Sub-carrier Switch Off in OFDM-Based Wireless Local Area Networks

Thomas Nitsche
IMDEA Networks
Avda. del Mar Mediterraneo 22
28918 Leganes, Spain
Email: thomas.nitsche@imdea.org

Joerg Widmer
IMDEA Networks
Avda. del Mar Mediterraneo 22
28918 Leganes, Spain
Email: joerg.widmer@imdea.org

***Abstract*—OFDM based wireless communication systems split the available frequency band into so-called sub-carriers, and data is transmitted on each of these sub-carriers in parallel. With frequency selective fading, sub-carriers may experience different channel qualities. Thus, choosing a different modulation and coding scheme (MCS) per sub-carrier improves performance. However, this comes at an increase in transceiver complexity and no current wireless system adapts the MCS at such a fine granularity. Some OFDMA based systems such as LTE allow to adapt the MCS per user, whereas wireless local area networks as specified by IEEE 802.11 use the same MCS on every sub-carrier.**

The performance of such wireless systems that use a single MCS in a frequency selective fading channel can be significantly improved through Sub-Carrier Switch Off (SSO), a simple but powerful alternative to adaptive MCS. SSO deactivates weak sub-carriers that excessively raise the error probability to improve the overall throughput. In this paper, we implement and test SSO in a software-defined radio testbed based on the Wireless Open Access Research Platform (WARP). We present a novel light-weight method for selecting the sub-carriers to be switched off based on the per-sub-carrier channel quality. The results we obtain from our measurements indicate that throughput increases of up to 250% are possible and thus SSO is a highly promising and very low complexity mechanism for future wireless local area networks.

I. Introduction

Orthogonal frequency division multiplexing (OFDM) based wireless systems divide the available frequency spectrum of a channel into several sub-carriers. These sub-carriers allow for the parallel transmission of multiple data streams with longer symbol duration compared to single channel schemes. In combination with a guard interval between symbols, this reduces inter-symbol interference (ISI) and thus lowers sensitivity to complex multipath environments. Additionally, OFDM maintains high throughput, due to multiple simultaneous data streams.

In an environment with frequency selective fading, the sub-carriers may experience very different fading characteristics. Adapting the modulation and coding scheme (MCS) on a per sub-carrier basis is one option to counter this effect. This method ensures high throughput while maintaining acceptable bit error rates but at the same time comes at an increase in transceiver complexity. In addition to the signal processing and computational overhead of such so-called loading algorithms [4], they also incur a considerable signaling overhead for the feedback of channel states and preferred MCS per sub-carrier. As a consequence, no current wireless system adapts the MCS at such a fine granularity. Wireless local area networks (WLAN) as standardized in IEEE 802.11, including the most recent 802.11ac standard, use the same MCS for each sub-carrier. Even sophisticated OFDMA based systems such as LTE only allow to adapt the MCS per user, but use the same MCS on all of the sub-carriers of a given user.

A simple but powerful alternative to adaptive MCS is to simply deactivate sub-carriers with severe fading. Such Sub-Carrier Switch Off (SSO) very effectively combats frequency selective fading and at the same time has much lower computational complexity and feedback overhead. Using SSO to deal with frequency selective fading has been first proposed in [12], [11] and [4]. In these papers, the merits of SSO in conjunction with power loading approaches are investigated by means of simulation.

In this work, we report on our implementation and experimental results of SSO. SSO lowers the average bit error probability and thus increases the packet reception probability. However, at the same time the data that formerly was to be transmitted on the disabled sub-carriers will be appended to the remaining sub-carriers, resulting in longer packets. An efficient SSO mechanism needs to carefully balance these two effects. We design an SSO scheme that enables or disables sub-carriers with respect to the sub-carrier channel estimation. Our implementation is based on the Wireless Open Access Research Platform (WARP) [9], an FPGA based software defined radio system. To the best of our knowledge, our implementation is the first to analyze the effect of SSO in an actual WLAN system.

From our measurements we observe that SSO provides throughput gains of up to 250%. Also, we found that SSO can decreases the necessary signal strength for successful data transmission. Under certain fading conditions we noticed a difference in the necessary signal strength of up to 5 dB. Thus, we consider SSO a highly promising mechanism to improve future wireless local area networks while maintaining low complexity.

The rest of this paper is structured as follows. Section II gives a short introduction into the basic principles of OFDM. It is followed by a description of SSO in an OFDM system. In

Section IV we outline the experimental setup used to evaluate SSO. Experimental results are presented in Section V. Finally, we give a short overview of related work in Section VI.

II. Overview of OFDM

The orthogonal sub-carriers of an OFDM System are spaced closely to each other inside the frequency range of the communication channel. Usually, transmission signals have to be spaced sufficiently far apart to allow them to be filtered properly by the receiver. Additionally the sidebands of every carrier will overlap and corrupt signal quality if no sufficient spacing is applied. OFDM circumvents these difficulties by spacing its sub-carriers reciprocal to the symbol period. With this the sub-carriers are orthogonal to each other and do not interfere, although being placed very closely in the spectrum. The base-band representation of the OFDM signal can be expressed using the inverse discrete Fourier transformation (IDFT) [3].[1]

$$x_t(n) = \text{IDFT}\{X(m+tN)\} = \sum_{k=0}^{N-1} X(k+tN)e^{j2\pi kn/N}$$
$$m = 0,1,...,N-1$$
$$n = 0,1,...,N-1$$
$$t = 0,1,...,\left\lceil \frac{D}{N} \right\rceil - 1 \qquad (1)$$

Here, N is the number of available sub-carriers, $X(k), k = 0,...,D-1$ is a sequence of D data symbols to be modulated onto these sub-carriers, and $x_t(n)$ is the time domain representation of the t-th symbol on the n-th sub-carrier. If necessary, zero-padding is added, i.e., $X(k) = 0, \forall k \geq D$.

On the receiver side, the received signal $y_t(n)$ is demodulated using discrete Fourier transformation to regain the per sub-carrier data symbols.

$$Y(k+tN) = \frac{1}{N}\sum_{n=0}^{N-1} y_t(n)e^{-j2\pi kn/N}$$
$$k = 0,1,...,N-1 \qquad (2)$$
$$t = 0,1,...,\left\lceil \frac{D}{N} \right\rceil - 1$$

The parallel transmission structure of an OFDM systems is its key advantage. In a typical wireless channel, the received signal consists of overlapping arbitrarily delayed copies of the original message due to multipath effects. Therefore, energy from one symbol may bleed into other symbols; this effect is known as inter-symbol interference (ISI) and lowers reception quality. The parallel transmission structure of OFDM with its increased symbol duration lessens this effect.

In a frequency selective fading channel, the different OFDM sub-carriers may experience different channel qualities. This results in certain sub-carriers being more or less error prone. It is important to take these differences in sub-carrier qualities into account. Neglecting the higher error probability of weak sub-carriers easily leads to packet loss and hence throughput degradation due to retransmissions.

[1] For simplicity, we omit guard interval, packet header, and pilot carriers.

III. Sub-carrier Switch Off

SSO provides throughput gains through adequate reaction to frequency selective fading. SSO shares some characteristics with adaptive modulation and coding, however instead of choosing a different MCS for every sub-carrier, SSO just has two options: it either transmits data on a sub-carrier using the MCS chosen for the packet or it makes no use of the sub-carrier. In the latter case, data that would have been transmitted on the deactivated sub-carriers is appended to the other active sub-carriers. This lowers the throughput of the transmission as it takes more time to transmit all the payload data on the remaining active sub-carriers. At the same time, it is possible to use higher order MCSs on the active sub-carriers, while maintaining a low packet and bit error rate (PER, respectively BER). In contrast to adaptive modulation and coding, the feedback overhead consists of only one bit per sub-carrier to signal whether it is active.[2]

The main principle of SSO is to insert a complex zero whenever a data symbol falls onto a deactivated sub-carrier. With χ as the set of deactivated sub-carriers, the base-band representation of SSO-enabled OFDM transmission can be expressed as:

$$x_t^{sso}(n) = \text{IDFT}\{X^{sso}(m+tN)\}$$
$$m = 0,1,...,N-1$$
$$n = 0,1,...,N-1 \qquad (3)$$
$$t = 0,1,...,\left\lceil \frac{D}{N-|\chi|} \right\rceil - 1$$

The number of symbols t that are needed to transmit the complete data stream increases with the number of deactivated sub-carriers. The data symbol stream modified for SSO can be expressed as follows.

$$X^{sso}(l) = \begin{cases} X(\gamma(l)), & \text{if } l \bmod N \notin \chi \\ 0, & \text{otherwise} \end{cases} \qquad (4)$$

Here, $\gamma(l)$ maps from a sub-carrier index in the SSO encoded OFDM transmission to the corresponding data symbol index that is to be placed in this sub-carrier.

$$\gamma(l) = l - ((\lfloor l/N \rfloor)|\chi| + |\{x \in \chi \,|\, x < l \bmod N\}|) \qquad (5)$$

A. Adaptive Sub-carrier Switch Off

When determining χ, the set of sub-carriers to deactivate for SSO, we aim to improve transmission quality and performance. This can be seen as a optimization problem. Therefore,

[2] For example, 802.11ac supports at most 512 sub-carriers corresponding for the maximum bandwidth of 160MHz, resulting in a signaling overhead of only 64 bytes. Certain of these sub-carriers are not used for data transmission and therefore do not need to be included in the signaling. In addition, this overhead can be reduced further due to the fact that neighboring sub-carriers tend to have similar channel quality, allowing for efficient aggregation of state information for neighboring sub-carriers.

we formulate the packet error rate P of a SSO enabled OFDM link as follows:

$$P(M,c) = 1 - \prod_{k=1}^{N}(1 - p(M,c,h(k),w(k))) \qquad (6)$$

Here, $p(M,c,h(k),w(k))$ is the error probability for sub-carrier k with channel $h(k)$ and AWGN noise $w(k)$, given modulation scheme M and coding scheme c. Furthermore, we define the throughput for a channel with conditions h and w as follows:

$$T_{h,w}(M,c,\chi) = b(M) \cdot r_{\text{fec}}(c) \cdot r_{\text{sso}}(\chi) \cdot (1 - P(M,c)) \quad (7)$$

The rate of the convolutional coder is given by r_{fec}, the reduction caused by switching of the sub-carriers in χ is $r_{\text{sso}}(\chi) = 1 - \frac{|\chi|}{N}$, and the throughput rate for modulation M is denoted $b(M)$.

The SSO optimization problem to be solved is thus to select the set of sub-carriers to switch off together with a modulation and coding scheme that maximizes throughput, given channel conditions $h(k)$ and $w(k)$.

$$\arg\max_{M,c,\chi}(T_{h,w}(M,c,\chi)) \qquad (8)$$

B. Threshold-based Adaptive Sub-Carrier Switch Off

For a practical solution, we simplify the general optimization problem given by Equation 8 by splitting the joint optimization of MCS and SSO into two separate optimization problems. On the one hand, we adapt the number of sub-carriers to be deactivated in order to increase the throughput given a packet MCS. On the other hand, we select the modulation and coding scheme that best utilizes the available channel. For the latter, we use one of the existing rate adaptation mechanisms, e.g., [6], [13].

To decide on the quality of a sub-carrier, we make use of the channel estimation techniques of OFDM systems that are used to handle frequency selective fading. This mechanism uses part of a wireless packet's preamble to predict the channel state of every sub-carrier. For the decoding of the payload, every symbol transmitted on a sub-carrier is equalized based on this estimate. This is done e.g. by multiplication with the inverse channel estimate $\hat{H}(k)$ [5].

$$\hat{X}(k) = \frac{Y(k)}{\hat{H}(k)}, k = 0,1,...,N-1 \qquad (9)$$

The influence of the multipath fading channel that causes the frequency selective fading can be described as:

$$y(n) = x(n) \otimes h(n) + w(n), \qquad (10)$$

Here, $w(n)$ is the additive white Gaussian noise (AWGN) and $h(n)$ is the channel impulse response. From the equalization process given in Equation 9, it can be seen that the AWGN noise $w(n)$ affecting a sub-carrier is scaled by the inverse of the channel estimate $\hat{H}(n)$ in the same way as the transmitted signal. This means a low magnitude of the channel estimate (highly attenuated channel) will result in amplification of $w(n)$ during equalization. By this, even in case of perfect channel estimation and presence of AWGN noise, a low magnitude for $\hat{H}(n)$ may lead to a significant distortion in the received sub-carrier signal. For this reason we consider $|\hat{H}(n)|$ as a suitable measure for the decision whether to deactivate a sub-carrier or not.

As all the sub-carriers are independent and their error rate only depends on the sub-channel quality, they will be switched off in the order of their channel qualities. The resulting optimization problem for adaptive sub-carrier switch off of the ξ weakest sub-carriers is:

$$\arg\max_{\xi}(T_{h,w}(M,c,min(\mathcal{N},\xi))) \qquad (11)$$

With $\mathcal{N}$ as set of sub-carriers and $min(\mathcal{N},\xi)$ describing the ξ weakest sub-carriers from $\mathcal{N}$.

With a measure for the sub-carrier channel quality we now propose a threshold-based heuristic solution to Equation 11. We define a channel quality threshold τ to decide whether a sub-carrier should be added to χ. This threshold specifies the sub-carrier quality at which the loss in throughput due to deactivating the sub-carrier equals the loss introduced due to the carrier's contribution to the error rate. The value for τ is equal for every sub-carrier but is MCS dependent. Therefore, for a channel estimate $|\hat{H}(k)|$ we define as a solution to Equation 11:

$$\xi = |\{k, |\hat{H}(k)| < \tau\}| \qquad (12)$$

It is obvious that SSO can not provide any gains for flat fading channels. The separate MCS adaptation mechanism ensures that a MCS is chosen for which the threshold τ is in range of the sub-carrier qualities, i.e., not all sub-carriers will be enabled as should be the case. Only in case signal quality is too low to reliably transmit data with any of the MCSs could it be that all sub-carrier qualities are below the threshold and all carriers would be disabled. This can be easily detected and the SSO functionality could be disabled for this extreme case.

IV. Hardware Platform and Experimental setup

For the experimental evaluation of OFDM sub-carrier switch-off we use the Wireless Open Access Platform (WARP). WARP provides an open source FPGA driven software defined radio platform, where the MAC and PHY layers can be modified in software. Our work builds upon the WARP reference implementation of an IEEE 802.11a-style OFDM PHY-Layer implementation that allows for real-time frame encoding and decoding.

A. PHY-Layer Structure

WARP's OFDM PHY-Layer provides a 10MHz bandwidth OFDM channel that consists of 64 sub-carriers as shown in Figure 1.[3] Out of these 64 sub-carriers, 49 can be used for actual data transmission. The rest are assigned to the DC carrier, pilot carriers, or are left empty according to the IEEE 802.11a specification (see Table I).

[3]Note that we use the sub-carrier numbering from 1 to 64 as given in Figure I for all sub-carrier related figures, i.e., the empty sub-carriers at the left and right border of the frequency spectrum appear in the middle of these figure.

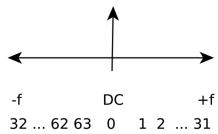

Fig. 1. Sub-carrier Allocation of the WARP OFDM-PHY Layer

Index	Utilization
0 (DC)	Must be empty
7,21,43,57	BPSK pilot tones
1-6, 8-20, 21-26, 38-42, 44-56, 58-63	Data Carriers
27-37	Must be empty

TABLE I
WARP OFDM-PHY SUB-CARRIER ASSIGNMENT

The data packets generated by the WARP OFDM implementation have the packet structure given in Figure 2, that is similar to the one of the IEEE 802.11 protocol family. It starts with a predefined preamble and training symbol sequence. This parts are needed to correct the carrier frequency offset, determine the signal strength and estimate the channel coefficients. After the training sequence follows a packet header. It is encoded with a robust MCS in order to transmit basic parameters for the packet, such as for example source and destination addresses, packet length or CRC. The rest of the packet contains the payload data. It is encoded with the full rate MCS that is a combination of one of the modulation schemes BPSK, QPSK, or QAM16 together with rate $1/2$, $2/3$ or $3/4$ punctured convolutional coding. Additionally, the convolutional coding can be deactivated. While the base rate is fixed and known to both sender and receiver, the full rate is signaled in the packet header and can vary. Therefor, the sender can decide on a full rate MCS while the receiver adapts to it during header decoding. In our experiments we choose the WARP default MCS of QPSK with rate 1/2 punctured convolutional coding for the packet header.

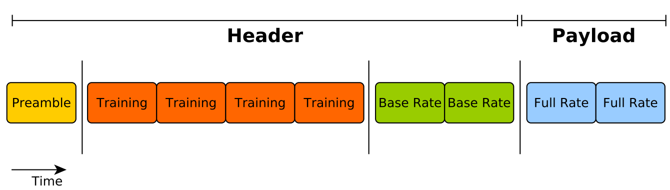

Fig. 2. WARP OFDM-PHY Layer Frame Format

B. WARPnet Measurement Setup

The intended scenario for our SSO algorithm is a wireless local area network with a slowly varying, pseudo-static channel, as experienced for example by notebooks in an office that usually keep their positions for long periods of time compared to the duration of a wireless packet transmission.

For our implementation we use the WARPnet framework. One key advantage of WARPnet is the ability to transfer large amounts of measurement data from the WARP boards via Ethernet to a commodity PC system for further processing. There, data analysis can be done without the restrictions of WARP's embedded FPGA environment.

WARPnet consists of three main components: a MAC layer running on the embedded WARP system, data co-processors running on the commodity PC system, and a python script to configure the experiment parameters. The MAC layer component receives control messages from the experiment control script via Ethernet to configure transmission parameters and to start and stop the packet exchange between the wireless nodes. On the receiver side, the MAC layer sends the received data together with other measurement values via Ethernet to the WARPnet co-processors for further processing. Figure 3 gives an overview of the WARPnet setup.

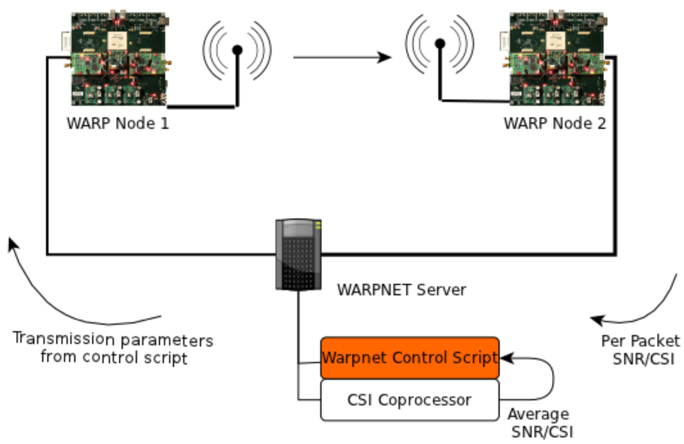

Fig. 3. General Experiment Setup

Although the WARPnet framework is very well suited for the testing of SSO, it has one drawback. WARP lacks the correction of phase offsets that arise from channel coefficient changes or sampling frequency offsets. As a consequence, the longer the duration of a packet transmission, the higher the risk that the phase offset results in packet loss. The maximum packet duration with WARP results from transmitting the maximum payload length of approximately 1400 bytes with the BPSK modulation and code rate $1/2$. To avoid the adverse effects of phase offsets, we limit the maximum number of sub-carriers that can be switched off. In our experiments, we use do not switch off more than 22 sub-carriers. That means, less than half of the available sub-carriers can be switched off, which roughly doubles the packet duration. In turn, we use a payload of approximately half the maximum payload length for all experiments, so that the maximum packet duration remains unchanged.

The throughput of the wireless links established in our testbed can be determined by Equation 7. For the WARP-Platform, values for the throughput base rates $b(M)$ and FEC throughput reduction are given in Tables II and III, respectively. The packet error rate P (see Equation 6) is determined from the number received packets with and without

TABLE II
MODULATION BASE RATES

Modulation	BPSK	QPSK	QAM16
Base Rate	5.72Mbit/s	11.44Mbit/s	22.88 Mbit/s

TABLE III
THROUGHPUT REDUCTION DUE TO CODING

Puncturing	1/2	2/3	3/4
Reduction Factor	0.5	0.67	0.75

payload errors, per SNR level and MCS. For the 10MHz Bandwidth OFDM implementation of the WARP Platform, throughput for the available modulations is defined according to Table II. The throughput reduction due to SSO is given by $r_{\text{sso}}(\chi) = 1 - \frac{|\chi|}{49}$.

C. Measured Values

To evaluate SSO performance we use WARPnet to collect PHY-Layer information for further analysis. This information can be separated into two groups:

- transmission quality measures and
- channel state information.

Transmission quality measures are used to determine throughput, and in particular the throughput gain achieved by SSO. Packets that are lost due to a corrupted header will be lost with or without SSO. Therefore, we ignore them in the performance comparison. We thus investigate packet error rate (PER) as the ratio between the successfully received packets and packets with an erroneous payload, given a certain payload MCS. Additionally, we measure every packet's signal strength. The WARP platform measures this value during packet preamble detection. Equal to Authors in [2] we consider this factor as a proxy to the SNR value, and in the following refer to it as the packet's SNR value.

The channel state information describes the characteristics of the channel and is used to decide on a suitable link configuration, including SSO. The channel state consists of two values. First, we consider the average signal strength of recently received packets as an indicator of the channel's overall attenuation. The second value that forms the channel state are the channel coefficients estimated from the preamble and training sequence. They consist of a complex fixed point number for every sub-carrier describing the magnitude of the channel estimation for this sub-carrier.

D. General Measurement Setup

Our basic test setup consists of two WARP boards running the OFDM PHY layer with the modified WARPnet MAC layer implementation. The nodes are placed in two neighboring testbed rooms so that no line of sight path exists. In order to analyze the behavior of SSO in a frequency faded channel we investigated a range of antenna configurations. Here, we report on results that were obtained for experiments where live monitoring of the channel estimation showed heavily faded carriers.

This kind of fading can be found in around 30%-40% of antenna settings. We also remark that WARPnet's bandwidth is limited to 10MHz, which is considerable less compared to current Wifi standards. IEEE 802.11ac for example uses channels with a bandwidth of up to 160MHz. For comparable indoor scenarios, channels with this bandwidth are likely to have a higher probability of deep faded sub-carriers.

The measurements were performed during night time to avoid that moving persons lead to varying channel conditions. With continues logging of the channel estimation we verified that the channel conditions were relatively stable. An experiment configuration is given by the transmit gain, modulation and coding scheme, set of deactivated sub-carriers χ, and packet length. For each such configuration we run an experiment for a duration of 20 seconds, and average the results over all of the packets transmitted during that time interval. This is necessary, as the estimated channel coefficients and detected SNR might vary slightly due to noise that affects the preamble detection process.

E. Feedback Mechanism

As we consider a pseudo-static channel (see IV-B) we use an implicit feedback mechanism. Therefore, the WARPnet control script decides which sub-carriers to switch off and sets the transmission parameters accordingly for the next time slot. An explicit feedback mechanism for a 802.11 like system could for example included the feedback in the acknowledgment packet that follows every received data packet. This way, no additional signaling packets need to be transmitted, reducing the feedback overhead solely to the bits added to the acknowledgment. For systems with 64 sub-carriers comparable to the WARP setup, this amounts to an additional overhead of only $8B/1460B = 0.5\%$, assuming maximum length data packets.

V. RESULTS

In this section we present the results for the SSO measurements that we conducted on our WARP testbed. The results are separated threefold. First we present throughput measurements that give insights on the gains achievable with SSO. Next we report on a series of measurements to determine the threshold described in section III-B. Finally these findings are validated with additional adaptive throughput measurements.

A. Throughput Gain for SSO

To determine the gains that SSO can achieve in different settings, we measure the transmission quality for the same frequency faded channel with different transmission configurations. To this end, we compare the throughput for the best performing MCS with and without SSO. In addition, this analysis provides us with insights about the influence of channel conditions on the performance of SSO.

As measuring the transmission quality for all possible configurations is cumbersome, we restrict ourselves to measuring a representative subset of channel configurations. For a first insight into the effect of SSO, we use different values for ξ and conduct experiments for all four possible coding rates in conjunction with WARP's three available modulation schemes. Furthermore, to determine the behavior over a range of signal reception strengths, we perform measurements for different transmit power settings.

In the following, we present the results of the throughput measurement for a frequency selective channel where SSO shows a significant throughput gain. In this experiment, we deactivate the 6, respectively 12, weakest sub-carriers. As a reference, the performance of the transmit configuration with all sub-carriers active is measured.

Figure 4 shows the corresponding channel coefficients over time according to the sub-carrier order described in Section IV. As can be seen from the figure, the basic fading pattern remains stable throughout the complete experiment run. Note that the spectrum shows considerable fading for the sub-carriers around sub-carrier number 15. The gap comprises more than 6 heavily faded carriers which suggests that the configuration disabling more than 6 sub-carriers might lead to higher gains.

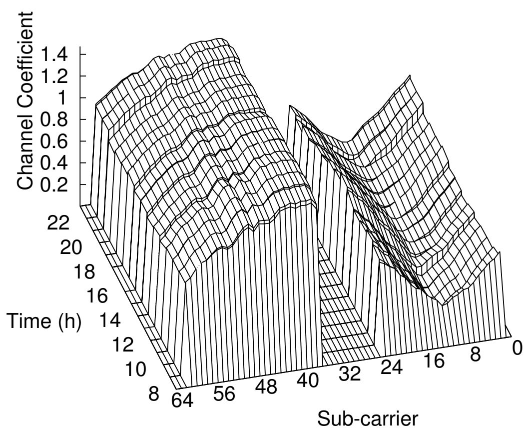

Fig. 4. Channel Coefficients

Figure 5, shows the performance of SSO for QPSK modulations in the described channel. [4] It can be seen, that QPSK still successfully delivers packets under this channel conditions even without SSO. However, the all active sub-carrier configuration only reaches the throughput level of the switched off configuration for high SNR values. Furthermore, for lower SNR values a successful data transmission without SSO is not possible. It can be seen, that SSO significantly decreases the minimum SNR that is needed for successful packet transmission. If we consider for example a throughput of 2Mbit/s we can see that SSO improves the reception by a value of roughly 5dBm. While a QPSK transmission with 1/2 coding and SSO still transmits with this throughput at a SNR of -81dBm the all sub-carrier active equivalent reaches this throughput at -76 dBm. Further, it is obvious that the order in transmission stability imposed by coding is preserved when SSO is applied. For all MCS, 1/2 coded transmission still are the most resilient while 3/4 coding are the lowest. As expected for this channel state deactivating 12 sub-carrier performs better than switching off 6.

[4]The results for BPSK basically give the same insights therefor, we foresee presenting them here.

Figures 6 shows throughput for the same channel using QAM16. It is obvious that for these channel conditions, QAM16 modulation without SSO is not feasible. In contrast, when disabling the weakest sub-carriers it is possible to successfully transmit data with QAM16. Note that disabling 12 sub-carriers performs significantly better than only deactivating 6 sub-carriers. We explain this with the higher error proneness of QAM16 modulations compared to QPSK and BPSK. Note that, in comparison to Figure 5 the SSO enabled QPSK transmissions performs better for lower SNR values while QAM16 provides higher throughput for higher SNR values.

In Figure 7 we show the maximal throughput for this experiment. This plot combines the data of all measured modulations. The throughput of the best configuration with all sub-carriers active compared to the best configuration with SSO is shown over the measured SNR range. The maximum absolute throughput gain can be found at -78dBm and has a magnitude of 4.1Mbit/s, where the conventional scheme achieves a throughput of 0.2Mbit/s, whereas SSO achieves 4.3Mbit/s. N

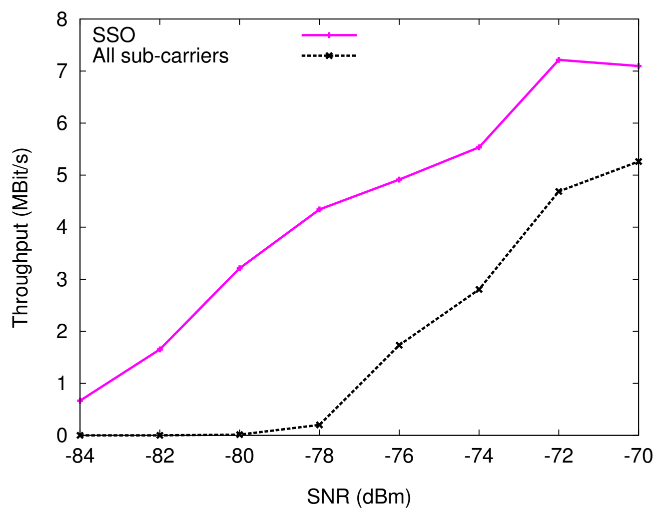

Fig. 7. SSO Throughput Improvement

B. Parameter Configuration for SSO

In a second step we experimentally determined the threshold τ described in section III-B. To this end we use the previously described setup for throughput measurements to collect the throughput for one specific MCS and all possible values for ξ. For every experiment we determine the MCS dependent value for τ from the best performing configuration as the channel state of the weakest active sub-carrier. We find that, the measured channel estimates do not correspond to the exact channel approximation $\hat{H}$. This effect is caused due to WARP's gain control which amplifies the packets before the

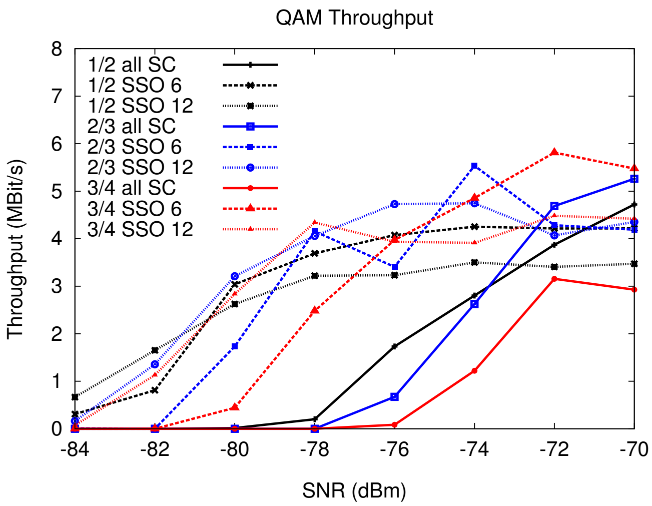

Fig. 5. QPSK Throughput

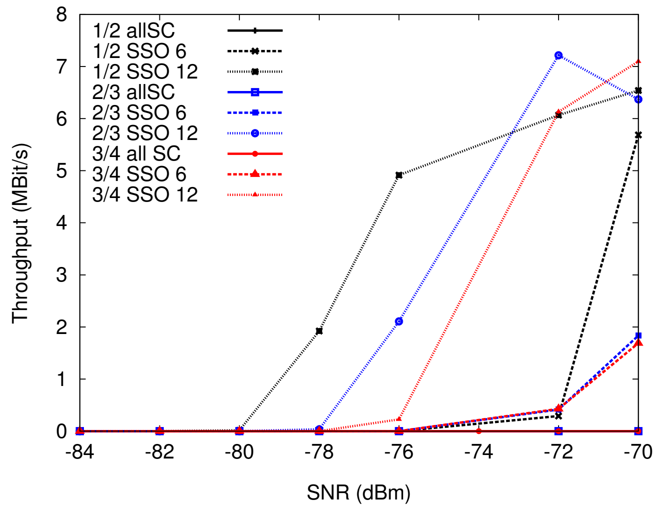

Fig. 6. QAM16 Throughput

channel estimation is done. Therefore, we extend τ_s to be a vector of thresholds depending as well as the receivers gain factor, which corresponds to the measured packet SNR.

We repeat the measurement procedure for several different pseudo-static channels and record the channel states and transmission quality measures. For every MCS, we average the values for τ_s over all conducted experiments. Our low-complexity adaptive SSO strategy determines the number ξ_s of weakest sub-carriers to deactivate depending on the signal strength and the vector of averaged thresholds τ_s as follows:

$$\xi_s = \{k, |\hat{H}(k)| < \tau_s\} \tag{13}$$

We then validate the calculated vector τ_s through further experiments. Hence, we determine the values for ξ that the adaptive SSO approach chooses when applied to the formerly saved experiment data. By comparing these values to the optimum values we can see the performance of the adaptive SSO approach with τ_s.

In the following, we exemplarily show the results of one of these experiments that measures the complete SSO configurations for the QAM16 1/2 MCS. Note that this experiment is run on a arbitrary channel, that is not related to the channel shown in the previous section on which the τ_s configuration is based. Figure 8 shows the channel state for this experiment.

Figure 9 shows the throughput graphs for this measured channel. For better readability we do not show all SSO configurations for higher values of ξ. The behavior of the curves that were left out is similar to that of the shown configurations. For high SNR values where all sub-carriers only suffer low error rates, switching off sub-carriers has negative impact on throughput, that is directly proportional to the number of deactivated carriers. This behavior changes at around -80dBm when the error rate of the weakest sub-carriers reaches the threshold τ. For even lower SNR values, the configuration with all sub-carrier active turns out to be the worst performing configuration with respect to throughput.

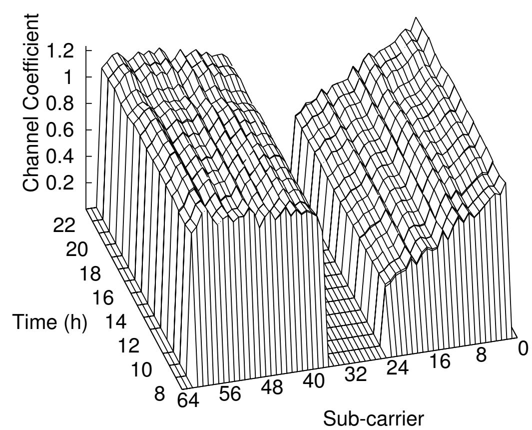

Fig. 8. Channel Coefficients

To verify the configuration of τ, we show an a posteriori analysis of the performance of our adaptive SSO strategy on the same data set in Figure 10. We compare the throughput that the adaptive strategy would have achieved, against the best performing transmit configuration, i.e., the maximum of the curves with a fixed number of sub-carriers switched off. For the former, for each data point we take the performance of the scheme with a fixed number sub-carriers switched off, where the deactivated sub-carriers conform to the threshold τ. As can be seen from the throughput curve of our adaptive strategy, it chooses configurations that are very close to the optimal decision. However, due to the averaging nature of our determination of τ, small deviations from the optimum exist.

C. Adaptive Sub-carrier Switch Off

In order to further confirm our findings for the threshold vectors, we apply the adaptive SSO strategy to live throughput measurements. In contrast to the validation approach in Section V-A, the adaptive strategy directly sets the SSO parameters based on the channel quality measurements reported

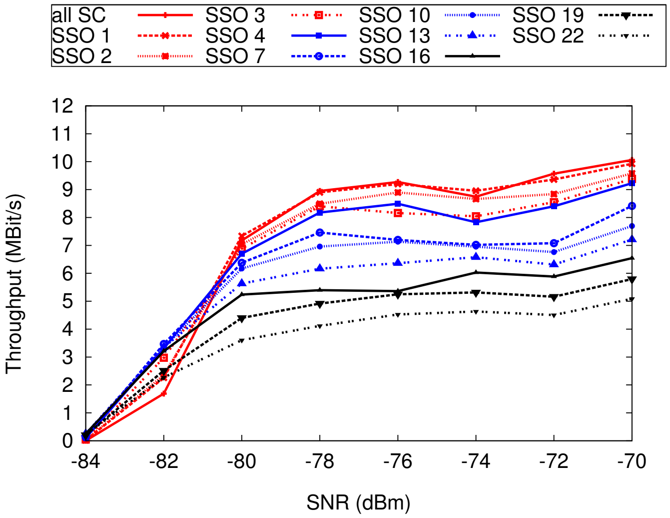

Fig. 9. QAM16 1/2 SSO Channel Reference Measurement

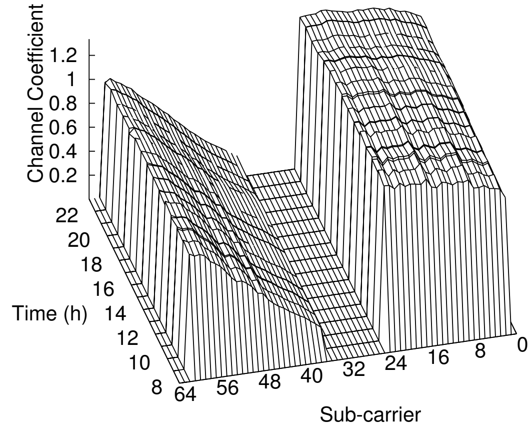

Fig. 11. Channel State for Adaptive SSO Evaluation Measurement

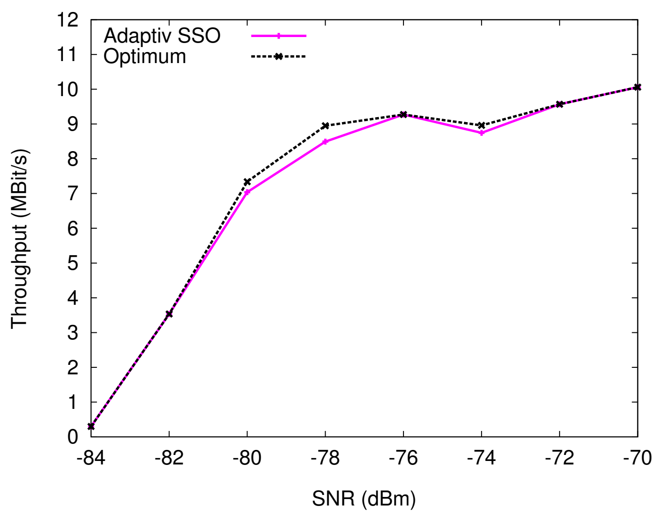

Fig. 10. Adaptive SSO Throughput Compared to Optimal Switch Off Throughput

very low packet error rates for the all active configuration are measured.

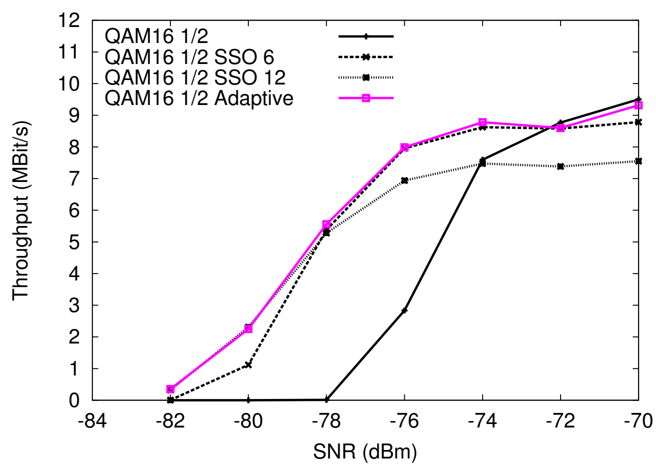

Fig. 12. QAM 1/2 Adaptive SSO Throughput Results

by the receiver. Furthermore, note that the channel for this measurement differs from the one in section V-B that was used for the determination of the threshold vector τ. A successful adjustment of the transmission parameter therefore shows that the adaptive approach also works for unknown channels. Figure 11 depicts the channel state for the adaptive throughput measurement.

Figure 12 shows an overview the results. We compare the adaptive strategy with the static switch off of 6 and 12 sub-carriers for the QAM16 1/2 MCS. As a further reference we also show the throughput for the all active configuration. It can be seen from the graph that the adaptive strategy always performs similar to the best static scheme. The small deviation from the all active configuration that is optimal for low SNR values results from the averaging in the threshold determination. It can easily be prevented by disabling SSO when

VI. Related Work

Sub-Carrier Switch Off is related to the sub-carrier channel quality dependent adaptive assignment of MCS. Cheong et. al [16] propose OFDM with adaptive modulation per sub-carrier to be a promising counter-measure for wireless channels with multipath induced frequency selective fading. They also refer to the possibility of not transmitting data on a deeply faded sub-carrier during the period of a symbol, similarly to the basic idea of sub-carrier switch off. Nonetheless, adaptive modulation systems include higher computational costs for determining the optimal MCS for every sub-carrier and higher transmitter and receiver complexity. Furthermore signaling back the desired MCS combination requires a certain band-width overhead, see ([10], [15], [14]).

A further method to combat frequency selective fading is to assign different transmit powers for every sub-carrier, referred to as power loading [4]. In [7] authors use SSO (referred to as 'sub-band blocking') in order to meet signal power limitations on the proposed power loading scheme. Also the combination of power loading and adapting the MCS has been analyzed, a technique known as bit loading [1]. The idea of SSO as a self contained frequency selective fading counter-measure has first been considered by Puñal and Gross in [12], [11] and [4]. As the newly released IEEE 802.11ac standard has channels with bandwidths of up to 160MHz, they argue that frequency selective faded sub-carriers will have even more impact in the future. Motivated by this, they compare different loading algorithms by means of simulation, namely power loading, bit loading, and SSO. They show that power loading has a minor influence compared to SSO and also show that their proposed combination of SSO and power loading comes close to the results achieved by bit loading.

Our results from a real testbed confirm and extend the findings on SSO that were investigated by Puñal et. al. by means of simulation. We show that in case of deeply faded sub carriers, SSO can archive significant performance gains even for a far smaller bandwidth of only 10 MHz and without the complexity of power loading or bit loading techniques. As the simulation results presented in the works [12], [11] and [4] do not include pure SSO results but results for a combination if SSO and powerloading, and the simulated packets are modulated using a higher modulation scheme (QAM-256), the simulation results are unfortunately not directly comparable to our measurements.

The idea to switch off certain sub-carriers of a OFDM transmission can also be beneficial to cognitive radio (CR) systems [8]. Here however, the idea is to disable sub-carriers that coincide with used parts of the spectrum, thus decreasing interference between CR transmitters and incumbent transmitters.

VII. Conclusion

In this paper we analyzed the gains for Sub-Carrier Switch Off in OFDM wireless local area networks. We designed a light-weight mechanism to select the set of sub-carriers to switch off based on the channel coefficients. We implemented our mechanism using the WARP software-defined radio platform and conducted a range of performance measurements. With the help of extensive channel measurements for all possible SSO configurations, we developed a simple but effective method to decide which sub-carriers to switch off. This method maximize throughput, while having very low computational complexity and feedback overhead.

In case of frequency selective fading, SSO turns out to be a very efficient measure to increase throughput and link stability. Under certain fading conditions we were able to transmit packets with a SNR range increase of up to 5 dBm, with corresponding significant increases in throughput from 0.2 Mbit/s without SSO to 4.3 Mbit/s with SSO. In summary, our analysis indicates that SSO is a very promising method to improve performance and resilience of future IEEE 802.11-style networks.

Our findings also open interesting questions for future work. The combination of rate selection and SSO leaves room for further improvements and measurements. Furthermore, the implementation for efficient feedback mechanisms of channel state information provides some interesting challenges. From a technical point of view, it would be highly interesting to combine SSO with pilot carrier relocation to further improve performance. In addition, SSO can also be applied to the packet headers, which might significantly improve the packet detection rate.

References

[1] A. Barreto and S. Furrer. Adaptive bit loading for wireless ofdm systems. In *Personal, Indoor and Mobile Radio Communications, 2001 12th IEEE International Symposium on*, volume 2, pages G–88 –G–92 vol.2, sep/oct 2001.

[2] J. Camp and E. Knightly. Modulation rate adaptation in urban and vehicular environments: cross-layer implementation and experimental evaluation. In *Proceedings of the 14th ACM international conference on Mobile computing and networking*, MobiCom '08, pages 315–326, New York, NY, USA, 2008. ACM.

[3] J. Cimini, L. Analysis and simulation of a digital mobile channel using orthogonal frequency division multiplexing. *Communications, IEEE Transactions on*, 33(7):665 – 675, jul 1985.

[4] J. Gross, M. Emmelmann, O. Puñal, and A. Wolisz. Dynamic single-user ofdm adaptation for ieee 802.11 systems. In *Proceedings of the 10th ACM Symposium on Modeling, analysis, and simulation of wireless and mobile systems*, MSWiM '07, pages 124–131, New York, NY, USA, 2007. ACM.

[5] M.-H. Hsieh and C.-H. Wei. Channel estimation for ofdm systems based on comb-type pilot arrangement in frequency selective fading channels. *Consumer Electronics, IEEE Transactions on*, 44(1):217 –225, feb 1998.

[6] A. Kamerman and L. Monteban. Wavelan-ii: a high-performance wireless lan for the unlicensed band. *Bell Labs Technical Journal*, 2(3):118–133, 1997.

[7] T. Keller and L. Hanzo. Sub-band adaptive pre-equalised ofdm transmission. In *Vehicular Technology Conference, 1999. VTC 1999 - Fall. IEEE VTS 50th*, volume 1, pages 334–338 vol.1, 1999.

[8] H. Mahmoud, T. Yucek, and H. Arslan. Ofdm for cognitive radio: merits and challenges. *Wireless Communications, IEEE*, 16(2):6–15, 2009.

[9] P. Murphy, A. Sabharwal, and B. Aazhang. Design of warp: A flexible wireless open-access research platform. In *Proceedings of EUSIPCO*, 2006.

[10] K. Pedersen, G. Monghal, I. Kovacs, T. Kolding, A. Pokhariyal, F. Frederiksen, and P. Mogensen. Frequency domain scheduling for ofdma with limited and noisy channel feedback. In *Vehicular Technology Conference, 2007. VTC-2007 Fall. 2007 IEEE 66th*, pages 1792 –1796, 30 2007-oct. 3 2007.

[11] O. Punal, H. Escudero, and J. Gross. Power loading: Candidate for future wlans? *2012 IEEE International Symposium on a World of Wireless, Mobile and Multimedia Networks (WoWMoM)*, 0:1–4, 2012.

[12] O. Punal and J. Gross. Combined subcarrier switch off and power loading for 80 mhz bandwidth wlans. In *Local Metropolitan Area Networks (LANMAN), 2011 18th IEEE Workshop on*, pages 1 –6, oct. 2011.

[13] S. Sen, N. Santhapuri, R. R. Choudhury, and S. Nelakuditi. Accurate: constellation based rate estimation in wireless networks. In *Proceedings of the 7th USENIX conference on Networked systems design and implementation*, NSDI'10, pages 12–12, Berkeley, CA, USA, 2010. USENIX Association.

[14] Z. Shen, J. Andrews, and B. Evans. Adaptive resource allocation in multiuser ofdm systems with proportional rate constraints. *Wireless Communications, IEEE Transactions on*, 4(6):2726 – 2737, nov. 2005.

[15] P. Svedman, S. K. Wilson, L. J. Cimini, and B. Ottersten. Opportunistic beamforming and scheduling for ofdma systems. *Communications, IEEE Transactions on*, 55(5):941 –952, may 2007.

[16] C. Y. Wong, R. Cheng, K. Lataief, and R. Murch. Multiuser ofdm with adaptive subcarrier, bit, and power allocation. *Selected Areas in Communications, IEEE Journal on*, 17(10):1747 –1758, oct 1999.