End-to-end performance in vehicular networks with an emphasis on safety and security applications

Final Report

Metrans Project 10-05

July 2012

Konstantinos Psounis

Ming Hsieg Department of Electrical Engineering

University of Southern California

Los Angeles, 90089, CA

Abstract

In wireless networks, the presence of interference among wireless links introduces dependencies among flows that do not share a single 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 interference range of each other. Specifically, control plane information needs to reach nearby nodes which often lie outside the communication range, but within the interference 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 reliable control data transmission at SNR values as low as 0 dB. It does this by employing a number of signal encoding techniques including adding redundancy to multitone modulation, shaping the spectrum to reduce inter-carrier interference, and the use of algebraic coding. Making use of a detection theory-based model we analyze the performance achievable by our modulation as well as the trade-off between the rate of the information transmitted and the likelihood of error. Using real-world experiments on an OFDM system built with software radios, we show that we can transmit data at the target SNR value of 0 dB 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 in which control-plane information must travel beyond the communication range of a node.

Contents

List of Figures

Disclosure: Project was funded in entirety under this contract to California Department of Transportation.

Final Report

End-to-end performance in vehicular networks with an emphasis on safety and security applications

1 Introduction

Vehicular networks (vanets) pose a number of challenges similar to the ones found in the design of ad hoc wireless networks: they require dedicated routing, neighbor discovery, transport and congestion control protocol and, all too often, significantly complicate the afore mentioned problems due to the intermittent and rapidly changing nature of their connectivity. The same challenges that intervene in the design of wireless protocols for distributed multi-hop networks are compounded in this more ephemeral environment: hidden terminal problems, the discrepancy between transmission range and interference range, the difficulty in finding a distributed scheduling protocol capable of approaching the theoretical limits of the capacity region of such a wireless network.

A number of solutions to these problems, in the form of effective distributed algorithms, have been proposed. However, in passing from a theoretical setting to a practical setting, many of these algorithms suffer from one of the above problems, the significant difference between the range at which information can be transmitted and the larger range at which a host interferes with other hosts. For many of these distributed algorithms, which use message passing as a primary tool of coordination, this problem is the main impediment toward deployability. However, most algorithms have low messaging overhead, compared to the amount of data transmitted as payload between the hosts, which prompts the question whether a trade-off between message transmission range and overhead is possible and whether the terms of this exchange are advantageous to the protocol designer. We seek to answer this question and provide an affirmative answer: at the physical layer, communication schemes which provide robust message passing in a setting as challenging as the one of vehicular networks are implementable with an added cost that is not significant.

The range of applications that our improved message passing scheme would support includes, but is not limited to:

1. Pre-crash warning: This application refers to a situation where a number of vehicles communicate to each other to warn their drivers that there is high possibility of a collision.

2. Post-crash warning: Once an accident occurs, approaching vehicles should be warned to prevent subsequent accidents, inform drivers to use alternate routes, etc.

3. Lane change warning: In this application two vehicles need to coordinate to inform their drivers about the high possibility of a collision due to a lane change. This could be either due to one car being on the blind spot of the other, or due to both cars merging to the same lane starting from dierent lanes.

4. SOS services: In this application a vehicle is periodically broadcasting a SOS message. Vehicles that come within range pick up the signal and forward it to other vehicles until one of them is within range of a roadside station. The message will eventually reach an emergency center, e.g. a police station.

5. Curve speed warning: This is another application that requires cooperation between vehicles and roadside stations. The stations monitor the speed of the cars and inform the drivers if they are approaching the curve with too much speed.

In a broader view, our scheme would also support routing and congestion control algorithms proposed for wireless mesh and vehicular networks, making the deployable in real-world settings.

In the following pages, we describe in detail the research problem that we set to solve, our proposed solution and the applications in which we have tested the impact of using this improved message passing scheme, along with the observed improvements.

In wireless networks, the presence of interference among wireless links introduces dependencies among flows that do not share a single 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 interference range of each other. Specifically, control plane information needs to reach nearby nodes which often lie outside the communication range, but within the interference 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 reliable control data transmission at SNR values as low as 0 dB. It does this by employing a number of signal encoding techniques including adding redundancy to multitone modulation, shaping the spectrum to reduce inter-carrier interference, and the use of algebraic coding. Making use of a detection theory-based model we analyze the performance achievable by our modulation as well as the trade-off between the rate of the information transmitted and the likelihood of error. Using real-world experiments on an OFDM system built with software radios, we show that we can transmit data at the target SNR value of 0 dB 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.

Many of the challenges encountered in the design of vehicular 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 vehicular network is commonly modeled as a set of temporary 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.

In our research 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, and minimizes the disruption of data transmissions due to the interference that it induces.

Our first contribution is the design 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. 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, tags are recognized using a receiver based on spectral analysis.

Our second contriboution is an analysis of the performance of tags at different noise and interference levels and when making different design choices. Starting from detection theory principles, we derive, under a sufficiently general propagation model, the detection likelihood/false alarm likelihood curves at different SNR levels.

Our third 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. The results of our evaluation support the conclusion that communication through tags is effective at SNR[1] values as low as 0

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

dB and is robust to the effects of interference.

Our fourth contribution is the use of Tag Spotting in applications, showcasing the performance improvements brought by the existence of a control plane able to reach all nodes within the interference 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 illustrate the way in which vehicular networks can benefit from the use of tags and give further examples of application tailor to vehicular connectivity that can be implemented using tags.

2 Related Work

Vehicular Networks. Most solutions for vehicular networks rely on the concept of delay-tolerant networking [1–3] which has been a topic of research within the last decade. While the fact that availablity of communication is not presumed makes delay tolerant networking a reasonable approach, our solution enhances this aspect of the problem, making communication possible at larger ranges and creating a more connected network infrastracture, capable of supporting more advanced routing and transport protocols. In this respect our solution provides a new setting in which the ideas of delay-tolerant networking should be tested.

Communication and Detection Theory. Tags employ a multicarrier spread-spectrum modulation. They are clearly related to MC-CDMA [4], however they use a non-coherent modulation and do not use orthogonal codewords. Like multitone FSK [5], tags use a combination of tones in order to transmit information.

The design of the tag detector presented in the following section is based on the detection theory of multipulse signals with constant amplitudes and unknown phases. While the classical detector for such a situation is well-studied and understood (see, for example [6, 7]), it requires a precise estimate of the background noise level in order to set appropriate detection thresholds. Interference from competing packet transmissions will confront tags with different levels of background noise, making a precise and timely estimate impossible. Our detector is independent of the level of background noise, requiring only a base SINR as prerequisite for the accuracy of the detection. In the appendix we apply a theoretical analysis similar to the one of the classical detector in order derive the detection/false alarm trade-off curves of our own detector.

Physical Layer Extensions. In the wireless networking world, carrier sense [8] can be seen as an example of a message passing mechanism operating beyond the data transmission range. Closely related is the use of dual busy tones [9] in order to signal channel occupancy. A recently proposed physical layer extension, CSMA with collision notification CSMA/CN [10], aims at reducing the impact of collisions through an early termination signal sent by the receiver of the colliding packet. The transmitter-based detector uses self-interference cancellation techniques in order to improve the SNR of the reciprocal channel and detects the termination signal using correlation, in a manner similar to [11]. However, as the authors of both these papers find out, a correlation based receiver cannot function without prior channel and frequency offset estimation, which prevents their use for broadcasts over arbitrary channels, as in the case of tags.

Carrier sense, dual busy tones and collision notifications are binary signaling mechanisms, not suited for transmitting numeric information, as required by message passing protocols. Another recent physical layer extension [12] aims at realizing a side-channel over spread-spectrum based protocols through perturbations of certain chips comprising a transmitted symbol.

The technology of software defined radios [13, 14] 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 [15], interference cancellation [16] or dynamic bandwidth adaptation [17]. In particular, [15] offers a different approach to the hidden terminal problem, using signal processing techniques in order to separate colliding transmissions. Perhaps the most similar technique to the one presented in this paper is the one of smart broadcast acknowledgments, introduced in [18], in which multitone modulation is used for the purpose of simultaneously conveying positive acknowledgments from multiple receivers. Optimal spectrum usage in the absence of sender synchronization is the topic of [19]. Spatial diversity is used in [20] in order to achieve interference alignment and cancellation.

Congestion Control and Scheduling. 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 [21, 22] in the tradition of TCP congestion control [23], or by regulating transmissions based on queue differentials [24] along the lines of back-pressure ideas [25].

In an effort to tackle the complexity of creating optimal schedulers, recent work on medium access for multihop networks has proposed distributed algorithms capable of approximating the optimal solution [26] [27] [28] [29] [30]. A common theme here is the use of local, neighborhood-centered information in achieving a global solution. The mechanism proposed in this paper offers an efficient way to implement the neighborhood-wide sharing of control information in the schemes mentioned.

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. The information sharing mechanism proposed in this paper eases the implementation of many of these ideas and improves their performance as control information will reach nodes outside the data transmission but within the carrier sense range in both mesh networks and vehicular networks.

3 Tag Spotting

High Level Design Overview. Tag Spotting uses a set of signals (tags) that are easily detectable in low SNR conditions. The design of tags is determined by a number of constraints. Firstly, due to their short timespan, the presence of multiple indistinguishable sources and the presence of varying levels of interference, the tag detector cannot perform accurate channel estimation, or achieve time, phase and fine frequency synchronization. Secondly, in order to protect competing data transmissions from further levels of interference, tags must abide by a maximal spectral power constraint, which prevents the use of a peaky transmission scheme such as multiple frequency shift keying (MFSK). Finally, a tag detector must perform identically in the presence of added interference, as long as the SINR remains unchanged.

To address these constraints, we employed a noncoherent communication scheme that spreads a tag's energy over the frequency domain. The number of carriers was increased in order to sharpen the signal's spectral footprint and ease spectral analysis through discrete Fourier transform even in the absence of complete frequency synchronization. The symbol sequences encoded over the different carriers were chosen according to an algebraic code, thus adding extra redundancy.

The current section presents the details of tag construction. The following section will present the reasoning behind the design decisions made throughout the construction.

Multitone Structure. Figure 3.1 presents a packet transmission in time domain representation as well as in two different time-frequency representations. The two time-frequency representations, pictured in the lower two panes, are realized by performing the spectral analysis of successive blocks of 512 samples (middle pane) or 64 samples (lower pane). As we will see, these lengths are natural choices for describing the structures of tags and packet payloads respectively. In the same representations, each vertical column corresponds to an analyzed block, while the horizontal line pattern present in each such column illustrates the block's power spectral decomposition. In order to make the representation more meaningful, we have stripped tags and data frames of their cyclic prefixes and we aligned the boundaries of the analysis periods with the boundaries of tags and data frames.

On the receiver side, the tag detector operates in the presence of small, tolerable but unknown frequency offsets, which may cause power spillage from the active wide carriers to the inactive carriers. As illustrated in Figure 3.2, we chose to send the entire signal power allotted to an active wide carrier using the central four of the eight thin carriers corresponding to this particular wide carrier. We also chose to encode the tones sent on these thin carriers using different random phases. The above choices reduce the amount of power leaked onto inactive wide carriers. A straightforward computation assuming a frequency offset distributed uniformly between zero and two thin carrier widths reveals that the expected power leaked is, in expectation, about 2.3% of the total power. Figure 3.3 further 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. By comparing the distribution of the received signal power in the frequency bins corresponding to wide carriers (middle pane) with the structure of the transmitted tag (lower pane), it can be seen that the received power is concentrated in those bins which correspond to active wide carriers. As a note, 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 their peak to average power ratio.

Constructing a Tag Detector. Let $T = \{t_1, t_2, \ldots, t_{60}\}$ denote the set of all tags and C_i be the set of all data 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

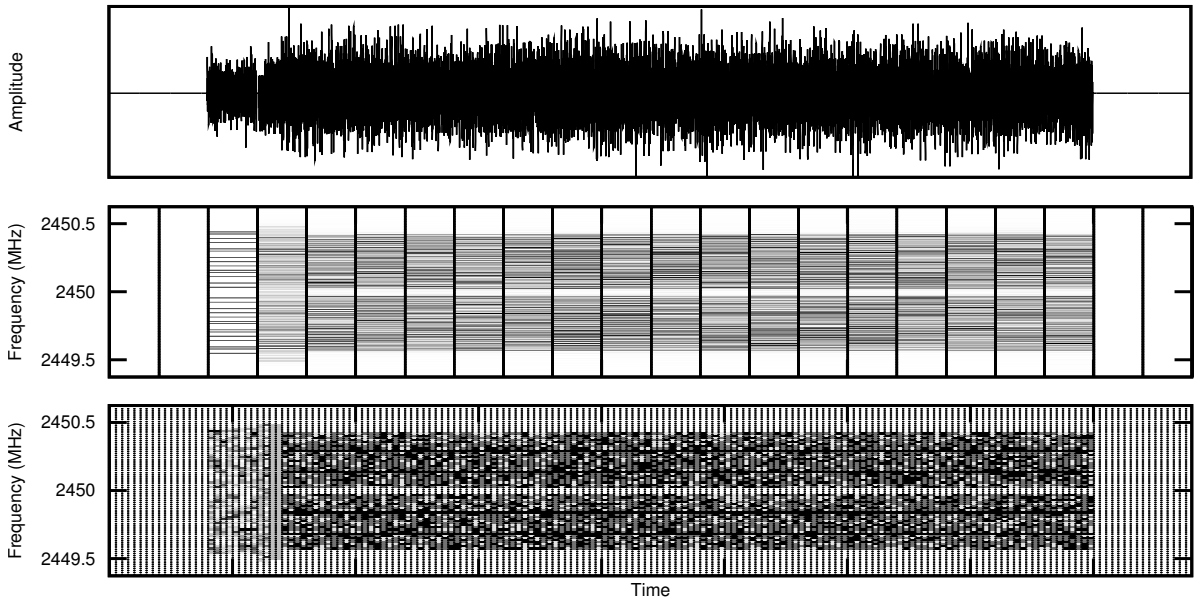

Figure 3.1: Packet transmission: the signal $s(t)$ (upper pane), a time frequency representation using 512 frequency bins (middle pane) and a time frequency representation using 64 frequency bins (lower pane).

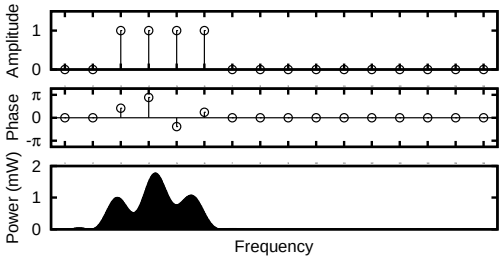

Figure 3.2: The discrete spectrum of a two-carrier group (amplitude and phase representation) and its continuous power spectrum.

we will name from now on tag strength:

$$\frac{\sum_{f \in C_i} r_f^2}{\sum_{\forall f} r_f^2}. \tag{3.1}$$

Tag strength is compared against a fixed threshold γ and in case the threshold is exceeded a possible tag observation is recorded. [1]

Detection intervals have the same length as a tag from which the cyclic prefix has been removed and are spaced one tag cyclic prefix length apart. It results that successive analyzed intervals have significant overlap. Every transmitted tag will completely cover at least one detection interval. The detector processes every interval by first computing the Fourier transform of the contained signal and then computing, based on the

[1] This equation is similar to the one of the low-SNR multipulse detector with n samples for a signal with unknown phase $\theta(t)$ varying at each pulse: $s(t) = A * \cos(2\pi f t + \theta(t))$, 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 [6, 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.

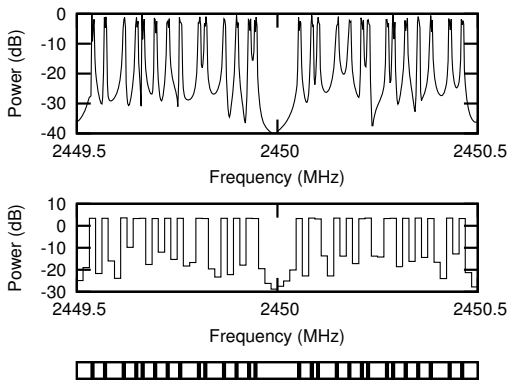

Figure 3.3: The spectrum of a received tag in the presence of a frequency offset. Upper pane: the 512 frequency bins Fourier transform of the corresponding detector interval. Middle pane: the 64 frequency bins representation used in the detection decision. Lower pane: the structure of the transmitted tag.

resulting spectrum, the strength of each tag according to Equation (3.1). In order for a tag recognition event to be recorded, the corresponding tag strength must, firstly, exceed the threshold value and, secondly, be maximal among all tag strengths (for all tags) derived from 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 default parameters, 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. In some control schemes some of the data messages will not require tags to be piggybacked, allowing for a lesser overhead.

4 Motivating the Design Choices

The previous section has presented without goint into too much detail the structure of tags. While the above description is complete, the decisions taken in the construction of tags may well seem arbitrary. The purpose of the current section is to motivate every design decision taken in tag construction.

Multitone Structure

The design space of tag signals is frequency space. At the lowest level, tags use the same form of modulation in frequency space, using orthogonal signals over a finite time interval, that is used in regular data transmission.

OFDM. In an opportunistic reception system that searches for the short occurrences of tags, the price of exact frequency synchronization and of exact channel estimation should be avoided. Tags are therefore transmitted and received without performing channel estimation or frequency synchronization. This raises a challenge in solving inter-carrier interference.

Remember that the choice of OFDM for data transmission is linked to the propagation behavior of wireless signals. Since sine signals are the eigenfunctions of the wireless channel, the use of orthogonal sines along with an appropriate cyclic prefix is meant to prevent any channel-caused interference among the different transmitted symbols. The limited timespan of the transmission interval precludes the use of actual sines or any other signals with narrow support in the frequency space. The actual signals used in the transmissions have a slowly-decaying spectral footprint. In OFDM data transmission, the inter-carrier interference which the slow spectral decay entails is avoided through exact frequency synchronization, using the fact that in frequency space the zeros of the base sine-like signals align with the peaks of all other signals in the base set. In contrast, for tags, the packing of noncoherent carriers into larger building blocks reduces the amount of power leaked among the frequency bins corresponding to thick carriers. This reduction in leaked power allows the system to function as intended even when the receiver is not frequency-locked onto the transmitter. The lack of frequency synchronization together with the lack of an estimate of the channel phase response at different frequencies also leads to uncertainty regarding the phase of the transmitted signal. The use of a noncoherent encoding allows us to overcome the lack of knowledge of the channel phase response without further complications. It could be argued that these channel characteristics should be measured in advance. However, our receptions are at best opportunistic, and the channel could be any one of a multitude of fast changing channels between any pair of hosts. Certainly, obtaining an exact estimate of the channel response and of the frequency offset, at the low SNR levels for which our system is designed, would greatly complicate the tag transmission problem.

Fading. Another characteristic of wireless channel transmission, frequency-selective fading, provides the rationale for the use of groups of two carriers as a encoding unit: since neighboring data carriers are likely to experience similar fading and since any of the codewords makes use of either one or the other of the two carriers in a two-carrier group in order to encode a bit value, fading over a two-carrier group will not induce a bias towards any of the hypotheses that a particular codeword has been transmitted. The received power and the likelihood of detection may well decrease due to fading. However, when considering a given overall signal to noise ratio, i.e. computed over all the carriers, the detection probabilities over fading and non-fading channels are quite similar, while the false alarm probabilities are the same. We can conclude that this particular design decision manages to overcome most of the difficulties that fading introduces in tag detection.

Constructing a Tag Detector

It is worth mentioning here a significant difference between the main purpose of a tag receiver and the purpose of a communication system receiver. While a communication system receiver is meant to accurately distinguish between a number of hypotheses corresponding to signal transmissions under the assumption that an actual transmission has occurred, the tag receiver listens for the most part to noise and background chatter. The main task of a tag receiver is therefore to detect, with sufficient confidence, a tag transmission when one occurs and, if possible, to correctly identify the transmitted tag. Tag detection is therefore a detection problem more than a communication problem and the design of the detector reflects this fact. The probabilities of false alarm allowed in the case of tags are well under the typical probabilities of misclassification allowed in a communication system, since the occurrence of tag transmissions are assumed a priori to be rather rare events. False alarms weigh in more heavily when compared to the total number of detected tags. Due to the fact that tag detection is essentially a detection theory problem, we choose the detection metric (probability of false alarm versus the probability of detection) to be the main measure of tags performance. The secondary metric considered will be the probability of misclassification of a transmitted tag. The experimental section will reveal that this probability is negligible due to the high threshold required for a positive tag detection, even when using a rather large number of codewords.

Detecting Patterns. In general, tag observations occur over short intervals of time and channel conditions change too frequently for the receiver to obtain and update an accurate noise and interference power estimate. The only assumption made in the following is that the spectral envelope of the noise and interference signals is flat, an assumption that can be justified in the case of data networks using OFDM-based encoding. We design therefore our modulation scheme and our receiver to use as a detection indication not the sheer amount of power received but rather the concentration of the received power into pre-determined frequency bins. The receiver detects a transmission event whenever the concentration of the received power (the ratio of the power received in the designed frequency bins to the overall received power) exceeds a certain threshold. Therefore, the receiver searches not just for the presence of a signal but for a certain spectral shape. The fact that a power ratio measurement is used as a detection metric guarantees an universal receiver in a wider sense: the probability of detection for a threshold value chosen as the receiver parameter will only increase with increasing SINR. A standard detector is denoted as universal when a similar guarantee exists in terms of the SNR.

Tags and FSK. The attentive reader might have noticed that a simpler encoding scheme might have provided a similar detection/false alarm performance trade-off without the use of an algebraic code. Frequency shift keying simply concentrates the available power into the frequency space corresponding to one of the available carriers, thus offering similar received power characteristics. However, FSK has a large power spectral density, due to the fact that all the transmitted power is effectively concentrated in one point of the frequency spectrum, which makes it undesirable in a network environment, where we would like to guarantee a certain flat envelope for the frequency spectrum of our transmission, with a fast decay outside the data band. Our choice of modulation limits the transmitted power at any given frequency, resulting thus in a flat spectrum, similar to the one corresponding to OFDM data transmissions. The experiments verify that the typical interference effects of tags on competing data transmissions are not worse than the interference effects caused by normal packet data transmissions sent at a similar overall power level, which would not be the case if tags were modulated using FSK.

5 Evaluation

Physical Layer Evaluation 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. In order to determine the likelihood 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. We have also evaluated through simulations the likelihood of misclassifications.

We used the following metrics throughout the evaluation:

Tag Strength is the quantity defined in Equation (3.1), 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 5.1d) and it is low enough to allow detection at the target SNR values. Using the naïve assumption that noise (and interference) power contribute 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.

Probability of Misclassification (P_m) is the likelihood of incorrect tag identification in the presence of a tag transmission.

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

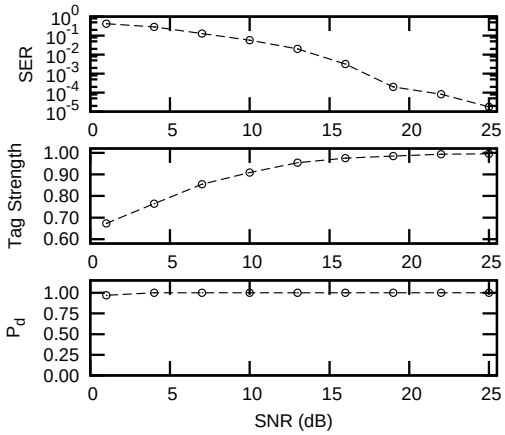

(a) Experimental results in the presence of channel noise.

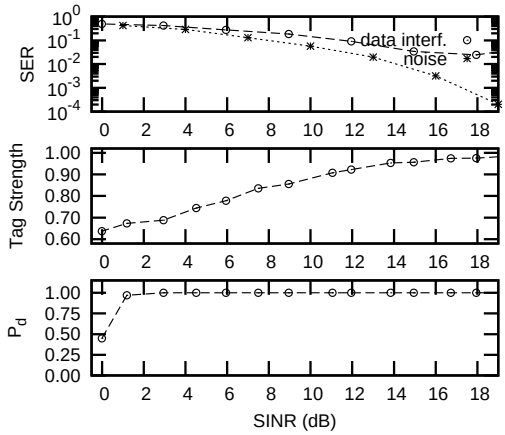

(b) Experimental results in the presence of data-like interference.

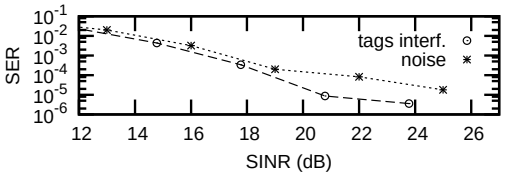

(c) Experimental results in the presence of tag-like interference.

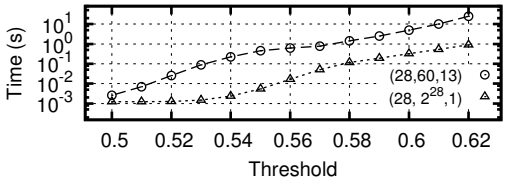

(d) Average time between false alarms at different thresholds.

Figure 5.1: Experimental Results.

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 3.3 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 5.1a 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 a 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 5.1b 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 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 5.1c presents the effect of tags on data transmissions. The SER curve is very close to the SER curve of Figure 5.1a, 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 5.1d presents the dependence of the average time 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 every 20 seconds.

6 Theoretical Performance Analysis

We have described the design of tags and introduced an universal detector[1] that does not require an estimate of the combined power of noise and background interference. However, the classic theoretical results on the detection and false alarm probability distributions are not readily applicable to our more complicated tag detector. In the following we will analyze, using rather conservative fading models, the performance achievable by a few particular tag families. At first we will consider the performance achievable when searching for a single tag signal, after which we will generalize to larger families of tags. The analysis will make use of two fading models, a narrowband fading model and a wideband fading model which assumes Rayleigh propagation. In the first model, which tries to approximate narrowband transmission, the fading across all carriers is identical, i.e. in the absence of noise the receiver would receive the same amount of power in the bandwidth intervals corresponding to different carriers. In the second model, approximating wideband transmission, fading will be modeled independently for every carrier using a Rayleigh propagation model. Due to the limited timespan of tags, these two models describe short-term fading effects only.

Single Tag. Let us consider a tag t. In what follows, we will denote by c the number of two-carrier groups used in the tag's construction. Assume that, for each thin carrier, the receiver noise has power n and its distribution can be modeled by a complex Gaussian random variable $\mathcal{N}(0, n)$. We further assume that for the active thin carriers the average signal power in the frequency band corresponding to any carrier is p. Under the assumption of narrowband transmission and without including the receiver noise contribution, the received signal obtained after demodulating one of the carriers can be modeled as a circularly uniform complex variable of constant amplitude $\sqrt{p}$, while in the case of wideband transmission the signal is modeled by a complex Gaussian random variable $\mathcal{N}(0, p)$. Assume that each thick carrier is composed of α thin carriers, out of which β are active. In the case of the system presented in the previous section $\alpha = 8$ and $\beta = 4$. Let P_t denote the set of thin carriers activated when tag t is transmitted, Q_t represent the thin carriers that, while they are not activated, belong to active thick carriers and R_t the thin carriers that belong to unactivated thick carriers. In accordance to the definition given when introducing the tag detector, we denote through $C_t = P_t \cup Q_t$ the set of all thin carriers belonging to activated thick carriers, regardless of whether they are active or not. Let r_f be the complex values obtained after computing a fast Fourier transform of the real and complex components of a sampled tag signal, i.e. the received signal values corresponding to the various thin carriers. Let $||\cdot||$ denote the l^2 norm.

Under the assumption of independent Rayleigh fading, it results that for each active thin carrier the amplitude of the received signal is distributed according to a complex Gaussian random variable $\mathcal{N}(0, p+n)$. Choosing a threshold value γ for the quantity defined in Equation 3.1, we can write the probability of detection as:

$$P_d = P\left(\frac{\sum_{f\in C_t} ||r_f||^2}{\sum_{f\notin C_t} ||r_f||^2} > \frac{\gamma}{1-\gamma}\right) \tag{6.1}$$

or

$$P_d = P\left(\frac{\sum_{f\in P_t} ||r_f||^2 + \sum_{f\in Q_t} ||r_f||^2}{\sum_{f\in R_t} ||r_f||^2} > \frac{\gamma}{1-\gamma}\right) \tag{6.2}$$

[1] a detector for which, for any chosen probability of error P_f, the corresponding probability of detection P_d can only increase when the SINR is increased. A universal detector is particularly suited to our purposes, given the unpredictable nature of interference power.

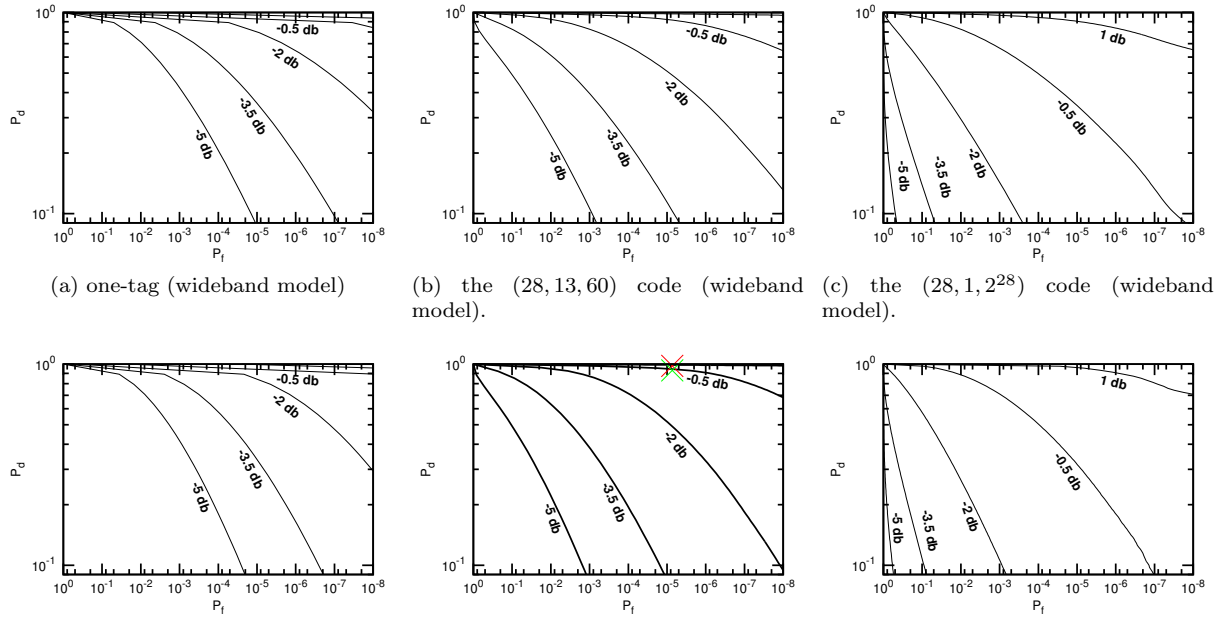

(a) one-tag (wideband model)

(b) the $(28, 13, 60)$ code (wideband model).

(c) the $(28, 1, 2^{28})$ code (wideband model).

(d) one-tag (narrowband model)

(e) the $(28, 13, 60)$ code (narrowband model).

(f) the $(28, 1, 2^{28})$ code (narrowband model).

Figure 6.1: Detection curves for different choices of code and propagation model.

It results from the previous paragraph and due to the independence of the circular random variables considered that the sums in the above equations can be written using the Chi-Square distribution with d components, denoted through χ^2_d. Namely $\frac{1}{p+n}\sum_{f\in P_t}||r_f||^2 \sim \chi^2_{2\beta c}$, $\frac{1}{n}\sum_{f\in Q_t}||r_f||^2 \sim \chi^2_{2(\alpha-\beta)c}$ and $\frac{1}{n}\sum_{f\in R_t}||r_f||^2 \sim \chi^2_{2\alpha c}$.

In the case of narrowband fading, the amplitude of the received signal for each active carrier is constant. Therefore the first of these sums can be written using the noncentral Chi-Square distribution with parameter $\lambda = 2\beta c\frac{p}{n}$. We write $\frac{1}{n}\sum_{f\in P_t}||r_f||^2 \sim \chi^2_{2\beta c}\left(2\beta c\frac{p}{n}\right)$.

The probability of false alarm in the case of a single tag (and a single tested hypothesis) can be obtained by setting $p = 0$ in the above formulas. Therefore, writing the ratio of the two Chi-Squared random variables using a random variable f that follows the Fisher-Snedcor F-distribution [31], $f \sim \mathcal{F}(2\alpha c, 2\alpha c)$,

$$P_f(t) = P\left(f > \frac{\gamma}{1-\gamma}\right) \tag{6.3}$$

Figures 6.1a and 6.1d present the detector's behavior at different SNR values in the case of wideband and narrowband signals, respectively. [2]

Choosing the number of active carriers. Consider in the following a problem mentioned previously, namely estimating the optimal number of thick carriers q that should be activated during a tag transmission in order to maximize the performance of the detector. Let N denote the total number of thick carriers. In order to obtain closed-form results, we use a simplified model of tags in which we set $\alpha = \beta$, that is active thick carriers will use all thin subcarriers for transmission. Let f' be a random variable generated using the corresponding Fisher-Snedcor distribution, $f' \sim \mathcal{F}(2\alpha q, 2\alpha(N-q))$. Under this assumption we can simplify the formulas for the probability of detection and false alarm in the wideband case to: $P'_d(t) = P\left(\left(1+\frac{p}{n}\right)f' > \frac{\gamma}{1-\gamma}\right)$ and $P'_f(t) = P\left(f' > \frac{\gamma}{1-\gamma}\right)$

[2] The SNR figures are computed using the power and noise figures for a thick carrier, that is $SNR = \frac{\beta p}{\alpha n}$.

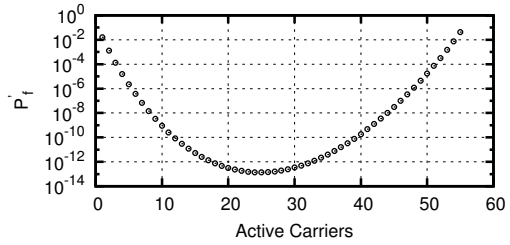

Figure 6.2: The probability of typical false alarms as a function of the number of active carriers. The total number of carriers used in tag construction is 56.

It results that the probability distribution of the receiver response corresponding to detections is just a scaled version of the probability distribution corresponding to false alarms. We introduce a new performance measure in order to characterize the performance change due to the choice of q. Let γ_0 be the value of γ for which $P_d^{'} = \frac{1}{2}$, at a SNR value of 0 dB. Figure 6.2 plots the behavior of $P_f^{'}$, for a detector with a threshold γ_0, for different values of q, when $N = 56$. The quantity plotted represents the value of the tail of the probability of false alarm in the typical detection region.

Families of Tags. The next point in our analysis will be considering the situation in which the detector searches for multiple hypotheses. Assume therefore that the tag t is a member of a family of tags T. It can be readily observed that both in the narrowband and wideband cases the probabilities of detection remain unchanged. The probability of false alarm can be rewritten as:

$$P_f(T) = P\left(\max_{t\in T}\left(\frac{\sum_{f\in C_t}||r_f||^2}{\sum_{f\notin C_t}||r_f||^2}\right) > \frac{\gamma}{1-\gamma}\right)$$

where $||r_f||^2 \sim \chi^2_2$.

Finding a proper bound on the maximum that would take into account the number of codewords chosen and the geometry of the codeword space is a non-trivial task. We had therefore to recur, for the time being, to Monte Carlo simulations in order to approximate this distribution.

Figures 6.1b and 6.1e present the detector's behavior, evaluated through Monte Carlo simulation, at different SNRs in the case of wideband and narrowband signals, respectively for the $(28, 13, 60)$ code mentioned previously . In particular, Figure 6.1e compares the theoretical predictions with the experimental results presented in the evaluation. The green cross at the top of the figure presents the experimentally measured detection probability for tags in the presence of background noise only, at an SNR value of 1 dB, as shown in in Figure 5.1a, while the red cross presents the theoretical value at the same SNR. The experimental data and the theoretical curve, which indicates a probability of detection nearing one, are in agreement. The model we are using assumes noise to be white and cannot accurately predict the performance of the detector in the presence of data-like interference. Figures 6.1c and 6.1f illustrate the same detection curves in the case of the simple $(28, 1, 2^{28})$ code.

The power levels on the carriers are being summed up in the hypotheses in different ways. Since the variables that are summed up are chosen from the same set, the probability that the maximum present in the function will cross any chosen threshold γ' is significantly lower than what the sum bound on the individual probabilities of errors associated with the different codewords would predict.

The simpler quantity $\max_{t\in T}\sum_{f\in C_t}||r_f||^2$ can be bounded using an initial symmetrization step [32] and deriving, using generic chaining [33], a Dudley-like inequality [34] on the probability that the maximum exceeds any given threshold. The resulting bound limits the increase of the necessary threshold, for any fixed probability of false alarm, to a quantity of the form $O(log(N))$ where N is the number of codewords used. For reasons of space we have not included the derivation of the bound.

For the code construction using the two-carrier groups presented previously , a simple upper bound on the probability of false alarm can be derived by considering the case of the simplest $(28, 1, 2^{28})$ code, which has the largest probability of false alarm of all possible codes since it includes all possible codewords .

Consider a set of pairs of Chi-Square distributed random variables $(x_{i,1}, x_{i,2})$ and let x_i^M and x_i^m represent the maximum and minimum, respectively, in each pair. The probability of false alarm can be thus written, for the afore mentioned code, as:

$$P_f(T) = P\left(\frac{\sum_{\forall i} x_i^M}{\sum_{\forall i} x_i^m} > \frac{\gamma}{1-\gamma}\right)$$

The formula above has been used in order to derive the detection curves for the code mentioned, which are presented in Figures 6.1c and 6.1f.

7 Applications: congestion control and scheduling for vanets

Congestion Control

In vehicular networks the same problems that plague the advance of wireless mesh networks, namely the difficulty in finding viable congestion control and scheduling schemes are amplified by the temporary nature of the wireless topology. While for most application reaching the capacity region is not a concern, as their data usage needs are far below the typical channel capacity of any particular connection, avoiding the negative effects of interference remains a primary concern. Our previous work in this field, the congestion control protocol WCP is readily adaptable to vehicular networks. In the following we will argue that tags improve significantly the performance of WCP and make it a viable option for deployment in realistic wireless scenarios.

We would like to emphasize that attaining a fair distribution of resources when using a wireless medium is, in our opinion, a neighborhood-centric problem. 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. At the medium access layer, all hosts must be able to gain access to the wireless channel. At the transport layer, all flows must be able to attain a sustained transfer rate comparable to the transfer rates of flows with which they compete, i.e flows that share the same wireless neighborhoods.

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 [23], increases in queue wait times [35] [36] or explicit congestion notifications [37].

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 [22] and [38], uses broadcasts of local information originating at every one of the network hosts. In these works, such information is piggybacked onto data/ack transmissions and broadcasted only to nodes within the communication range. Tag Spotting is capable of passing this type of control information to all nodes within the interference range.

With this in mind, we extend WCP [22], 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 by each node in data and ack packets: a congestion bit that indicates congestion events to its neighbors and the maximum of the round trip times (RTTs) of flows traversing it, a metric used in achieving a max-min fair allocation. 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 Both loss rates and delays experienced by competing flows passing through a congested neighborhood may vary widely, significantly more than in wired networks. This causes the senders' AIMD controllers to increase their rates at significantly different paces following a congestion event, unless a common loop duration is used. For more details on WCP, the interested reader is referred to [22].

We have simulated the performance of WCP using the Qualnet network simulator [39]. We have extended Qualnet's physical layer simulation in order to also handle the likelihood of tag detection using the detection probabilities measured in the evaluation part . The content of tags is composed of one congestion bit and

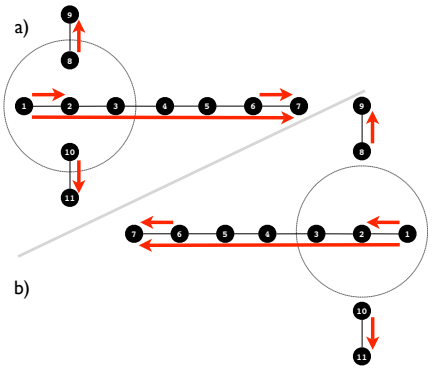

Figure 7.1: Chain-cross topology with all competing flows separated by at most one transmission range (upper) and with some competing flows separated by more than the transmission range (lower).

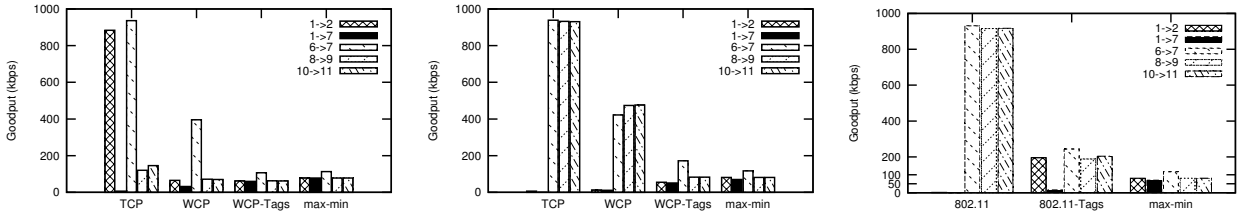

(a) for the topology in Figure 7.1a.

(b) for the topology in Figure 7.1b.

(c) under the modified scheduling policy for the topology in Figure 7.1b.

Figure 7.2: Goodput Results

a field that encodes, on a logarithmic scale with base $\sqrt[3]{2}$, the value, in milliseconds, 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 and increase the likelihood of tag reception.

Figure 7.1a 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 7.2a 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, above all, the fairness achieved and secondly the throughput. 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 [40]. Any increase in the rate of the long flow in the figure will involve a drastic reduction 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 [41] the max-min rate allocation for these flows and compare it to the other allocations. It is evident that both 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 7.1b 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 7.2b. As before, WCP-Tags yields rates which are close to the max-min optimal allocation.

Scheduling

A large class of scheduling algorithms for vehicular networks and wireless multi-hop networks are centered on ideas such as queue equalization using backpressure [26] [42], broadcasting local congestion indications [21] or creating a computationally tractable approximation of an optimal schedule [27] [28]. 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.

Consider again the topology illustrated in Figure 7.1b. The results of the previous section 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 15 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 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 7.2c 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.

One reason for this was discussed in the previous section : AIMD controllers require not only congestion notifications at a neighborhood level, but must also compensate for the different periods of their feedback loops due to significant differences in the round trip times of interfering flows.

As the results in Figure 7.2c show, the improved scheduler alleviates flow starvation. The shortest of the two starved flows in the figure reaches a rate similar to the ones of the other three short flows. 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 15 kbps.

Other Applications

Vehicular networks, as previously mentioned, pose a number of extra problems when compared to multihop wireless networks. In the following we will try to illustrate some of the applications of tags to a large number of problems that appear in vehicular networks. In a sense, we are creating a cookbook of algorithm implementations, using tags as one of the prime ingredients.

Neighbor discovery is a principal problem in realizing a wireless topology. While many of the algorithms present today are only able to pass neighbor identity information withing the data transmission range, a large enough family of tags can be used in order to pass identities within a wider range. Covering the entire interference range allows host to be aware of their complete wireless neighborhood, avoiding thus hidden terminal or blind interference situations.

While realizing an optimal distributed scheduling algorithm which offers a fair allocation of resources is by no means a trivial task, the afore mentioned complete depiction of the wireless neighborhood, coupled with the ability to send control messages that pass significant information, such as queue sizes of wait times, can assist the protocol designer in implementing such an algorithm.

Vehicular security applications, for example collision avoidance, lane change warnings and crash warnings make use of a small set of control messages, which can be easily dispatched using tags, with high reliability and at a very small overhead. Especially in this type of applications, the long range, high robustness characteristics of tags can provide an increase in range and prevent the need for increased transmitter power.

The practical realization of transport protocols suitable for inter-vehicle communication is still being awaited. As our results show, tags are a valuable tool in the realization of such protocols. While WCP focuses on achieving a fair distribution of rates in every wireless neighborhood, other more mundane tasks associated with the transport layer, such as flow control or the transmission of negative acknowledgments and of interference and collision notifications can be left to tags. The low SNR required by tag transmission allows, in combination with self-cancellation techniques, for the realization of partial duplex wireless channels, in which the receiver can provide useful feedback to the transmitter.

8 Implementation

The previous section discussed in detail practical applications of the research findings for congestion control and scheduling in VANETS, as well as for supporting VANET-specific applications like collision avoidance, lane change warnings and crash warnings.

Implementing the research findings in a real, commercial system, is certainly possible and it may benefit the performance and robustness of the overall system. Auto makers are already experimenting with automated applications that will increase vehicular safety. In their systems, they are experimenting with suitable physical layer techniques, MAC layer protocols like the 802.11p standard, as well as routing and transport techniques. They are also implementing applications like crash warnings, and are testing them in experimental setups. The engineers working on these prototypes can use the findings of this research project and implement into their prototype boxes the use of TAGS to improve transport and scheduling performance as discussed in the previous section. They can also implement TAGS and use them to pass application-specific control messages.

Such prototyping work may take place within an industrial research lab setup as mentioned above, or within a university under additional funding from funding sources like METRANS. The PI is already discussing with CISCO Systems, the largest networking company in the world, about the possibility of building a prototype box as a joint venture of USC and CISCO.

Bibliography

[1] K. Fall, "A delay-tolerant network architecture for challenged internets," in *SIGCOMM*.

[2] Delay tolerant networking research group., "http://www.dtnrg.org."

[3] N. Banerjee, M. D. Corner, D. Towsley, and B. N. Levine, "Relays, Base Stations, and Meshes: Enhancing Mobile Networks with Infrastructure," in *Mobicom*, 2008.

[4] N. Yee, J.-P. Linnartz, and G. Fettweis, "Multi-carrier CDMA in indoor wireless radio networks," 1993.

[5] 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.

[6] R. N. McDonough and A. D. Whalen, *Detection of Signals in Noise*. Academic Press, Inc., 1995.

[7] H. L. V. Trees, *Detection, Estimation, and Modulation Theory: Radar-Sonar Signal Processing and Gaussian Signals in Noise*. Melbourne, FL, USA: Krieger Publishing Co., Inc., 1992.

[8] M. Z. Brodsky and R. T. Morris, "In defense of wireless carrier sense," in *SIGCOMM*, 2009.

[9] Z. J. Haas and J. Deng, "Dual busy tone multiple access (dbtma) - a multiple access control scheme for ad hoc networks," *IEEE Transactions on Communications*, 2002.

[10] S. Sen, R. Roy Choudhury, and S. Nelakuditi, "Csma/cn: carrier sense multiple access with collision notification," in *MobiCom*, 2010.

[11] 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.

[12] K. Wu, H. Tan, Y. Liu, J. Zhang, Q. Zhang, and L. Ni, "Side channel: bits over interference," in *MobiCom*, 2010.

[13] Ettus Research, "USRP board," http://www.ettus.com/.

[14] S. Gupta, C. Hunter, P. Murphy, and A. Sabharwal, "Warpnet: Clean slate research on deployed wireless networks," in *Mobihoc*, 2009.

[15] S. Gollakota and D. Katabi, "Zigzag decoding: combating hidden terminals in wireless networks," in *SIGCOMM*, 2008.

[16] D. Halperin, T. Anderson, and D. Wetherall, "Taking the sting out of carrier sense: interference cancellation for wireless LANs," in *MobiCom*, 2008.

[17] R. Chandra, R. Mahajan, T. Moscibroda, R. Raghavendra, and P. Bahl, "A case for adapting channel width in wireless networks," in *SIGCOMM*, 2008.

[18] A. Dutta, D. Saha, D. Grunwald, and D. Sicker, "SMACK: a SMart ACKnowledgment scheme for broadcast messages in wireless networks," in *SIGCOMM*, 2009.

[19] L. Yang, B. Y. Zhao, and H. Zheng, "The spaces between us: setting and maintaining boundaries in wireless spectrum access," in *MobiCom*, 2010.

[20] S. Gollakota, S. D. Perli, and D. Katabi, "Interference alignment and cancellation," in *SIGCOMM*, 2009.

[21] K. Xu, M. Gerla, L. Qi, and Y. Shu, "Enhancing TCP fairness in ad hoc wireless networks using neighborhood red," in *MobiCom*, 2003.

[22] S. Rangwala, A. Jindal, K.-Y. Jang, K. Psounis, and R. Govindan, "Understanding congestion control in multi-hop wireless mesh networks," in *MobiCom*, 2008.

[23] V. Jacobson, "Congestion avoidance and control," in *SIGCOMM*, 1988.

[24] A. Warrier, S. Janakiraman, S. Ha, and I. Rhee, "DiffQ: Differential Backlog Congestion Control for Multi-hop Wireless Networks," in *INFOCOM*, 2009.

[25] 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.

[26] 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.

[27] 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.

[28] K. Jain, J. Padhye, V. N. Padmanabhan, and L. Qiu, "Impact of interference on multi-hop wireless network performance," in *MobiCom*, 2003.

[29] 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.

[30] 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.

[31] N. L. Johnson, K. Samuel, and B. Nikhil, *Continuous Univariate Distributions, Vol. 2.* New York: Wiley, 1995.

[32] M. Ledoux and M. Talagrand, *Probability in Banach Spaces.* Springer, 2006.

[33] M. Talagrand, "Majorizing Measures: The Generic Chaining," *The Annals of Probability*, vol. 24, no. 3, pp. 1049–1103, 1996. [Online]. Available: http://www.jstor.org/stable/2244967

[34] R. M. Dudley, "The sizes of compact subsets of hilbert space and continuity of gaussian processes," *Journal of Functional Analysis*, vol. 1, no. 3, pp. 290 – 330, 1967.

[35] 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.

[36] 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.

[37] S. Floyd, "TCP and explicit congestion notification," *SIGCOMM Comput. Commun. Rev.*, vol. 24, no. 5, pp. 8–23, 1994.

[38] B. Radunović, C. Gkantsidis, D. Gunawardena, and P. Key, "Horizon: balancing TCP over multiple paths in wireless mesh networks," in *MobiCom*, 2008.

[39] Scalable Network Technologies, "Qualnet simulator: http://www.qualnet.com/."

[40] J. Li, C. Blake, D. S. De Couto, H. I. Lee, and R. Morris, "Capacity of ad hoc wireless networks," in *MobiCom*, 2001.

[41] 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.

[42] 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.