

Tag Spotting: communicating beyond carrier sense range

Horia Vlad Balan
USC
hbalan@usc.edu

Konstantinos Psounis
USC
kpsounis@usc.edu

USC CENG-2010-5

ABSTRACT

A fundamental difference between wireless and wired networks is the presence of interference among wireless links, which introduces dependencies among flows that do not share a link or node. As a result, when designing a resource allocation scheme, be it a medium access scheduler or a flow rate controller, one needs to consider the interdependence among nodes within carrier sense of each other. In other words, the control plane information needs to reach nearby nodes which often lie outside the communication range, but within the carrier sense range of a node of interest.

But how can one communicate control plane information well beyond the existing communication range? To address this fundamental need we introduce tag spotting. Tag spotting refers to a communication system which allows information decoding at SNR values as low as 0 dB. It does this by employing a number of tricks including adding redundancy to multitone modulation, shaping the spectrum to reduce inter-carrier interference, and the use of algebraic coding. Using real-world experiments on an OFDM system built with software radios, we show that we can decode data at the target SNR value with a 6% overhead; that is, 6% of our packet is used for our low-SNR decodable tags (which carry up to a couple of bytes of data in our testbed), while the remaining 94% is used for traditional header and payload data. We also demonstrate via simulations how tag spotting can be used in implementing fair and efficient rate control and scheduling schemes in the context of wireless multi-hop networks, while pointing out that the idea of tag spotting is useful in the context of any wireless network where control-plane information must travel beyond the communication range of a node.

Categories and Subject Descriptors

C.2.2 [Computer System Organization]: Computer Communication Networks

General Terms

Design, Experimentation, Performance

Keywords

Wireless, Tag Spotting, Software Radios, Resource Allocation

1. INTRODUCTION

Many of the challenges encountered in the design of wireless networks with multiple transmission and reception points stem from the quirks of wireless signal propagation. Using currently prevailing transmission techniques, wireless signals cannot be focused exclusively towards their intended recipient, making wireless an inherently shared medium. Wireless transmissions are local in their coverage, and, in general, no sender or receiver will have access to complete channel state information. Because of these characteristics, a wireless network is commonly modeled as a set of links among which interference may occur depending on the particular choice of senders transmitting at the same time. The effects of wireless interference are far reaching, affecting all network layers, from physical layer and medium access to flow control and user satisfaction. They extend beyond the space of a single host or a single link, as flows that do not share any hosts or links in their paths might in fact find themselves competing for resources. Its direct consequence is unfairness leading to flow starvation and underutilization of available resources. A study of the exact mechanisms through which interference leads to unfairness reveals problems at multiple network layers. The most general statements of these problems frequently preclude finding a decentralized and optimal solution. However, interference is a local disruption, and therefore leaves hope that a local, if imperfect, solution may be found.

Distributed algorithms often make use of local exchanges of information. This creates a need for a communication backplane capable of connecting each host to the set of hosts affected by its transmissions. This requirement is more cumbersome than it might seem at first sight, for successful data transmission at common data rates requires rather large signal to noise ratios. The capacity of links is however affected even by interferers reaching them at far lower signal levels. Connecting the recipients of interference with transmitters requires thus a communication backplane capable of operating over channels offering low signal-to-noise ratios. This, in turn, implies that constructing such a com-

munication backplane will require designing a physical layer different from the standard physical layers used for high rate data transmissions. Further differences arise from the fact that, in a wireless environment designed to support primarily data transmissions over short links operating at strong signal levels, long range communication is at best opportunistic. Backplane communication receivers are therefore required to be able to discriminate actual backplane transmissions from high levels of background chatter. Moreover, in order to offer a significant improvement without further aggravating the interference problem, communications along the backplane should not create new interference constraints.

In this paper we propose a signaling scheme enabling the creation of a communication backplane which meets all the above requirements. Our scheme induces a low per-packet overhead, is resilient to high levels of noise and interference, aims to minimize its disruption of data transmissions through added interference and does not require destination hosts to be reachable at all times.

Our first contribution is the design of the Tag Spotting scheme. This scheme makes use of tags, members of a set of signals designed to be easily detectable and recognizable in the presence of high levels of noise and interference, in the absence of time and phase synchronization and with only approximate frequency synchronization. Since tag transmissions affect concurrent data transmissions, the distribution of transmitted power of tags has been designed to be similar to the distribution of power for data transmissions, in both time and frequency domains. This apparently simple constraint prevented us from directly using a modulation known to give good results in low SNR conditions, such as, for example, MFSK [1] [2]. Tag signals are designed to be detectable at distances approaching and exceeding the carrier sense range¹. Their increase in range over regular data transmissions is obtained in part through added redundancy. Tag signals are modulated using multitone modulation over a time duration that is larger than the duration of a regular data-transmitting tone. A tag is a distinct superposition of several tones whose frequencies are chosen according to the codewords of a binary algebraic code. On the receiver side, in the absence of synchronization, tags are recognized using a receiver based on spectral analysis.

What about interference cause by tags on data payload? OFDM, the prevalent modulation for data transmission in today's wireless networks, used in this study as well, exhibits flat time and frequency power densities. Each tag makes use of a limited subset of the tones used

for data transmission without changing their respective power spectral densities. We will prove that interference disruptions caused by tags are comparable to the ones caused by data transmissions.

Our second contribution is the implementation and testing of Tag Spotting through experiments performed using a software radio platform in a testbed comprising senders, receivers and interferers. This series of experiments made use of a tag family composed of 60 different tags, capable of conveying about 5.9 bits of information. Using channel simulations we also test our design with larger tag families allowing more bits of information per tag. The results of these experiments and simulations support the conclusion that communication through tags is effective at SNR² values as low as 0 dB and is robust to the effects of interference.

Our third contribution is the use of Tag Spotting in two applications, showcasing the performance improvements brought by the existence of a control plane able to reach nodes beyond carrier sense range. Specifically, we use tags to efficiently implement a state of the art congestion control scheme for multi-hop networks which requires neighboring nodes, i.e, nodes that interfere with each other, to exchange control information in an effort to fairly share the available bandwidth. We also use tags in order to design and test a simple MAC-layer signaling mechanism meant to prevent the starvation of TCP flows in multi-hop wireless networks.

This paper is organized as follows: In Section 2 we give an overview of related work and indicate some congestion control and scheduling mechanisms that would benefit from the use of tags. Section 3 introduces multitone modulation along with a simple model for estimating the spectral footprint of multitone signals and discusses techniques for limiting inter-carrier interference. Armed with the conclusions of Section 3, we proceed in Section 4 to describe in detail the structure of tags and construct a tag detector capable of distinguishing tags from interference. Section 5 experimentally evaluates, using a software radio platform, the communication range of tags as well as the rate of false detections. It is shown experimentally that tags can be reliably identified at SNR values as low as 0 dB while the likelihood of false detections can be sufficiently limited. The same section evaluates the impact of noise, tag transmissions, and data transmissions, on each other. In Section 6 we discuss a number of important issues like the range over which tags can be decoded and the ability to send more bits per tag using other encoding schemes. Two examples of using tags to facilitate the implementation and improve the performance of congestion control and scheduling are presented in Sections 7

¹Note that carrier sensing range obviously depends on the threshold used by wireless cards. In practice carrier sense detection may occur reliably even when using SNR thresholds lower than 0 dB. Our statement is accurate for the standard thresholds used today in commercial cards.

²Throughout this presentation we understand the noise part of the SNR figure to also include interference power, unless specifically noted otherwise.

and 8. Finally, Section 9 concludes the paper.

2. RELATED WORK

Tag Spotting is closely related to a classic research topic of communication theory, namely information transmission at low signal-to-noise ratios. The motivation of this research has varied from securing the transmitted data such as in the case of spread-spectrum communication [3] to protection against interference in the case of the widely used CDMA standard [4] and achieving long-range transmission [2].

In the wireless networking world, carrier sense [5] can be seen as an example of a message passing mechanism operating beyond the data transmission range. In the context of the 802.11 distributed coordinated function (DCF), carrier sense acts as a binary form of channel state signaling.

Our work is partly motivated by an observation made in a recent study [6], stating that local wireless networks will not be able to achieve fairness while still using CSMA/CA and conventional AIMD-based rate controllers as long as the overhearing range of wireless transmissions will not surpass their transmission range and approach the carrier sense range. Another motivation is the recent development of a number of congestion control and scheduling schemes for multi-hop wireless networks that are based on the local sharing of information, such as, for example [7] [8] and [9]. While these schemes append control plane information to data packets and rely on packet overhearing, it has been recognized that this information needs to reach all nodes within the carrier sense range of a node of interest.

Prior works on congestion control for multi-hop wireless networks differ in the way in which congestion is reported to the source. One class of schemes sends implicit or imprecise feedback by dropping or marking packets [7, 10] in the tradition of TCP congestion control [11], or by regulating transmissions based on queue differentials [12] along the lines of back-pressure ideas [13]. Another class of schemes [7, 14] explicitly computes available channel capacity and then sends precise rate feedback, in the spirit of wired network congestion control mechanisms such as XCP [15] and RCP [16]. The information sharing mechanism proposed in this paper eases the implementation of many of these ideas and improves their performance (see Section 7), as control information will reach nodes outside the data transmission but within the carrier sense range.

In an effort to tackle the complexity of creating optimal schedulers, recent work on medium access for multi-hop networks has proposed distributed algorithms capable of approximating the optimal solution [17] [18] [19] [20] [21]. A common theme of these efforts is the use of local, neighborhood-centered information in achieving a global solution. Local information is at the heart

of several other wireless multi-hop problems: neighbor discovery [22], capacity estimation [10] and signaling congestion and starvation. The mechanism proposed in this paper offers an efficient way to implement the neighborhood-wide sharing of control information in all the schemes mentioned.

The technology of software defined radios has acted as an enabler for some of the recent advances in multiuser wireless network research. It allowed, for example, the experimentation of techniques such as zigzag decoding [23], interference cancellation [24] or dynamic bandwidth adaptation [25]. Perhaps the most similar technique to the one presented in this paper is the one of smart broadcast acknowledgments, introduced in [26], in which multitone modulation is used for the purpose of simultaneously conveying positive acknowledgments from multiple receivers.

While tag construction makes use of algebraic codes, the topic of coding is beyond the scope of this work. We refer the interested reader to the excellent textbooks [27] and [28]. Another related topic that is not discussed in depth is the construction and analysis of codes that can operate under frequency-selective fading in wideband channels, a subject discussed among others by [29].

3. AN OFDM PRIMER

The current section presents the basics of OFDM modulation and introduces notation that will be used throughout the rest of the paper.

An OFDM encoder uses an inverse discrete Fourier transform in order to encode a sequence of symbols into a set of tones over a finite time interval, from here on named either a frame or an OFDM symbol. Consider the sequence x of length $N = 2^k$ whose elements are chosen from a signal constellation, arriving for encoding at an IFFT (inverse fast Fourier transform) frame generator. The resulting signal will be generated according to the formula:

$$X(t) = \sum_{\forall k \in \{0, N-1\}} x_k e^{ik2\pi t}.$$

While the different modes have constant amplitude one, the randomly weighted sum of modes has variable amplitude. In the worst case, this leads to a large peak to average power ratio, requiring system components to behave linearly over a high dynamic range. This problem worsens as the number of carriers increases. However, it can be shown that only a small number of input symbol sequences produce signals that exhibit high peak powers [30].

Passing to the continuous Fourier transform exposes a windowing effect and provides us with an insight into the spectral footprint of the generated signal. The discrete spectrum of an OFDM frame is illustrated in the




Figure 1: The discrete and continuous Fourier transforms of an OFDM frame.

upper half of Figure 1 and corresponds to the original encoded symbol sequence x . Since the summed exponentials are bounded in time, their continuous transforms are sinc functions and the spectrum of the encoded signal is the sum of these sinc functions, as depicted in the lower half of Figure 1. The orthogonality of different signals is preserved: the peak of each sinc function is aligned with the zeros of all other sinc functions.

In Section 4 we use an increased number of carriers in order to obtain a fast spectral decay and limit inter-carrier interference. The reasoning behind this approach is encompassed in Equations (1) and (2) below, which describe the effects of selfinterference on OFDM signals and give the speed of their spectral decay. In the presence of a small frequency offset δ , the interference signal added by a carrier with unitary amplitude to a carrier placed k positions away is

$$\begin{aligned} l(k, \delta) &= \int_{[0,1]} e^{2\pi i(k+\delta)t} dt \\ &= \frac{\sin(\pi\delta)}{\pi(k+\delta)} e^{i\pi\delta}, \end{aligned}$$

and in power terms,

$$p(k, \delta) = \text{sinc}^2(\pi(k+\delta)) \leq \frac{1}{(k+\delta)^2}. \quad (1)$$

It can be easily verified that the magnitude of the interference among two carriers does not depend on the absolute difference of carrier frequencies but only on the number of carrier positions separating them. Consider now a carrier at carrier position zero and an infinite block of carriers starting at carrier position k . The average power leaked by this block of carriers into the zero-positioned carrier can be bounded according to

Equation (1) to be:

$$p_b(k) = \sum_{c=k}^{\infty} \frac{1}{c^2} = \frac{\pi^2}{6} - \sum_{c=1}^{k-1} \frac{1}{c^2}, \quad (2)$$

where we assume that symbols on different carriers are independent. The series in Equation (2) is rapidly converging. Based on the observations above we conclude that an increase in the number of carriers will sharpen the spectrum of a multitone signal and accelerate its spectral decay. Moreover, the insertion of a relatively small number of null-carriers will significantly reduce interference between two neighboring active regions of the spectrum.

The system design employed in our experiments uses a choice of parameters similar to 802.11a: there are 64 carriers, 12 of which are designated as null carriers while a further 4 serve as pilot carriers, leaving 48 carriers available for data transmission.

4. TAG SPOTTING

High Level Design Overview. Tag spotting uses a set of signals that are easily detectable and identifiable in low SNR conditions, in the absence of timing and phase synchronization and without perfect frequency synchronization. In the following such header signals will be identified as tags. A tag can be inserted at the beginning of any packet transmission in order to broadcast a small piece of information to neighboring hosts.

Tags are designed around a two-layer structure. At signal level, tags are generated using multitone modulation with a larger number of carriers than the one used in generating regular data frames and with a correspondingly longer time duration. The increase in the number of carriers is used in order to sharpen the spectral footprint of tag signals and ease spectral analysis through a discrete Fourier transform even in the absence of complete frequency synchronization. At a second layer, the symbol sequences encoded over the different carriers are chosen according to an algebraic code in order to enhance tag separation.

The lack of phase, time and to some extent frequency synchronization is the primary constraint in constructing a tag detector. A further hindrance, not found in the case of standard detectors, is the possible presence of significant levels of interference from data transmissions adding to regular receiver or channel noise. The detector design introduced here is based on a Fourier transform-based analysis of the power levels in the different frequency bins corresponding to the 64 data carriers.

Multitone Structure. Figure 2 presents a sampled packet transmission, in time domain representation as




Figure 2: Packet transmission: the signal $s(t)$ (upper pane), a time frequency representation using 512 frequency bins (middle pane) and a time frequency representation using 64 frequency bins (lower pane).

well as in two different frequency domain representations, using 512- and 64-points Fourier transforms. As we will see these are the natural scales for analyzing tags and data frames, respectively. The block start times of the successive frequency analysis intervals have been aligned with the start time of the tag as well as with the start times of the OFDM frames comprising the actual packet. For both tags and data frames the cyclic prefixes have been discarded in order to simplify the frequency domain representation. The packet transmission begins with a tag, transmitted, for reasons to be explained in the following, at 3 dB lower total power than the rest of the packet. The transmission continues with six 64-carrier PN-sequence based frames that are used in order to achieve block boundary and frequency synchronization as well as to assess channel gain. The rest of the transmission is composed of a large number of 48-carrier frames that encode the packet payload.

The frequency spectrum of the tag can be observed in the third column of the middle pane. The tag itself spans 512 samples (the duration of eight regular OFDM symbols) to which a cyclic prefix spanning two further OFDM symbol durations is added in order to compensate for the lack of timing synchronization. The added redundancy brought by this eight-fold increase in the number of samples translates, as it is well-known, into a 9 dB SNR gain.

The spectrum of a tag is shaped around regular carriers, which are activated or deactivated, i.e., are used

for signal transmission or left vacant, according to an on/off pattern. The following explains the construction of this pattern; as mentioned previously the system design uses 64 different carriers. For tag construction purposes, eight of those have been designated null carriers, as opposed to the twelve null carriers used in payload data transmission. The remaining 56 carriers are to be grouped into 28 groups of two neighboring carriers. In each such small group only one of the two carriers will be activated. Each two-carrier group can be in one of two states, depending on which one of the two carriers is activated. Every tag can thus be naturally mapped to a 28-bit string in which each bit marks the choice of state for a corresponding group. The resulting modulation scheme is akin to multitone FSK [31] in that it uses a combination of tones in order to transmit information.

On the receiver side, the tag detector operates in the presence of small, tolerable but unknown frequency offsets. In order to prevent frequency offset-caused spillage of power onto neighboring non-active carriers, the density of the carriers is increased by a factor of eight, transforming the 56 non-null carriers into $56 * 8$ “thin” subcarriers. The structure of a tag in frequency domain is depicted in the lower pane of Figure 4 while a two-carrier subgroup is depicted in Figure 3.

As illustrated in Figure 3, each group of eight thin subcarriers uses the four central thin carriers in order to deliver the whole power of the original carrier. The central thin carriers, however, encode these tones using




Figure 3: The discrete spectrum of a two-carrier group (amplitude and phase representation) and its continuous power spectrum.

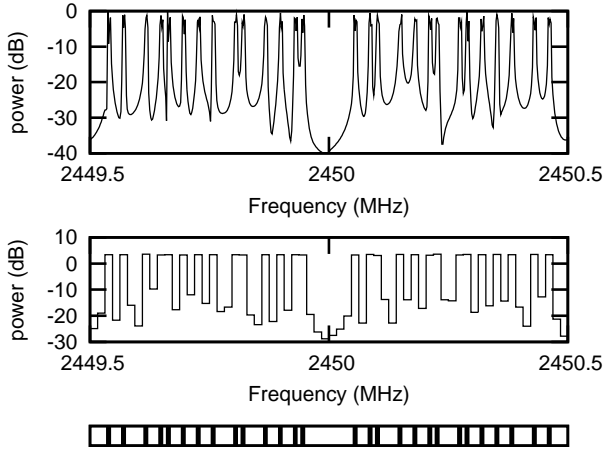


Figure 4: The spectrum of a received tag in the presence of a frequency offset. Upper pane: the 512 frequency bins Fourier transform of the corresponding detector interval. Middle pane: the 64 frequency bins representation used in the detection decision. Lower pane: the structure of the transmitted tag.

different random phases. Looking at Equation (2) it can be readily seen that this results in a faster decay of the original carrier's spectrum and reduces the power leaked onto non-activated carriers. A straightforward computation assuming a uniform frequency offset smaller than the width of two thin carriers reveals that the expected power leaked is, in expectation, about 2.3% of the total power. Figure 4 illustrates this aspect by presenting the spectrum of a received tag in the presence of a frequency offset equal to 10% of a regular carrier width. The use of random phases in signal construction has one further advantage: it allows sampling tags from a larger signal set in order to limit the peak to average power ratio of tags, a problem described in Section 3.

Encoding. As it is common in communication system design, modulation is supplemented by a coding layer. The purpose of this layer is to create a subset of maximally differentiable signals, allowing for the con-

struction of a (tag) detector able to operate at low SNR levels.

The tag signal construction detailed above produces 2^{28} different basic tags, too many for efficient detection and insufficiently distinguishable from each other. A further restriction makes use of a $(28, 60, 13)$ ³ nonlinear code proposed by Sloane and Seidel [32]. Out of the basic tags only those having binary representations corresponding to the 60 codewords of this code are preserved. Every two different tags remaining use different symbols in at least 26 carrier positions.

Regular data frames dissipate the transmitted power over 48 carriers while tags make use of only 28 carriers. Since the average spectral densities of used carriers in tags and data transmissions are equal, it results that the transmitted power in the case of tags is lower than the transmitted power in the case of data, as it was observed in Figure 2.

This choice of code is not optimal, i.e., there might be other codes with larger codebooks, providing more bits per tag, which are capable of fulfilling the detection criteria at the same target SNR values. We consider the use of two other codes, the $(28, 2^8, 11)$ and $(28, 2^{13}, 7)$ extended BCH codes [28] in Section 6.

Constructing a Tag Detector. Let $T = \{t_1, t_2, \dots, t_{60}\}$ denote the set of all tags and C_i be the set of all carriers activated when transmitting tag t_i . Let r_f denote the power of the received signal in the frequency bin corresponding to the f -th carrier. Our detector does not assume the channel phase response to be uniform and can therefore be used in a wideband scenario. We compute the following quantity which we will name from now on tag strength:

$$\frac{\sum_{f \in C_i} r_f^2}{\sum_{\forall f} r_f^2}. \quad (3)$$

Tag strength is compared against a fixed threshold γ and in case the threshold is exceeded a possible tag observation is recorded.⁴

³The notation for codes used here is in the form (N, M, D) where N is the binary codeword length, M denotes the number of distinct codewords and D denotes the minimal Hamming distance between any two codewords.

⁴This equation is similar to the of the low-SNR multipulse detector with n samples for a signal with unknown phase $\theta(t)$ varying at each pulse: $s(t) = A * \cos(2\pi ft + \theta(t))$, at a given SNR value $\alpha = \frac{A}{2N}$. Denote by H_S the hypothesis that signal s has been sent and by H_n the hypothesis that no signal has been sent. That detector is based on the equation $\log \frac{p(r|H_S)}{p(r|H_n)} \approx \sum_{i=1}^N r(t_i)^2 > \gamma' \frac{N}{\alpha}$ where γ' is a constant (see [33, p. 293]); in our case, the correction factor $\sum_{\forall f} r_f^2$ can be seen as an approximation of $\frac{A^2}{4} + N = (\alpha + 1)N$, meant to remove the linear dependence of the threshold on the noise power N , allowing thus for added noise-like interference.

A detection interval has the same length as a tag from which the cyclic prefix has been removed. The tag detector component of the receiver analyzes detection intervals starting at time instants spaced one tag cyclic prefix length apart. It results that the successive detection intervals analyzed have significant overlap. Every transmitted tag will completely cover at least one detection interval. The detector processes every detection interval by first computing the Fourier transform of the contained signal, and then producing, based on the resulting spectrum, a vector containing the strength of each tag, according to Equation (3). In order for a tag recognition event to be recorded for any of the tags, the corresponding tag strength element of the vector must, firstly, exceed the threshold value and, secondly, be maximal not only among the other vector elements but also among the elements of all tag strength vectors corresponding to the detection intervals that overlap the current interval. A further detection metric effective in filtering off-band interference is computed for each detection interval by weighting the power levels in different carriers through the carrier's position in the frequency band, summing the resulting values and afterward dividing the result through the total interval power. As long as the resulting "center of mass" is placed in the central quarter of the frequency band, the tag observations are considered valid, otherwise they will be attributed to off-band interference.

In order to reduce the number of intervals analyzed and the likelihood of false alarms, a simple carrier sense scheme is employed. The receiver maintains a running estimate of channel noise and processes only those intervals for which the SNR exceeds -1 dB.

Overhead. Adding a tag to a packet incurs a transmission time overhead. Assume for now that only data packets are tagged and that a typical data packet has a payload of about 1500 bytes. Encoded using the parameters presented in Sections 3 and 5, the payload will span 125 data frames, to which six synchronization frames are appended. A tag spans the equivalent of eight data frames and therefore its overhead is about 6.1% in terms of the normal packet duration.

5. EVALUATION

5.1 Experimental Setup

The current section presents the experimental evaluation of tag spotting.

(a) Hardware and Software Setup. The experiments were conducted using two USRP boards [34], one transmitter and one receiver, in an indoor environment without line of sight and with multiple concrete walls between the sender and the receiver. The USRP boards were fitted with RFX2400 daughterboards capable of

transmitting and receiving in the 2.4 - 2.483 GHz wireless band. The GNU Radio software suite was used for signal transmission and recording. All signal processing was done using the Octave numerical software suite.

(b) System Parameters. The carrier frequency chosen was 2.45 GHz and the bandwidth of the system was set to 1 MHz. In order to obtain a linear channel suitable for OFDM transmission without distortion effects introduced by the transmitter mixer and the receiver, the signal has been oversampled by a factor of four on both transmitter and receiver sides. The received signal has been filtered in order to preserve only the frequency band of interest. The interpolation factor at the transmitter was set to 32 while the decimation factor at the receiver was set to 16. Measurements in indoor environments with and without line of sight connectivity did not reveal variations of the amplitude of channel response within the used bandwidth.

(c) Configuration. The OFDM system structure resembles the one of 802.11a. The system uses 64 carriers out of which 12 are designed as null carriers and 4 are pilot carriers. The pilot carrier information is known in advance to both the transmitter and the receiver. The system uses the Schmidl and Cox algorithm [35] for packet detection and initial CFO estimation and the Tufvesson [36] algorithm for block boundary detection. These two algorithms make use of two frames in the preamble. A further four frames in the preamble are reserved for measuring the channel's amplitude response. The phase response, more likely to undergo a long excursion in a short interval of time, is estimated for each frame using the four known pilot symbol values. The measured phase responses for the pilot symbols serve as inputs to a linear interpolator which estimates the phase responses of all carriers. While many OFDM systems make use of the phase estimators in order to update their carrier frequency offset estimates, it was found in our setup, after examining measurement results, that the frequency jitter was too small to warrant performing any updates of the CFO estimate during a packet duration. The symbol constellation used was 16-QAM. Packet payload was encoded using a rate 1/2 punctured Trellis code. The bandwidth used was 1 MHz, giving a resulting link speed of 2.4Mbps.

While the design of an efficient OFDM transmission and reception system is not the focus of this presentation, having a viable system was a precondition for showing that any interference effects produced by our set of header tags are comparable to interference effects caused by data transmissions and by environmental noise.

(d) Experiment Description. We performed three series of experiments intended to evaluate the impact of

decreasing the signal to noise ratio on the effectiveness of tag spotting, the impact of rising interference power on tag spotting and the disruption caused to data transmission by interference in the form of tags. A further series of experiments studied the dependence of the rate of false alarms on the detection threshold.

(e) Metrics

Tag Strength is the quantity defined in Equation (3), the primary metric for deciding whether a tag observation will be recorded. It is a measure of the ratio of power contained in the frequency bins allotted to a given tag and the total received power. In order for a tag observation to be recorded one of the necessary conditions is that the tag strength must exceed a threshold value γ . In all experiments presented $\gamma = 0.62$ was used. This choice of threshold accomplishes two goals: it is high enough to correspond to a low rate of false alarms, as verified through the experimental results presented in the current section (see Figure 8) and it is low enough to allow detection at the target SNR values. Using the naive assumption that noise (and interference) power contributes equally to the power levels detected in the different frequency bins, the SNR value that corresponds to this threshold can be derived to be about 0.4 dB.

Symbol Error Rate (SER) is measured for the payload of all correctly identified packets, that is, packets for which the packet detection, block boundary start estimation and CFO estimation succeed. It is the primary metric for estimating the effects of various noise and interference levels on data transmission. This metric was considered more fundamental than the bit error rate (BER), which is heavily dependent on the type of coding employed, a system design parameter that varies largely in current designs.

Probability of Detection (P_d) is defined as the probability that a header tag will be correctly detected and identified at different SNR and SINR levels. It is the primary metric for the success of tag spotting.

Probability of False Alarm (P_f) is defined as the likelihood that, in any given detection interval, noise and interference will cause a spurious tag detection and identification in the absence of a tag transmission, denoted as false alarms of the first kind, or an incorrect tag identification in the presence of a tag transmission, denoted as false alarms of the second kind. A high γ threshold value decreases significantly the probability of false alarm and false alarms of the first type are unlikely to occur in regular experiments. In order to determine the likelihood of this type of false alarms, we have conducted

a further series of experiments using half minute-long recorded signal sequences containing ambient radio noise pertaining to standard 802.11b/g transmissions in an office building occupied by numerous wireless networks in order to measure the detector's robustness to different kinds of radio interference. For the chosen threshold value we were not able to generate false alarms of the second kind.

(f) Practical Considerations. The tag detector presented in Section 4 is unable to compensate for frequency variations between the sender and the receiver. The structure of tags and the detection method described makes it possible to tolerate frequency offsets of up to two thin carriers, or about 4 kHz, without perceivable performance impacts. We have found that the clock jitters of the USRP radios are well within this limit, however different radios have initial frequency offsets of up to 200 kHz, necessitating a supplementary calibration step before each experiment. We note that most WiFi radios in deployment today manifest comparable tolerances and that compensating for frequency and phase drift is one of the main tasks of the wireless system designer. In a practical scenario the clock components could be replaced with more accurately designed/packed counterparts, and therefore we conclude that constructing self-sufficient tag detectors is possible.

Even when using perfect signal generators, complete frequency alignment between the sender and the receiver would be impossible in a mobile environment due to the Doppler effect. The current work assumes a static situation in which the relative movements of hosts generate frequency shifts that are negligible when compared to the carrier bandwidth.

5.2 Experimental Results

(a) Impact of Noise. The first series of experiments tries to quantify the range effectiveness of tag spotting in the presence of different levels of noise, in an interference-free environment.

The transmitter was configured to send sequences of 100 packets with random header tags. On the receiver side the transmitted sequence was decoded and the sequence of detected tags was compared to the original transmitted sequence, in order to obtain an estimate of the detection probability P_d . The decoded symbol payload of received packets was compared with the known symbol payload on the transmitter side in order to estimate the SER. The transmission's SNR was estimated for each detected packet using a low-pass filter-based average power estimator.

The power level of the transmitter was varied between levels spaced 3 dB apart, resulting in different channel SNR values.




Figure 5: Experimental results in the presence of channel noise.

Figure 4 presents the spectrum of a received header tag, along with the transmitted signal. It can be readily seen that the power leaked onto non-activated carriers is negligible when compared to the total signal power.

Figure 5 illustrates the results obtained. The upper pane shows the Symbol Error Rate (SER) for the payload as a function of the Signal to Noise ratio (SNR). The curve is typical for a receiver employing 16-QAM modulation, however the receiver appears to exhibit an error floor at the higher SNR values measured. At SNR values of 20-25 dB, the system can sustain data transmission, when using an typical error-correcting code. This curve serves as a reference for the next experiments, in which noise-based disruptions will be replaced with data-like interference and tag-like interference.

The middle pane shows the Tag Strength as a function of the SNR. The curve decreases steadily as the SNR decreases, reaching the threshold value γ around 0 dB. Finally, the lower pane shows the probability of detection as a function of the SNR. It can be seen that the probability of detection is close to one over the entire range considered.

(b) Impact of Interference

Figure 6 present the results of the same experiment in the presence of a second source transmitting an uninterrupted stream of payload-like data. As before the upper pane plots the Symbol Error Rate, the middle pane the Tag Strength, and the lower pane the probability of detection, all as a function of the SNR. The SER has a slightly different behavior in this case, due to




Figure 6: Experimental results in the presence of data-like interference.




Figure 7: Experimental results in the presence of tag-like interference.

the presence of a different type of interference, as can be seen when comparing the SER curve in the presence of data interference with the SER curve in the presence of just noise. The other quantities of interest, tag strength and the probability of detection P_d remain essentially unchanged. The probability of detection climbs a steep curve and quickly settles close to one. We conclude that the tag detector acts almost identically in the presence of pure noise or noise combined with temporary interference.

(c) Impact of Tag Interference on Data

Figure 7 presents the effect of tags on data transmissions. The SER curve is very close to the SER curve of Figure 5, demonstrating that interference from tags does not increase the error likelihood beyond the error likelihood in the presence of comparable levels of noise.

(d) Likelihood of False Alarm

Figure 8 presents the dependence of the average time




Figure 8: Average time between false alarms at different thresholds.

in-between false alarms on the threshold γ , when analyzing recordings of ambient WiFi traffic. Carrier sense has been disabled in this experiment and every input detection interval is analyzed. These results support our choice of detection threshold, since false alarms occur at a rate of less than once during every 20 seconds of analyzed data.

6. DISCUSSION

The experiments presented in the previous section have shown that tag spotting is effective at SNR values as low as 0 dB, in the presence of receiver noise as well as interference from data transmissions. Moreover, it was shown that tags at a certain signal level affect data in a manner similar to similar levels of noise. Finally, it was shown that, with a suitable choice of detection thresholds, false tag spottings occur at a low rate.

Tag range. The possible use of tags in wireless multi-hop networks prompts us to calculate the increase in range, expressed in distance terms, that transmission through tags brings over regular data transmission. It is well-known that power decay exponents are highly dependent on specific environments. Measurements of signal propagation in the 2.4 GHz band described in [37] reveal a dependence of received power on distance of the form $p \propto \frac{1}{r^d}$ where the exponent d varies from 3 for line of sight propagation to about 6 for non-line of sight propagation. Regular data transmissions using the 16-QAM modulation present in our system necessitate a SNR value of about 20 dB while tags are detectable at a SNR value of 0 dB. It results that the ratio of the range of tag communication to the range of the data transmissions can be, depending on the power decay coefficient utilized, anywhere between 2.15 and 4.65. The carrier sense threshold is usually set about 10-20 dB above the noise level [5] [38], which substantiates our claim that Tag Spotting can communicate information beyond the carrier sense range.

More bits per tag. The system we implemented in the testbed makes use of a (28, 60, 13) code and is capable of conveying about 5.9 bits of information within every tag. As we mentioned already, this choice of code is not optimal, i.e., there might be other codes with

larger codebooks, and thus providing more bits per tag, which are capable of fulfilling the detection criteria at the same target SNR values. With this in mind, we have simulated our Tag Spotting system while using two different codes, the (28, 2⁸, 11) and the (28, 2¹³, 7) extended BCH codes, under a frequency-independent fading assumption. With these more complex codebooks, tags were able to carry about 2 bytes of information while still being decodable at 0 dB SNR. While our system deployment uses a relatively narrow band of 1 MHz thereby justifying the frequency-independent fading assumption in our simulation, an algebraic code with a higher Hamming distance will offer superior benefits in an environment that presents frequency-selective fading. It is beyond the scope of this paper to investigate which codes are optimal for different wideband fading models. More generally, it is beyond the scope of this work to optimize the physical layer technology used to send the tags. Our purpose is to show that it is possible to build an efficient low-SNR communication system based on multitone modulation, which can be easily used in today’s systems to communicate control information at, or even beyond, the carrier sense range.

7. CONGESTION CONTROL

In this section and the next one we present the use of Tag Spotting in two applications aimed towards obtaining a fair resource allocation in multi-hop wireless networks. We explore the structure, the granularity and the rate of information that hosts within a neighborhood should exchange in order to solve the fairness problem, keeping in mind that the fairness issue impacts the design of both the medium access and transport layers of the networking stack.

Congestion control requires a supporting feedback loop that delivers information about congestion events to the sender. In wired networks such congestion events are indicated by queue overflows resulting in packet losses [11], increases in queue wait times resulting in increased round trip times [39] [40] or explicit congestion notifications [41]. In a wireless setting, congestion is not always primarily experienced by the flow that causes it and a neighborhood-wide signaling mechanism becomes necessary. Previous work on congestion control in wireless networks, for example [7] and [42], offers important insights into the nature of the information to be shared. Further, the afore mentioned works suggest that all information necessary in order to perform efficient congestion control can be disseminated through broadcasts of local information originating at every one of the network hosts. In particular, such information is piggybacked onto data/ack transmissions and broadcasted to all nodes within communication range of the originating node.

But, as mentioned earlier, this information should

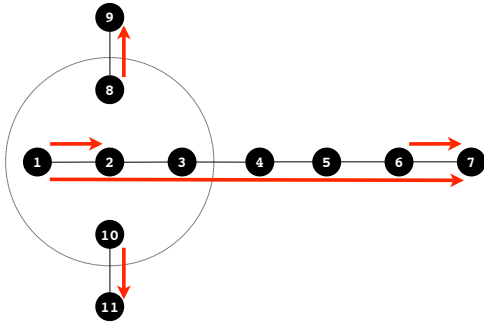


Figure 9: Chain-cross topology with all competing flows separated by at most one transmission range. The large circle indicates the size of the data transmission range.

reach all nodes which potentially interfere with the node of interest, rather than just those within data transmission range. Tag Spotting is capable of passing this type of control information to all neighboring nodes. With this in mind, we extend WCP [7], a recent AIMD-based scheme from the congestion control literature, by using Tag Spotting for communicating congestion notifications, and we assess the achieved performance of both the original WCP protocol and its extended derivative through simulations.

In the original WCP, there are two types of information broadcasted via data and ack packets. Firstly, each node broadcasts a congestion bit to indicate congestion events to its neighbors. These notifications travel back to the sources of the neighborhood flows, which, in turn, reduce the rate of the flows. Secondly, neighboring nodes exchange the round trip times (RTTs) of flows traversing them, in order to compute the maximum RTT of a flow via a neighborhood. This maximum RTT is then used to pace the rate increases of the AIMD controllers which set the rates of the flows traversing the neighborhood. The reasoning behind this is as follows: Since the different traffic flows may not share a single queue, the loss rates as well as the delays experienced by competing flows passing through a congested neighborhood may vary widely, significantly more than in wired networks. This difference would lead the senders’ AIMD controllers to increase their rates at significantly different paces following a congestion event, unless a common time convention like the one mentioned above is used. For more details on WCP, the interested reader is referred to [7].

We have simulated the performance of WCP using the Qualnet network simulator [43]. Qualnet provides the possibility of physical layer simulation using the physical interference model [44]. We have extended Qualnet’s physical layer simulation in order to also handle the likelihood of tag detection using the detection probabilities measured in Section 5. The content of tags




Figure 10: Chain-cross topology with some competing flows separated by more than the transmission range. The large circle indicates the size of the data transmission range.

is composed of one congestion bit and a field that encodes, on a logarithmic scale with base $\sqrt[3]{2}$, the value, in milliseconds, of the longest RTT of all flows traversing the tag emitter. We call the tag-based implementation of WCP, WCP-Tags. For the original WCP we have used a broadcasting mechanism that shares the same information as WCP-Tags, however the limit SNR for broadcast detection has been set at the same value at which successful payload data decoding occurs, since the original WCP broadcasts are inserted into the payload of data packets. All hosts use the regular 802.11 MAC for ad-hoc networks with default simulator values, with the only modification that the number of allowed MAC layer retransmissions has been doubled from its default value in order to decrease the rate of packet drops.

Figure 9 illustrates a textbook configuration for evaluating congestion control protocols in wireless environments. The two short flows on the outside of the central chain of nodes are within the transmission range of node 2, and we can therefore expect that the two variants of WCP will have similar performance. Figure 11 illustrates the rates obtained by TCP, original WCP (“WCP”) and the tag-based implementation of WCP (“WCP-Tags”). It can be readily observed that, while TCP leads to starvation of the central flow, both WCP and WCP-Tags manage a fairer rate allocation.

Discussing the results of these experiments requires taking into account more than the order of magnitude of throughputs achieved by individual flows or the combined throughput of all flows involved. It is well known that supporting a long flow in a wireless multi-hop network is possible only at rates significantly lower than the maximal link speed [45]. Any increase in the rate of the long flow in Figure 9 will involve a drastic reduc-




Figure 11: Goodput results for the topology in Figure 9.

tion in the rates of the other, shorter flows. To make this point more precise, we compute using brute force simulations and the theoretical framework in [46] the max-min rate allocation for these flows and compare it to the other allocations. It is evident that the original WCP and WCP-Tags yield rates which are close to the max-min optimal rate allocation. Thus, it is not possible for the long flow to avoid starvation unless the rates of the short flows are significantly reduced from their TCP allocation. This follows directly from the definition of max-min fairness which states that a rate allocation is max-min fair if it is not possible to increase the rate of a flow from its current value without decreasing the rate of another flow whose current rate is the smaller of the two. Formally, if r_m is the max-min rate vector and r is some other feasible rate vector, then if for some flow a $r(a) > r_m(a)$, there exist some other flow b with $r_m(b) \leq r_m(a)$, for which $r(b) < r_m(b)$.

Figure 10 illustrates a variation of the previous topology in which the original WCP cannot effectively signal congestion between the involved hosts, due to the fact that some hosts are within the interference range but outside the data transmission range of each other. In particular, under the original WCP node 2 cannot inform nodes 8 and 10 that it is congested and the long flow is almost starved, similarly to what happens under TCP. In contrast, WCP-Tags does not starve the long flow as nodes 8 and 10 reduce their rates once they receive notifications through tags that node 2 is congested. The rate allocations achieved by the three protocols as well as the max-min rate allocation for this topology are illustrated in Figure 12. As before, WCP-Tags yields rates which are close to the max-min optimal allocation.

8. SCHEDULING

A large class of scheduling algorithms for wireless multi-hop networks are centered on ideas such as queue equalization using backpressure [17] [9], broadcasting local congestion indications [10] or creating a computationally tractable approximation of an optimal schedule [18] [19] [20] [21]. While some of the proposed sched-




Figure 12: Goodput results for the topology in Figure 10.

ulers assume a perfect knowledge of the distribution of packets in a network and a few others require only information available at every host on which a portion of the algorithm is running, we find that a large number of them could make do with small amounts of information broadcasted by hosts to their wireless neighborhoods. An extensive presentation of all the approaches taken in constructing distributed schedulers, their benefits and their implementation details lies well outside the scope of this paper; however, some of the benefits of schemes that use local communication will be illustrated in the current section through the evaluation of a rather simple scheduling scheme, designed as an extension to the ad-hoc mode of the the 802.11 MAC. The simple mechanism presented here targets some of the unfairness effects introduced by the 802.11 MAC which may lead to flow starvation. The solution presented does not address all scenarios under which starvation can occur and serves only to illustrate the potential of tags in constructing communication-enhanced scheduling schemes. In particular our solution is complementary to RTS/CTS and does not always solve the hidden-terminal induced problem of congestion at the receiver.

Consider again the topology illustrated in Figure 9. The results of Section 7 have already shown that using a standard 802.11 MAC in conjunction with TCP drives the longest flow in this topology into starvation. Our solution preserves TCP as the transport protocol but seeks to relieve such severely disadvantaged flows by enhancing the scheduling algorithm. The key to achieving this goal is an exchange of tags conveying a meaningful measure of starvation, namely the average delay of the packets currently enqueued for transmission.

The hosts observe all detected tags and decide that a host in their neighborhood is starved for medium access whenever they receive a tag conveying an average queueing time which is at least 32 times larger than their own average queueing time. In this case the tag receiver will enter a silence period of 100 milliseconds, allowing the starved host to gain access to the channel and transmit its packets. These numbers are not particularly optimized since our focus is to showcase the




Figure 13: Goodput results under the modified scheduling policy for the topology in Figure 9.

benefits of using tags in scheduling with a very simple addition to 802.11, rather than to provide a fully optimized and tested solution. All other medium access activity proceeds according to the normal 802.11 specification. We call this scheme 802.11-Tags. The only other MAC-layer modification applied to both vanilla 802.11 and 802.11-Tags consists in doubling the number of MAC-layer retries performed in case of collision, just as in the previous section.

Interpreting the results in Figure 13 requires looking beyond the performance of individual links. In order for the long flow to achieve a transmission rate on the order of tens of kilobytes, all other flows must lower their rates far below the hundreds of kilobytes available on individual links. As shown in the previous section, a fair distribution of rates is associated with a drastic decrease of overall throughput. As expected, this simple MAC layer modification cannot achieve a max-min fair distribution of rates when used in conjunction with a standard AIMD rate controller like TCP at the transport layer. However, as the results in Figure 13 show, the improved scheduler alleviates flow starvation. The long flow in the figure increases its achieved rate from a level that cannot prevent connection timeouts and interruptions to a sustained rate of about 20 kbps.

9. CONCLUSION

This paper has proposed a mechanism for sharing control information in wireless networks, able to function in low SNR conditions and without introducing new interference constraints. It experimentally evaluated its performance using software radio platforms. It was found that this communication scheme is effective at SNR values as low as 0 dB, effectively covering the carrier sense range of a wireless host. Two algorithm implementations that make use of this scheme have been evaluated through simulations, confirming its applicability in protocol design.

10. REFERENCES

- [1] J. Proakis, *Digital Communications*, 4th ed. McGraw-Hill Science/Engineering/Math, 2000.
- [2] H. Robin, D. Bayley, T. Murray, and J. Ralphs, "Multitone signalling system employing quenched resonators for use on noisy radio-telprinter circuits," *Proceedings of the IEE*, vol. 110, no. 9, pp. 1554–1568, 1963.
- [3] M. Simon, J. Omura, R. Scholtz, and B. Levitt, *Spread spectrum communications handbook*. McGraw-Hill New York, 1994.
- [4] A. J. Viterbi, *CDMA: principles of spread spectrum communication*. Addison Wesley Longman Publishing Co., Inc., 1995.
- [5] M. Z. Brodsky and R. T. Morris, "In defense of wireless carrier sense," in *SIGCOMM*, 2009.
- [6] Y. Jian and S. Chen, "Can CSMA/CA networks be made fair?" in *MobiCom*, 2008.
- [7] S. Rangwala, A. Jindal, K.-Y. Jang, K. Psounis, and R. Govindan, "Understanding congestion control in multi-hop wireless mesh networks," in *MobiCom*, 2008.
- [8] M. Chiang, "Balancing transport and physical layers in wireless multihop networks: jointly optimal congestion control and power control," *IEEE JSAC*, vol. 23, no. 1, pp. 104–116, 2005.
- [9] L. Chen, S. H. Low, M. Chiang, and J. C. Doyle, "Cross-layer congestion control, routing and scheduling design in ad hoc wireless networks," in *INFOCOM*, 2006.
- [10] K. Xu, M. Gerla, L. Qi, and Y. Shu, "Enhancing TCP fairness in ad hoc wireless networks using neighborhood red," in *MobiCom*, 2003.
- [11] V. Jacobson, "Congestion avoidance and control," in *SIGCOMM*, 1988.
- [12] A. Warriar, S. Janakiraman, S. Ha, and I. Rhee, "DiffQ: Differential Backlog Congestion Control for Multi-hop Wireless Networks," in *INFOCOM*, 2009.
- [13] L. Tassiulas and A. Ephremides, "Stability properties of constrained queueing systems and scheduling policies for maximum throughput in multihop radio networks," *IEEE Transactions on Automatic Control*, vol. 37, no. 12, pp. 1936–1948, 1992.
- [14] K. Tan, F. Jiang, Q. Zhang, and X. Shen, "Congestion control in multihop wireless networks," *IEEE Transactions on Vehicular Technology*, vol. 56, no. 2, p. 863, 2007.
- [15] D. Katabi, M. Handley, and C. Rohrs, "Congestion control for high bandwidth-delay product networks," in *SIGCOMM*, 2002.
- [16] N. Dukkipati, M. Kobayashi, R. Zhang-Shen, and N. McKeown, "Processor sharing flows in the internet," in *Thirteenth International Workshop*

- on *Quality of Service (IWQoS)*. Springer, 2005.
- [17] A. L. Stolyar, "Maximizing queueing network utility subject to stability: Greedy primal-dual algorithm," *Queueing Syst. Theory Appl.*, vol. 50, no. 4, pp. 401–457, 2005.
- [18] X. Lin and N. Shroff, "Joint rate control and scheduling in multihop wireless networks," in *Decision and Control, 2004. CDC. 43rd IEEE Conference on*, vol. 2, 2004, pp. 1484–1489 vol.2.
- [19] K. Jain, J. Padhye, V. N. Padmanabhan, and L. Qiu, "Impact of interference on multi-hop wireless network performance," in *MobiCom*, 2003.
- [20] V. S. A. Kumar, M. V. Marathe, S. Parthasarathy, and A. Srinivasan, "Algorithmic aspects of capacity in wireless networks," *SIGMETRICS Perform. Eval. Rev.*, vol. 33, no. 1, pp. 133–144, 2005.
- [21] C. Joo, X. Lin, and N. B. Shroff, "Understanding the capacity region of the greedy maximal scheduling algorithm in multihop wireless networks," *IEEE/ACM Transactions on Networking*, vol. 17, no. 4, pp. 1132–1145, 2009.
- [22] S. Vasudevan, D. Towsley, D. Goeckel, and R. Khalili, "Neighbor discovery in wireless networks and the coupon collector's problem," in *MobiCom*, 2009.
- [23] S. Gollakota and D. Katabi, "Zigzag decoding: combating hidden terminals in wireless networks," in *SIGCOMM*, 2008.
- [24] D. Halperin, T. Anderson, and D. Wetherall, "Taking the sting out of carrier sense: interference cancellation for wireless LANs," in *MobiCom*, 2008.
- [25] R. Chandra, R. Mahajan, T. Moscibroda, R. Raghavendra, and P. Bahl, "A case for adapting channel width in wireless networks," in *SIGCOMM*, 2008.
- [26] A. Dutta, D. Saha, D. Grunwald, and D. Sicker, "SMACK: a SMart ACKnowledgment scheme for broadcast messages in wireless networks," in *SIGCOMM*, 2009.
- [27] E. Berlekamp, *Algebraic coding theory*. McGraw-Hill, 1968.
- [28] F. MacWilliams and N. Sloane, *The theory of error-correcting codes*. North-Holland Amsterdam, 1978.
- [29] I. Telatar and D. Tse, "Capacity and mutual information of wideband multipath fading channels," *IEEE Transactions on Information Theory*, vol. 46, no. 4, pp. 1384–1400, 2000.
- [30] G. Li and G. Stuber, *Orthogonal frequency division multiplexing for wireless communications*. Springer, 2006.
- [31] C. Luo and M. Medard, "Frequency-shift keying for ultrawideband - achieving rates of the order of capacity," in *In 40th Annual Allerton Conference on Comm., Control, and Computing*, 2002.
- [32] N. J. A. Sloane and J. J. Seidel, "A New Family of Nonlinear Codes Obtained from Conference Matrices," *New York Academy Sciences Annals*, vol. 175, pp. 363–365, 1970.
- [33] R. N. McDonough and A. D. Whalen, *Detection of Signals in Noise*. Academic Press, Inc., 1995.
- [34] Ettus Research, "USRP board," <http://www.ettus.com/>.
- [35] T. Schmidl and D. Cox, "Robust frequency and timing synchronization for OFDM," *IEEE Transactions on Communications*, vol. 45, no. 12, pp. 1613–1621, 1997.
- [36] F. Tufvesson, O. Edfors, and M. Faulkner, "Time and frequency synchronization for OFDM using PN-sequence preambles," in *IEEE Vehicular Technology Conference, VTC 1999 - Fall*, vol. 4, 1999, pp. 2203–2207.
- [37] T. Paul and T. Ogunfunmi, "Wireless LAN Comes of Age: Understanding the IEEE 802.11n Amendment," *IEEE Circuits and Systems Magazine*, vol. 8, no. 1, pp. 28–54, 2008.
- [38] A. Vasan, R. Ramjee, and T. Woo, "ECHOS-enhanced capacity 802.11 hotspots," in *INFOCOM*, 2005.
- [39] D. X. Wei, C. Jin, S. H. Low, and S. Hegde, "FAST TCP: motivation, architecture, algorithms, performance," *IEEE/ACM Transactions on Networking*, vol. 14, no. 6, pp. 1246–1259, 2006.
- [40] L. Brakmo and L. Peterson, "TCP Vegas: end to end congestion avoidance on a global Internet," *IEEE JSAC*, vol. 13, no. 8, pp. 1465–1480, 1995.
- [41] S. Floyd, "TCP and explicit congestion notification," *SIGCOMM Comput. Commun. Rev.*, vol. 24, no. 5, pp. 8–23, 1994.
- [42] B. Radunović, C. Gkantsidis, D. Gunawardena, and P. Key, "Horizon: balancing TCP over multiple paths in wireless mesh networks," in *MobiCom*, 2008.
- [43] Scalable Network Technologies, "Qualnet simulator: <http://www.qualnet.com/>."
- [44] P. Gupta and P. Kumar, "The capacity of wireless networks," *IEEE Transactions on Information Theory*, vol. 46, no. 2, pp. 388–404, 2000.
- [45] J. Li, C. Blake, D. S. De Couto, H. I. Lee, and R. Morris, "Capacity of ad hoc wireless networks," in *MobiCom*, 2001.
- [46] A. Jindal and K. Psounis, "Characterizing the Achievable Rate Region of Wireless Multi-hop Networks with 802.11 Scheduling," *IEEE/ACM Transactions on Networking*, 2009.