frames in the main channel, so as to facilitate connection establishment and transmission control in the side channel.

5.1 Connection Establishment

We enable lightweight connection establishment over the side channel by exploiting the idle time periods between data frames in the main channel. For example, the 802.11 standard requires a DCF Interframe Space (DIFS) of $34\mu\text{s}$ between every two WiFi frames [17, 29]. More specifically, each OFDM-based 802.11 data frame starts with Short Training Sequences (STSs), which consists of 10 duplicate short symbols spanning 16 samples, i.e., $0.8\mu\text{s}$ in a 20 MHz bandwidth system each. The frame detection, together with the Automatic Gain Control (AGC) and timing synchronization, is achieved by computing the auto-correlation of the STS. Hence, as shown in Figure 5(a), our approach indicates the existence of the side channel by placing a special preamble, as a set of STSs, to the idle period between two frames in the main channel.

Since a STS is the beginning indicator of a WiFi frame, it can be efficiently detected by the embedded WiFi AGC via auto-correlation without any additional hardware. Such auto-correlation exploits the cyclic characteristic of STSs and takes the power level of the received signal into consideration [19]. In WiFi systems, a new data frame is considered to be detected once the auto-correlation coefficients $c[n]$ of N consecutive samples exceeds the threshold C_{th} . The special preamble indicating the side channel existence, hence, can be detected in the same way. From Figure 5(b) which demonstrates such auto-correlation process with 10 dB of channel SNR, we can see that this special preamble can be efficiently distinguished from the data frame in the main channel, even in low-SNR scenarios.

5.2 Side Channel Addressing

The destination of a side channel may be different from that of the main channel. Hence, we encode the information about the traffic destination of side channel as the relative time location (L) of the special preamble in the idle

(a) Preamble design

(b) Preamble detection via auto-correlation

Figure 5: Design and detection of side channel preamble

period between data frames. The address information is represented as

$$Address = (L - L_{guard} - L_{Preamble})/D, \quad (3)$$

where L_{guard} is the guard interval between the special preamble and the regular WiFi frame, $L_{Preamble}$ is the duration of the special preamble, and D is the resolution of preamble detection. When $L_{guard} = L_{Preamble} = 0$ and D equals to 10 WiFi samples ($0.5\mu\text{s}$), the WiFi receiver can distinguish $34\mu\text{s}/0.5\mu\text{s} = 68$ different time locations in the idle period, allowing $\lceil \log_2 68 \rceil = 6$ data bits being encoded and a maximum number of $2^6 = 64$ devices to be interconnected at the same time. These data bits could also be used to transmit data acknowledgements that are necessary for reliable data transfer.

The actual amount of data bits being encoded also depends on the values of L_{guard} and $L_{Preamble}$. A small L_{guard} or a large $L_{Preamble}$ would mistakenly indicate an immediate start of regular frame, leading to frame detection error in the main channel. On the contrary, a large L_{guard} reduces the number of data bits being encoded, and a small $L_{Preamble}$ makes the preamble detection prone to channel noise. By experimentally investigating with the off-the-shelf WiFi receivers, we found the preamble consisting of 2 STSs with a guard time of 10 samples is the most suitable design choice for the wireless side channel.

6. PERFORMANCE EVALUATION

In this section, we evaluate the performance of our proposed side channel design. Our proposed work is implemented over USRP/GNU Radio platforms. The metrics being used in our evaluations include the data transmission delay, channel throughput, and the detection rate of patterned interference.

6.1 System Implementation and Experiment Setup

We implemented our design over USRP SDR boards with the GNURadio toolkit, which realize an 802.11a WiFi transceiver

Figure 6: Data transmission delay

over the 5.85 GHz frequency band. USRP N210 motherboards and UBX 40 RF daughterboards are used, which operates in the range between 10 MHz and 6 GHz. In the main channel, the frame generation follows the IEEE 802.11a standard [1]: the first OFDM symbol of each frame stores 3-byte Physical Layer Convergence Protocol (PLCP) header indicating the length and data rate of the frame, followed by the MAC Protocol data unit (MPDU) where a random sequence of strings is located onto its frame body. For each OFDM symbol, 32 subcarriers with index -24 to 11 are exploited for the side channel, excluding the DC (index 0) and pilot (index -21 and ± 7) subcarriers.

We use this implementation to conduct experiments in a $5m \times 5m$ office which contains rich multipath fading effect. The sender and receiver are placed out of line of sight and 4 meters away from each other. We tune the channel SNR by adjusting the USRP parameter of Tx Gain, which represents the RF transmit gain within a range from 2 dB to 30 dB. The maximum Tx power then corresponds to 13 dBm. In each experiment, 250 WiFi packets of 1 kB are synthesized and sent with an interval of 2 seconds, ensuring the timely executions of BER computation and other statistical operations on the PC host connected to USRP boards. In addition, the noise invariance is estimated every 5 frames.

Besides the synthesized packets being sent at a consistent rate, we also evaluate our system with realistic real-time WiFi traffic, more specifically, peer-to-peer Voice over IP (VoIP) traffic. To ensure smooth conversations between endpoints, the VoIP traffic has strict requirements on the network transmission delay: any delay exceeding 400 ms will seriously impair the quality of conversation [11]. We choose to use the VoIP traffic generated by Skype, which transfers an uplink UDP packet every 20 ms. Due to the difficulty of directly interacting between the GNURadio platform and the Skype application, we use the network sniffer Wireshark to capture the uplink Skype packets from a laptop, during which we upload large files to Google Drive to simulate congested wireless network scenarios.

6.2 Data Transmission Delay

We first evaluate the data transmission delay over the wireless side channel. In particular, we measure the round-trip delay for the sender to “ping” the receiver. To emulate the network congestion in the main channel, we generate multiple UDP sockets in the main channel between the sender and receiver. In each UDP socket, we continuously transmit packets of 1 kB every 100 ms. At the PC host operating the USRP board, each UDP socket is operated in a separate application thread and the time for each transmission is individually recorded. Meanwhile, the same type of packets are sent through the side channel with the same

Figure 7: Evaluation with VoIP traffic

interval between consecutive packets.

The round-trip data transmission delay through the main channel and the side channel, with different number of UDP sockets in the main channel, is shown in Figure 6(a). When the number of UDP sockets increases, the main channel is getting congested due to their contention for channel access, and the data transmission delay increases linearly. However, the transmission delay over the side channel remains constantly at a low level in all cases, because data delivery will never be delayed in the side channel and only experience the link propagation delay. It's noted that there is a slightly descending trend of the side channel delay with a larger number of UDP sockets in the main channel. The major reason of this descending trend is the multi-threaded decoding operations at the PC host, which affects the packet processing delay. The OS at the PC host would provide more than enough resources when multiple threads are used.

Besides, we evaluate the data transmission delay under different channel SNR scenarios. We use 4 UDP sockets in our experiment to emulate the main channel traffic, and tune the RF Tx gain to simulate different SNR conditions. The main channel data rate is configured to be 24 Mbps. As shown in Figure 6(b), when the channel SNR drops, the data transmission delay in the main channel quickly increases due to the higher chance of packet corruption and retransmission. Comparatively, data traffic in the side channel is only slightly affected and experiences a delay increase less than 10 ms. Such resistance to SNR degradation and fluctuation is another major advantage of our side channel design, and is important to support real-time data traffic.

The realistic VoIP traffic has also been exploited to evaluate the effectiveness of our design in supporting real-time applications in practice. In our experiment, the data rate in the main channel is set to be 18 Mbps. The characteristics of the UDP VoIP traffic and concurrent TCP traffic, which are recorded as described in Section 6.1, are shown in Figure 7(a). After packet extraction and reconstruction, we set one USRP board as the sender to deliver the recorded packets concurrently to another two USRP boards, which are set to be the recipients of the VoIP packets and other traffic packets, respectively.

The VoIP packets are carried by the side channel and the main channel in two separate experiments for fair comparison, and the transmission delays under these two cases are shown in Figure 7(b). The results in Figure 7(b) shows that the side channel can effectively reduce the transmission delay by 60% (reducing the average transmission delay from 75.1 ms to 29.7 ms). In addition, since the side channel is dedicated to the real-time traffic, packets being transmitted in the side channel will not experience any channel contention from other ongoing traffic, and hence experience a much small delay jitter.

(a) Erasing one subcarrier (b) Erasing two subcarriers
Figure 8: Detection rate of patterned interference

6.3 Detection Rate of Patterned Interference

We first evaluate the detection rate of patterned interference, when the subcarrier containing patterned interference is detected as the one with the minimum level of energy. The detection rates of such patterned interference, when one and two subcarriers are erased in an OFDM symbol, are shown in Figures 8(a) and 8(b), respectively. When the channel SNR is higher than 10 dB, the detection rate is nearly 100% with BPSK and QPSK. When higher-order modulation schemes such as 16-QAM and 64-QAM are being used, the detection rate would improve as the SNR increases, and 90% detection rate could be achieved under when the SNR is higher than 20 dB.

Figure 8 shows that the detection rate of patterned interference is tightly correlated with the SNR condition. Such correlation can be better illustrated by Figure 9 which plots the constellation diagrams of the received symbols in different SNR conditions. When the SNR is good as shown in Figure 9(b), all the side channel symbols are near the origin on the constellation diagram and can be accurately separated from the adjacent four constellation points used by the main channel. However, when the SNR drops, we notice from Figure 9(a) that the symbols of the main channel and the side channel are mixed together and become harder to be separated. In addition, the detection rate is also affected by the modulation scheme being used. When a higher-level modulation scheme such as 64-QAM is applied, the distance between constellation points becomes smaller, making it more difficult to differentiate between the side channel and the main channel signals.

Table 1: Side channel detection rate with a 9 Mbps data rate

SNR (dB)	4	6	8	10	12	14	16
Basic approach (1)	0.22	0.45	0.72	0.90	0.98	0.99	1.00
Prob. approach (1)	0.34	0.61	0.84	0.93	0.98	0.99	1.00
Basic approach (2)	0.07	0.24	0.56	0.86	0.96	0.98	1.00
Prob. approach (2)	0.18	0.42	0.74	0.90	0.97	0.98	1.00

Based on these results, we further evaluate the performance of our proposed approach in Section 4.3 to probabilistic detection of patterned interference. The evaluation results are summarized in Table 1, where (1) and (2) denote the cases in which one and two subcarriers are erased in each OFDM symbol, respectively. Generally speaking, the probabilistic approach of patterned interference detection is able to improve the detection rate by 20%, especially in low SNR scenarios.

6.4 Data Throughput of the Side Channel

Built on precise detection of patterned interference, we

(a) 10 dB (b) 20 dB
Figure 9: Constellation diagrams under different SNR (QPSK 3/4)

evaluate the data throughput of the side channel in practical network scenarios with different channel SNR conditions. Intuitively, the higher the channel SNR is, the more accurate the patterned interference could be detected, and hence the higher data throughput could be realized in the side channel. Table 2 describes the throughput of the side channel with one subcarrier being erased in each OFDM symbol. The side channel could achieve a throughput higher than 1 Mbps with BPSK when the SNR is higher than 8 dB. The maximum data rate of 1.25 Mbps can be approached when the SNR is higher than 16 dB.

Table 2: Side channel throughput (Mbps) with one subcarrier being erased in each OFDM symbol

SNR (dB)	4	8	12	16	24	30
BPSK	0.781	1.170	1.213	1.233	1.235	1.235
QPSK	0.784	1.176	1.237	1.226	1.235	1.225
16-QAM	0.397	0.807	1.116	1.186	1.188	1.210
64-QAM	0.350	0.620	0.888	1.101	1.157	1.200

Furthermore, when two subcarriers are erased in each OFDM symbol, the minimum SNR required to achieve 95% of the maximum throughput of 2.5 Mbps is shown in Table 3. Compared to Table 2, at least another 3 dB is required in the channel SNR to maximize the data throughput.

Table 3: SNR requirement with two subcarriers being erased in each OFDM symbol

Modulation type	BPSK	QPSK	16-QAM	64-QAM
Minimum SNR (dB)	13	15	20	27

6.5 Impacts on the Main Channel

The encoded information conveyed by the subcarriers being erased is entirely discarded, which will inevitably have impacts on the capability of noise resistance of the main channel. To evaluate the impact of our side channel design on the main channel, we introduce two SNR values, denoted as $SNR_{0.01}$ and $SNR_{0.001}$, under which the main channel BER reaches 1% and 0.1%, respectively. Then, when the side channel is enforced to the main channel with a BER of 1%, we evaluate the increase of the main channel BER due to the existence of the side channel, and the results are shown in Figure 10. Since our proposed design is able to appropriately control the amount of patterned interference being applied to the main channel, we are able to limit the increase of the main channel BER within 0.6%, and hence retain the performance of the main channel in various practical scenarios.

The BER increase under different SNR conditions in the main channel is shown in Figure 10(b). Such increase could be efficiently controlled between 2% to 4% for all the eight

Figure 10: Impact on the main channel: (a) BER increase under 1% main channel BER; (b) BER increase under different SNR conditions

possible data rates being used in WiFi networks. By comparing Figure 10(b) with Figure 1, we notice that the BER increase only exceeds 1% when the channel SNR is lower than the required SNR to support the corresponding data rate. For example, the BER increase for 64-QAM (54 Mbps) is only noticeable when the channel SNR is lower than 17.5 dB, which is much lower than the required 26.7 dB to support 64-QAM [23]. Therefore, in most of practical scenarios, the existence of the side channel produces only negligible impacts on the main channel.

6.6 Accuracy of Preamble Detection

In this section, we evaluate our proposed design of connection establishment and transmission control in Section 5. In our experiment, we adjust the parameter D from 10 samples ($0.5\mu s$) to 30 samples ($1.5\mu s$). The mis-detection ratio, as shown in Figure 11(a), can be efficiently controlled below 1% when SNR is higher than 10 dB, and only increases to more than 20% when SNR is lower than 5 dB.

On the other hand, we also evaluate the false addressing ratio, which is defined as the percentage of incorrect interpretation of the address information being encoded in the preamble. The evaluation results are shown in Figure 11(b). The addressing error is mainly resulted from the timing inaccuracy of the internal clock at the WiFi receiver, and hence can be eliminated by increasing D . When the value of D increases, the system is more tolerant to channel noise and reduces the false addressing ratio, at a cost of a smaller addressing space being supported. Also, the false addressing ratio can be efficiently suppressed below 1% when SNR exceeds 10 dB. Given that 10 dB is the minimum SNR required to maintain a WiFi connection in practice [23], we prefer to use $D = 10$ in our real application and hence allow 6-bit addresses being encoded into the preamble.

6.7 Power Consumption of the Side Channel

Theoretically, subcarrier erasure reduces the RF power consumed by the WiFi frontend, and the side channel, hence, incurs limited additional overhead to wireless network operations. In practice, since it is difficult to directly measure the power consumed by the RF frontend on USRP daughterboard, we instead measure the total power consumed by a USRP board that also includes the power consumed by power ICs, FPGA cores and so on. We use a DC power supplier HY3005F-3 to power up the USRP with 6 V fixed voltage supply, and evaluate the power consumption by measuring the current. Our experimental results show that, when a USRP board is in idle status, it draws 2.29A from the power supply. A running 802.11a system between two USRP boards, then, increases such energy consumption to 2.62A.

Figure 11: Accuracy of preamble detection

However, when we further apply the side channel onto the WiFi transceiver and increase the number of erased OFDM subcarriers in each symbol from 0 to 3, there is not any noticeable change on the power consumption. Since the resolution of power measurement of the power supply we use is 10mA, we conclude that the power consumption if the side channel is negligible compared to the main channel ($2.62A - 2.29A = 330mA$). In the future, we will adopt a testing device with a higher measurement resolution to investigate the exact power consumption of the side channel.

7. DISCUSSIONS

7.1 Reducing the Main Channel BER

In our current design of the wireless side channel, erasing OFDM subcarriers will inevitably result in data loss and BER increase in the main channel. Our experiment results in Section 6.5 have shown that the existence of the side channel could potentially increase the BER of the main channel by up to 1% in practical scenarios. To better understand such increase of the main channel BER, we categorize the patterned interference being applied on the main channel into two groups, i.e., the destructive interference and non-destructive interference. In destructive interference design such as Flashback [5] and our current design, the entire OFDM subcarrier is interfered, and we are unable to estimate the original waveform in the interfered subcarrier from the interfered signal. Although this design reduces the complexity of interference detection and the operations of the side channel, it raises additional challenge to reduce the main channel BER.

On the other hand, a design based on non-destructive interference allows estimation of the original signal waveform in the interfered subcarrier from the received signal. A viable solution to such estimation is to interpolate weak interference to an OFDM subcarrier. Such interpolation, however, requires the design of a highly robust signal and channel detector, especially in low SNR scenarios. Such robustness could be possibly achieved by exploiting interference cancellation techniques which cancel the intended interference and recover the information being interfered [26]. Exploitation of such interference cancellation and development of non-destructive patterned interference approaches will be our future work.

7.2 Further Improvement of Side Channel Throughput

Another limitation of our current design is that, the designated throughput of the side channel remains constant in different SNR scenarios. In other words, the increase of channel SNR does not benefit the side channel, whose data rate is limited by the elapsed time of an OFDM symbol. However when the channel SNR is sufficiently high, interfer-

ence that is weaker than energy erasure could also be possibly detected. The throughput of the side channel, hence, could be further increased by appending more interference.

To further increase such throughput, a modulation approach in the side channel based on non-destructive interference is desired, and a new data rate adaptation mechanism should be introduced in the side channel to take the channel SNR condition and the modulation scheme in the main channel into account. Under good SNR conditions, a viable solution is to utilize minimum but detectable non-destructive interference to intentionally interfere each sub-carrier in an OFDM symbol, which yields a big increase in the side channel throughput.

8. CONCLUSION

In this paper, we present a novel design of high-throughput wireless side channel, which efficiently supports real-time wireless traffic without consuming any additional wireless spectrum resource. The basic idea of our side channel design is to exploit the SNR margin in the main channel to encode data as patterned interference, and realize such patterned interference as energy erasure over OFDM subcarriers. We have implemented and evaluated the side channel design over practical SDR platforms. The experiment results verify the effectiveness of the side channel in reducing the data transmission latency and providing a data throughput higher than 1 Mbps, with minimum impact on the performance of the main wireless channel. Our future work will incorporate the proposed side channel design into latest wireless network standards, such as 802.11n and 802.11ac.

Acknowledgments

We would like to thank anonymous reviewers for their insightful comments and helpful suggestions. This work was supported in part by the National Science Foundation (NSF) under grant number CNS-1456656, CNS-1526769 and CNS-1553395. This work was also supported in part by the Army Research Office (ARO) under grant number W911NF-15-1-0221.

9. REFERENCES

- [1] Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications. *IEEE Standard 802.11-2012*, March 2012.
- [2] A. Amiri Sani, K. Boos, M. H. Yun, and L. Zhong. Rio: a system solution for sharing I/O between mobile systems. In *ACM MobiSys*, 2014.
- [3] L. Chen and W. B. Heinzelman. QoS-aware routing based on bandwidth estimation for mobile ad hoc networks. *IEEE Journal on Selected Areas in Communications*, 2005.
- [4] D.-M. Chiu and R. Jain. Analysis of the increase and decrease algorithms for congestion avoidance in computer networks. *Computer Networks and ISDN systems*, 1989.
- [5] A. Cidon, K. Nagaraj, S. Katti, and P. Viswanath. Flashback: Decoupled lightweight wireless control. *ACM SIGCOMM*, 2012.
- [6] M. Ergen, S. Coleri, and P. Varaiya. QoS aware adaptive resource allocation techniques for fair scheduling in ofdma based broadband wireless access systems. *IEEE Trans. on Broadcasting*, 2003.
- [7] W. Gao, Y. Li, H. Lu, T. Wang, and C. Liu. On exploiting dynamic execution patterns for workload offloading in mobile cloud applications. In *IEEE ICNP*, 2014.
- [8] Y. Geng, W. Hu, Y. Yang, W. Gao, and G. Cao. Energy-efficient computation offloading in cellular networks. In *IEEE ICNP*, 2015.
- [9] K. Ha, Z. Chen, W. Hu, W. Richter, P. Pillai, and M. Satyanarayanan. Towards Wearable Cognitive Assistance. In *ACM MobiSys*, pages 68–81, 2014.
- [10] Z.-n. Kong, D. H. Tsang, B. Bensaou, and D. Gao. Performance analysis of IEEE 802.11e contention-based channel access. *Selected Areas in Communications, IEEE Journal on*, 2004.
- [11] J. F. Kurose and K. W. Ross. *Computer Networking: A Top-Down Approach (6th Edition)*. Pearson, 2012.
- [12] L. B. Le, E. Modiano, and N. B. Shroff. Optimal control of wireless networks with finite buffers. *IEEE/ACM Trans. on Networking*, 2012.
- [13] K. Lee, D. Chu, E. Cuervo, A. Wolman, and J. Flinn. Demo: DeLorean: Using speculation to enable low-latency continuous interaction for mobile cloud gaming. In *ACM MobiSys*, 2014.
- [14] Y. Li and W. Gao. Code offload with least context migration in the mobile cloud. In *IEEE INFOCOM*, 2015.
- [15] Z. Liu, Y. Xin, and G. B. Giannakis. Linear constellation precoding for OFDM with maximum multipath diversity and coding gains. *IEEE Trans. on Communications*, 2003.
- [16] R. Love, R. Kuchibhotla, A. Ghosh, R. Ratasuk, B. Classon, and Y. Blankenship. Downlink control channel design for 3GPP LTE. In *IEEE WCNC*, 2008.
- [17] E. Magistretti, K. K. Chintalapudi, B. Radunovic, and R. Ramjee. WiFi-Nano: Reclaiming WiFi efficiency through 800ns slots. In *ACM MobiCom*, 2011.
- [18] D. R. Pauluzzi and N. C. Beaulieu. A comparison of SNR estimation techniques for the AWGN channel. *IEEE Trans. on Communications*, 2000.
- [19] T. M. Schmidl and D. C. Cox. Robust frequency and timing synchronization for OFDM. *IEEE Trans. on Communications*, 45(12):1613–1621, 1997.
- [20] W.-L. Shen, K. C.-J. Lin, S. Gollakota, and M.-S. Chen. Rate adaptation for 802.11 multiuser MIMO networks. *IEEE Trans. on Mobile Computing*, 2014.
- [21] O. Simeone, Y. Bar-Ness, and U. Spagnolini. Pilot-based channel estimation for OFDM systems by tracking the delay-subspace. *IEEE Trans. on Wireless Communications*, 2004.
- [22] H. Sugiyama and K. Nosu. MPPM: a method for improving the band-utilization efficiency in optical PPM. *Journal of Lightwave Technology*, 1989.
- [23] J. Thomson, B. Baas, E. M. Cooper, J. M. Gilbert, G. Hsieh, P. Husted, A. Lokanathan, J. S. Kuskin, D. McCracken, B. McFarland, et al. An integrated 802.11a baseband and MAC processor. In *IEEE International Solid-State Circuits Conference*, pages 126–451, 2002.
- [24] L. Tong and W. Gao. Application-aware traffic scheduling for workload offloading in mobile clouds. In *IEEE INFOCOM*, 2016.
- [25] T. Wild, F. Schaich, and Y. Chen. 5G air interface design based on universal filtered (UF-)OFDM. In *IEEE DSP*, 2014.
- [26] K. Wu, H. Li, L. Wang, Y. Yi, Y. Liu, Q. Zhang, and L. Ni. HJam: Attachment transmission in WLANs. In *IEEE INFOCOM*, 2012.
- [27] K. Wu, H. Tan, Y. Liu, J. Zhang, Q. Zhang, and L. Ni. Side channel: Bits over interference. In *ACM Mobicom*, 2010.
- [28] B. Yang, K. Letaief, R. S. Cheng, and Z. Cao. Channel estimation for OFDM transmission in multipath fading channels based on parametric channel modeling. *IEEE Trans. on Communications*, 2001.
- [29] X. Zhang and K. G. Shin. E-MiLi: energy-minimizing idle listening in wireless networks. In *ACM MobiCom*, 2011.
- [30] X. Zhang and K. G. Shin. Gap sense: Lightweight coordination of heterogeneous wireless devices. In *IEEE INFOCOM*, 2013.