Evaluating Dynamic OFDMA Subchannel Allocation for Wireless Mesh Networks on SDRs

Robin Klose
TU Darmstadt
Secure Mobile Networking Lab
Mornwegstr. 32, 64293
Darmstadt, Germany
robin.klose@cased.de

Adrian Loch
TU Darmstadt
Secure Mobile Networking Lab
Mornwegstr. 32, 64293
Darmstadt, Germany
adrian.loch@cased.de

Matthias Hollick
TU Darmstadt
Secure Mobile Networking Lab
Mornwegstr. 32, 64293
Darmstadt, Germany
matthias.hollick@cased.de

ABSTRACT

Orthogonal Frequency-Division Multiple Access (OFDMA) has emerged as an advanced technique to enhance resource utilization and efficiency in infrastructure-based networks. However, its performance in wireless mesh networks is mostly unexplored. In this paper, we practically study the benefits of OFDMA in a scenario with multiple co-located transmitters and receivers without centralized controller by means of the software-defined radio platform *WARP*. We propose five different dynamic subchannel allocation strategies and compare their performance to that of OFDM as a baseline. Four of these strategies are constrained to be fair with respect to the number of subchannels allocated per communication link, while the fifth always allocates a subchannel to its best possible communication link. By means of testbed experiments with software-defined radios, we show that the overall bit error rate can be reduced by a factor of ten, while the overall channel capacity can locally be enhanced by 10% to 30%. Further, we use the *Subchannel Avoidance Gain* as a metric to quantify the ability of a dynamic subchannel allocation strategy to avoid subchannel allocations resulting in poor channel conditions.

Categories and Subject Descriptors

C.2.1 [Computer-Communication Networks]: Network Architecture and Design—*Wireless Communication*

Keywords

OFDMA; wireless mesh networks; resource allocation

1. INTRODUCTION

Traditional routing protocols for wireless multi-hop networks are typically built upon physical layers designed for infrastructure-based networks like IEEE 802.11 and are hence limited to the underlying capabilities of these physical layers. Most importantly, such designs cannot exploit concurrency

at the physical layer in decentralized scenarios. However, enhancing concurrency at the physical layer is the fundamental key to make wireless multi-hop networks scale [7].

Novel routing protocols take the broadcasting capability of the wireless medium into account in order to increase efficiency. For instance, broadcast messages are received differently by each node in range of the transmitter due to spatial diversity. Message fragments which arrive corrupted at one node might arrive correctly at a neighbor. Thus, nodes can cooperate to forward partial messages [8]. While such a scheme only allows one node to send data at a time, we aim at harnessing spatial diversity even further by allowing for concurrent transmissions based on OFDMA.

OFDMA is a multi-user variation of OFDM which has emerged as a new key technology to improve efficiency in advanced infrastructure-based wireless networks. It combines OFDM on the physical layer with Frequency Division Multiple Access (FDMA) on the MAC layer, allowing to assign different subcarriers to different nodes in order to increase concurrency. Adjacent subcarriers often experience similar channel conditions and are thus grouped to *subchannels*. The aggregate bandwidth of such a subchannel has to be less than the coherence bandwidth of the channel, which determines how close two subcarriers need to be in order to be similar. There is a large body of work on OFDMA techniques for infrastructure-based networks, as it plays an important role for the development of cellular networks, such as LTE. However, applying OFDMA in wireless multi-hop networks is more challenging due to their decentralized network topology. While in the infrastructure-based case a central base station coordinates traffic, in a wireless multi-hop network OFDMA transmissions can take place concurrently from several transmitters to several receivers.

Different links in a network experience different channel conditions due to spatial diversity. We refer to a realization of the wireless channel between a single transmitter and a single receiver as a *communication link*. In addition, different subchannels of such a communication link may also experience different conditions due to the following reasons. First, the fading characteristics of a wireless channel are inherently frequency selective. For instance, some subchannels on a communication link may experience excellent channel conditions while others suffer from a deep fade. Second, channels may be impaired by interference, which may be frequency-selective as well. This diversity can be exploited at the physical layer by assigning OFDMA subchannels to the communication links which provide the best signal-to-

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

SRIF '13 August 12, 2013, Hong Kong, China.

Copyright 2013 ACM 978-1-4503-2181-5/13/08 ...\$15.00.

interference-plus-noise ratio (SINR). Likewise, the impact of channel impairments might be mitigated by avoiding the allocation of subchannels which result in poor performance.

The denser the network, the more communication links are available and thus the higher is the probability that an allocation strategy can find a link which provides good conditions for a certain subchannel. Hence, OFDMA is particularly promising in combination with wireless multi-hop networks, which typically feature a large number of links.

In this work, we study the performance gains practically achievable by employing OFDMA at the physical layer in a local network with multiple concurrent transmitters and receivers. Our local network can be considered to model a subset of a greater wireless mesh network (WMN), such that all our nodes are in communication range of each other. We also refer to such a portion of a network as a *small-scale network*. We have implemented an OFDMA transceiver on the software-defined radio (SDR) platform WARP [1] in conjunction with WARPLab, a Matlab framework which allows to access WARP resources from the Matlab workspace. While OFDMA for WMNs has been studied in theory and simulation, to the best of our knowledge our system is the first *practical* implementation of OFDMA with *multiple* concurrent transmitters and receivers on the WARP platform. Our contributions are as follows:

- We evaluate the performance of OFDMA in a small-scale network with a practical implementation on SDRs.
- We propose five different subchannel allocation strategies which apply different criteria and constraints.
- We study the performance of the proposed strategies for multiple modulation schemes.
- We study the performance of the proposed strategies under conditions of artificially created interference.
- We show that the small-scale network's overall channel capacity can be increased by 10% to 30%, while the overall bit error rate can be reduced by a factor of ten compared to a traditional OFDM physical layer.

The remainder of this paper is organized as follows. Related work is discussed in Section 2. We outline our SDR measurement technique and metrics in Section 3 and define the proposed dynamic subchannel allocation strategies in Section 4. In Section 5 we describe our experimental setup. We present two different experiments, one to examine the subchannel allocation strategies' ability to exploit excellent channel conditions in Section 6, and one to examine their ability to avoid bad channel conditions under conditions of interference in Section 7. Finally, we discuss the results in Section 8 and conclude the paper in Section 9.

2. RELATED WORK

There exists a large body of work on subcarrier allocation mechanisms for OFDMA. It has been extensively discussed both in theory as well as in simulation. The infrastructure-based case has attracted most attention, but also decentralized settings such as mesh networks have been considered. However, practical evaluations of such schemes are limited. The adoption of OFDMA in WiMAX and LTE has motivated practical studies in infrastructure-based scenarios, but to the best of our knowledge the *decentralized* case has

not been considered yet. In this section, we briefly give an overview on the theoretical and simulative work with a focus on WMNs. Then we survey the existing practical work.

OFDMA in theory/simulation. Strategic interference management and optimal resource allocation strategies in OFDMA-based WMNs are studied by means of cooperative game theory and simulations in [6].

A decentralized fair scheduling scheme is presented in [10]. It decouples the global control problem into two subproblems. A mesh router allocates subcarriers to communication links while the mesh clients negotiate their power allocations themselves. Even though this approach might work for power allocation, it might not take advantage of multi-user channel diversity as subcarrier allocation is coordinated by a centralized device. Channel state information might change too rapidly in practical systems for being taken into account for subchannel allocations by such a centralized approach.

A cross-layer design for joint rate control and OFDMA scheduling for wireless mesh networks is presented in [5]. This work solves the network utility maximization problem analytically, but is based on partly unrealistic assumptions. Most importantly, the relay nodes in the mesh network are assumed to be able to simultaneously receive and transmit via disjoint subcarriers. This assumption does not hold for practical systems since off-the-shelf radio network interfaces can only be operated in half-duplex mode.

The authors of [13, 14, 15] employ simulations to evaluate OFDMA-based PHY and MAC mechanisms for wireless mesh networks. They propose a protocol called *Concurrent Transmission or Reception Multiple Access* [13], which assigns multiple subchannels to each link in the network depending on traffic demands and the number of interfering links. Furthermore, they introduce a MAC protocol called CoCo-MAC [14, 15], which allocates groups of subcarriers within pre-assigned subchannels for concurrent communication with multiple neighbors. Besides increasing concurrency, CoCo-MAC also aims at exploiting multi-user channel diversity by assigning subcarriers to a node's neighbors in a way that selects the subcarriers with the best channel conditions within the pre-assigned subchannels.

OFDMA in practice. Practical issues have been considered for the implementation of OFDMA both for WiMAX and LTE. While WiMAX also considered mesh networking [2], such decentralized scenarios were deprecated in a later revision of the standard. In [4], the authors present a demonstration of a downlink OFDMA WiMAX transceiver based on the WARP platform. However, it only deals with the infrastructure-based case. Similarly, practical real-time SDR testbeds for LTE have been studied in [3]. Also 802.11a has been extended to support OFDMA and implemented on the SORBAS SDR [11]. Again, our work stands apart from both implementations as we specifically exploit the large diversity available in decentralized scenarios.

Further practical work not directly related to standards also exist. The authors of [12] implement OFDMA in the real-time mode of the WARP platform to evaluate a mechanism that distributes power among subchannels according to CSI measurements in order to improve performance. Nevertheless, they do not deal with subcarrier allocation, but focus on power allocation. Hence, we conclude that a significant amount of work has been done to understand OFDMA, but when it comes to the practical case in *decentralized* scenarios, there is a dearth of implementations.

3. EVALUATION TECHNIQUES

We evaluate our subchannel allocation strategies on the WARP software-defined radio. The WARPLab framework enables us to implement strategies flexibly in Matlab, but our setup results in a non-real-time system. Essentially, instead of processing data on the SDR itself, it is sent via Ethernet to a computer and processed in Matlab. Note that this approach only relocates processing from the WARP board to Matlab, but measurements are still online. However, sending data from Matlab to the SDRs and back incurs a significant delay. Hence, channel state information (CSI) might be outdated by the time it is processed.

3.1 CMSP Technique

Since subchannel allocation is based on CSI measurements, the non-real-time nature of WARPLab poses a strong limitation. In order to circumvent this problem, we resort to a technique called *Concurrent Measuring of Subchannel Performances* (CMSP) [9] which exploits the coherence bandwidth of channels. We briefly introduce how it works, but details can be found in [9]. Essentially, the technique allows to measure all possible subchannel allocations in one single transmission, which eliminates the outdated CSI problem.

CMSP is designed to be used in a small-scale network, i.e., a network with m transmitters and n receivers where all nodes are in range of each other. In a traditional real-time approach to OFDMA subchannel allocation, a two step process is needed. First, the m transmitters measure CSI to all n receivers and coordinate to decide which transmitter gets which subchannel to communicate with which receiver. Second, the actual communication takes places according to the agreed allocation. In CMSP both steps are merged. The key to enable this technique is exploiting that all subcarriers within a subchannel experience virtually the same channel conditions. A subchannel is a group of adjacent subcarriers whose aggregate bandwidth does not exceed the coherence bandwidth B_c . Traditionally, all subcarriers in a subchannel are assigned as a block to one link. On the contrary, in CMSP at least one subcarrier in the subchannel is assigned to each transmitter statically. This means that on each subchannel all receivers get data from all transmitters. Hence, after the transmission, each receiver can *extrapolate* how the performance on a certain subchannel would have been if all subcarriers had been allocated to any of the m transmitters, since it knows how the performance was for one (or more) subcarrier of that transmitter in the subchannel.

In other words, we use a static allocation of *subcarriers* which allows us to infer after the transmission how any possible allocation of *subchannels* would have been. While we analyze the performance of multiple allocations after the actual transmission has taken place, note that this is *not* equivalent to an evaluation based on traces. CMSP does not only collect CSI in form of traces to use them for an offline simulation, but sends actual data which is then used as a basis for computing bit error rates, symbol error rates and other relevant metrics. No simulation is used at all. CSI is collected exclusively to be used as an input for the allocation strategies, which use it in order to decide how to distribute subchannels. Further, note that we use CMSP as a *measurement* technique to circumvent the non-real-time nature of WARPLab and evaluate our strategies. In a productive, real-time environment, the aforementioned traditional two-step approach involving CSI feedback would be used.

3.2 Metrics

We evaluate the performance of dynamic subchannel allocation strategies in terms of the symbol error rate (SER), the bit error rate (BER) and the channel capacity according to Shannon's theorem (CAP). Additionally, we use the *Subchannel Avoidance Gain* (SAG) [9], which is a metric that uses the BER and the SER to determine if symbol errors become less severe due to a certain allocation s compared to a baseline b . Equation 1 shows how it relates both metrics.

$$\frac{\text{BER}(s)}{\text{BER}(b)} = (1 - \text{SAG}(s, b)) \frac{\text{SER}(s)}{\text{SER}(b)} \quad (1)$$

Essentially, the SAG metric exploits the characteristics of Gray coding. Severe symbol errors typically cause the received sample to fall in a decision region which is non-neighboring with the actual region the symbol belongs to. Hence, when using Gray coding, the more severe the error is, the more bits of the symbol are wrong. If the symbol errors caused by strategy s are similarly severe to the ones caused by the baseline b , the BER and SER ratios are equal and the SAG is zero. If s still causes the same amount of symbol errors than b , but errors are less severe and thus the BER is reduced to, e.g., half its value, the SAG is 50%.

To reduce the severeness of errors, allocation schemes avoid subchannels with bad channel conditions. However, due to fairness constraints, in turn they might have to cede subchannels which experience good conditions. For example, if subchannel i_1 experiences a deep fade on link l_1 , the allocation scheme may reassign it to link l_2 , where symbol errors still occur, but are less severe. In turn, the scheme may have to assign subchannel i_2 on link l_2 to link l_1 due to fairness. If subchannel i_2 on link l_1 is worse than i_2 on l_2 , but not as bad as i_1 on l_1 , the overall BER might improve as the deep fade has been avoided, while the SER might increase. The SAG metric aims at quantifying precisely this trade-off, which is neither captured by the BER nor the SER alone.

Note that the error rates of both strategies s and b must be greater than 0 in order to be able to calculate the SAG metric. Moreover, the result of the SAG metric might be negative if the allocation strategy chooses poor links.

4. ALLOCATION STRATEGIES

In this section we introduce our allocation strategies, which assign subchannels to links in a wireless mesh network in order to improve transmission quality. The problem setting is as follows. We consider a small-scale network, i.e., a set of m transmitters and n receivers which are all in transmission range of each other. Hence, there are $m \times n$ links available in the network. All transmitters shall send data at the same time, but using disjoint sets of OFDMA subchannels, which translates into an interference-less transmission. Thus, the goal is finding the aforementioned sets of subchannels which maximize the performance in terms of BER, SER and CAP. We assume that all transmitters have data for all receivers.

Note that subchannels are not just distributed among the m transmitters, but among the $m \times n$ available links, i.e., the strategies also take into account to which receiver data is delivered. The reason for this is that each link of a transmitter might experience a different quality. For example, while a subchannel on the link from transmitter t_1 to receiver r_1 might support a high modulation scheme, the same subchannel on the link from t_1 to r_2 might be very poor.

Figure 1 shows the problem setting. There are $m \times n$ links and i subchannels. Each subchannel must be allocated to only one link. The larger the SINR on a certain subchannel of a link, the better its performance.

4.1 Reference Strategy: fair_{OFDM}

In order to measure the performance improvement achieved by allocating OFDMA subchannels to different transmitters, we compare our results to a reference strategy based on OFDM, which we call fair_{OFDM}. In this case, *all* available i subchannels shown in Figure 1 are assigned to the same link. Hence, the reference strategy cannot avoid the subchannels of the chosen link which have a low SINR.

In our $m \times n$ scenario, we measure the performance of fair_{OFDM} as follows. As OFDMA is not used, different nodes cannot send data at the same time. Hence, in order to be fair, transmitters t send one at a time to each of the receivers r , until all links have been used once. The performance of fair_{OFDM} is the average of all $m \times n$ transmissions. Note that when using the CMSP technique introduced in Section 3.1, the $m \times n$ transmissions can be measured in just one step. The SER, BER and CAP performance can then be calculated as shown in Equations 2, 3 and 4, respectively.

$$\text{SER}(\text{fair}_{\text{OFDM}}) = \frac{\sum_{t=1}^{N_{\text{tx}}} \sum_{r=1}^{N_{\text{rx}}} \text{SER}_{\text{OFDM}}(t \rightarrow r)}{N_{\text{tx}} \cdot N_{\text{rx}}} \quad (2)$$

$$\text{BER}(\text{fair}_{\text{OFDM}}) = \frac{\sum_{t=1}^{N_{\text{tx}}} \sum_{r=1}^{N_{\text{rx}}} \text{BER}_{\text{OFDM}}(t \rightarrow r)}{N_{\text{tx}} \cdot N_{\text{rx}}} \quad (3)$$

$$\text{CAP}(\text{fair}_{\text{OFDM}}) = \frac{\sum_{t=1}^{N_{\text{tx}}} \sum_{r=1}^{N_{\text{rx}}} \text{CAP}_{\text{OFDM}}(t \rightarrow r)}{N_{\text{tx}} \cdot N_{\text{rx}}} \quad (4)$$

4.2 Dynamic Subchannel Allocation Strategies

Our five proposed dynamic subchannel allocation strategies follow a common scheme, but apply different criteria. Figure 2 shows an overview of how the strategies work. The goal is to allocate each available subchannel to one of the available links. Essentially, a strategy chooses one of the $m \times n$ links according to its criteria in a first step, and allocates the subchannel on which that link experiences the best SINR to it in a second step. The process is repeated until no more subchannels are left.

If a strategy is *fair*, each link gets the same number of subchannels. Once a link reaches its fair share, it is removed from the list of initially $m \times n$ links out of which the strategy can choose from in each iteration. Note that fairness is *per*

Figure 1: Problem setting. The i available subchannels must be distributed among $m \times n$ links.

link, and thus each sender transmits data to each receiver over the same number of subchannels, i.e., there is no sender which only sends to some of the receivers but not to others.

Moreover, a strategy can be *dynamic* or *static*. A dynamic strategy allocates only one subchannel to the chosen link in each iteration. Hence, the same link might be chosen again in a later iteration if it still has not reached its fair share. As opposed to that, a static strategy allocates in each iteration the complete fair share, i.e., the link is directly assigned the maximum number of subchannels out of its best ones. Hence, the link is not chosen again in subsequent iterations.

We compare five different strategies for choosing links. The best_{dmax} strategy is not fair and hence assigns each subchannel to the link which experiences the highest SINR on that subchannel. The other four strategies are fair and differ in how they choose the next link out of the $m \times n$ available ones for which they are going to allocate a subchannel.

1. fair_{dmax} is dynamic and chooses the link which features the subchannel with the highest SINR, which is then allocated.
2. fair_{dmin} is dynamic and chooses the link which features the subchannel with the lowest SINR, but then allocates the subchannel of the chosen link with the highest SINR. Hence, it ensures that links with bad subchannels get the best subchannels they have.
3. fair_{smin} is static and also chooses the link which features the subchannel with the lowest SINR.
4. fair_{rand} is dynamic and chooses a random link out of the $m \times n$ available ones. It then allocates the subchannel with the highest SINR of the chosen link.

5. EXPERIMENT SETUP

We carry out experiments in a small-scale network with four transmitters and four receivers. Hence, each subchannel is allocated to one out of 16 available links. Figure 3 depicts an overview of our setup. We connect the four transmitters to the same WARP board, but the data of each transmitter is handled as if each antenna were an independent node. The same is true for the four receivers, which are connected to a second WARP board. The shared boards are used exclusively to achieve synchronization, which is a strong requirement in OFDMA. We abstract from the synchronization issue as it is unrelated to the performance of allocation algorithms and thus out of scope of this paper.

We place the antennas in an empty conference room, thus obtaining a line of sight (LOS) environment. Distances are as specified in Figure 3. Both boards are connected to a computer running Matlab and experiments are carried out using CMSP, as described in Section 3.1. We conduct two series of experiments, whose parameters are summarized in Tables 1 and 3. N_{tx} is the number of transmitters, N_{rx} is the number of receivers, N_{subchn} is the number of subchannels i to be shared among all links, N_{pil} is the number of pilot symbols used to measure CSI, N_{dat} is the number of data symbols used to measure performance in terms of the metrics presented in Section 3.2, BPS is the number of bits per symbol, T_s is the duration of an OFDMA symbol, T_G is the guard space between OFDMA symbols, Δf is the sub-carrier spacing, f_L is the lowest baseband frequency, f_H is the highest baseband frequency and B_{PB} is the passband bandwidth.

Figure 2: Common scheme for all subchannel allocation strategies.

Table 1: Transmission parameters

N_{tx}	N_{rx}	N_{subchn}	N_{pil}	N_{dat}	BPS
4	4	32	20	20	variable
T_S [μs]	T_G [ns]	Δf [kHz]	f_L [kHz]	f_H [MHz]	B_{PB} [MHz]
6.8	400	147.059	73.529	9.485	18.824

6. EXPLOITING EXCELLENT CSI

In this section, we examine the performance of our proposed dynamic subchannel allocation strategies for multiple modulation schemes in order to study the strategies' ability to provide allocations resulting in excellent channel conditions. Different modulation schemes require different SINR levels as a minimum to be reliably operable. The experiment consists of five series of transmissions with BPS levels 1, 2, 4, 6, and 8, while all other parameters stay unaltered.

The transmission parameters are given in Table 1. Since channel conditions change over time, recording the five series of transmissions one after another may introduce a bias with respect to channel quality. Hence, we perform the series of transmissions in an interleaved order, i.e., we first record the first transmissions of each series, then the second transmissions of each series, and so on. Each of the five series contains 100 records of transmissions. The experiment is conducted within a timeframe of 96 minutes.

6.1 Observations

The mean performance of the strategies throughout the respective series of transmissions is depicted in Figures 4 and 5. For one bit per symbol, there are virtually no errors for all strategies, as can be seen in Figure 4. With increasing BPS; however, the OFDM strategy has the steepest growth in terms of error rates. The best SER and BER performances are achieved with $best_{dmax}$. It has virtually no symbol errors for up to six bits per symbol. Among the fair subchannel allocation strategies, $fair_{smin}$ and $fair_{dmin}$ provide the best performance in terms of error rates.

In Figure 5, the mean channel capacity is almost constant for each of the strategies over different levels of BPS. This is an indication that the mean channel conditions do not change significantly between different series of transmissions. Variations in CAP, as indicated by the error bars, are due to variations in channel quality over time, which are inherently

Figure 3: Experimental setup for measurements.

given by the wireless channel. The fair strategies perform almost equally and achieve on average about 10% more CAP than OFDM. The $best_{dmax}$ strategy performs best in terms of CAP. It provides on average about 30% more channel capacity than OFDM.

6.2 Evaluation of the Results

Intuitively, the $fair_{dmax}$ strategy seems most promising for high BPS values, since it ignores links with bad conditions and focuses on allocating subchannels to links with excellent conditions, instead. The rationale behind such an intuition would be that links which experience low SNRs might anyhow be unsuitable for high modulation schemes. Nevertheless, $fair_{dmax}$ performs worse than the $fair_{smin}$ and the $fair_{dmin}$ strategies in terms of SER, which first select communication links subjected to bad channel conditions on some of their subchannels. This is a clear indication that communication links that are subjected to subchannels with very bad channel conditions still experience also good channel conditions on other subchannels, which are at least in part so good that no symbol errors occur even for higher values of BPS. While $fair_{smin}$ and $fair_{dmin}$ manage to avoid allocating subchannels with poor quality, $fair_{dmax}$ might be forced to allocate subchannels with poor quality due to the fairness constraint, after having allocated excellent subchannels.

Figure 4: Mean error rates as a function of BPS

Figure 5: Mean channel capacities as a function of BPS: Error bars indicate 5% and 95% quantiles.

The SAGs of the different subchannel allocation strategies with respect to OFDM are given in Table 2. There are some NaN (*Not a Number*) entries in the table when no symbol errors occur at all. In these cases, it is not possible to evaluate whether a subchannel allocation strategy manages to reduce the impact of severe symbol errors in addition to its performance. The results for BPS = 8 can be considered to be most relevant since in that case the SAG metric can be calculated for all strategies. This last row clearly shows that $\text{best}_{\text{dmax}}$ performs best in terms of SAG with a value around 22%. The $\text{fair}_{\text{smin}}$ and the $\text{fair}_{\text{dmin}}$ strategies are specifically designed to avoid bad channel conditions and still manage to allocate subchannels with good conditions. They are yet constrained to be fair and thus achieve only about 15% SAG. These two strategies also perform quite similarly for BPS

Table 2: SAG($\cdot$, OFDM)

BPS	$\text{best}_{\text{dmax}}$	$\text{fair}_{\text{dmax}}$	$\text{fair}_{\text{rand}}$	$\text{fair}_{\text{smin}}$	$\text{fair}_{\text{dmin}}$
1	NaN	0%	NaN	NaN	NaN
2	NaN	9.67%	9.67%	NaN	NaN
4	NaN	2.42%	-12.20%	17.14%	17.14%
6	NaN	-2.08%	6.30%	15.97%	16.51%
8	22.29%	0.43%	11.26%	15.06%	14.80%

Figure 6: Mean error rates as a function of B_n

Table 3: Transmission parameters

N_{tx}	N_{rx}	N_{subchn}	N_{pil}	N_{dat}	BPS
2	4	64	20	20	6
T_S [μs]	T_G [ns]	Δf [kHz]	f_L [kHz]	f_H [MHz]	B_{PB} [MHz]
6.8	400	147.059	73.529	9.485	18.824

values 4, 6 and 8. The $\text{fair}_{\text{dmax}}$ and $\text{fair}_{\text{rand}}$ strategies perform worst. They achieve SAGs between about -10% and +10% by chance and have quite unstable SAGs at different levels of BPS. Since $\text{fair}_{\text{dmax}}$ is tied to the fairness constraint, it can be forced to allocate subchannels with bad conditions after having allocated subchannels with excellent conditions. The $\text{fair}_{\text{rand}}$ strategy inherently selects communication links in random order and may select communication links which experience good or bad channel conditions by chance.

7. AVOIDING INTERFERENCE

In this section we examine our proposed dynamic subchannel allocation strategies under conditions of artificially created interference. In this experiment, T3 and T4 are not used for data but for noise transmissions, in order to deliberately distort OFDMA transmissions of T1 and T2. We use artificial noise to emulate interference in a WMN where further away transmitters are not in communication range, but still in mutual interference range of each other. The noise interferers transmit noise with the lowest possible baseband and RF gains adjustable in WARPLab, while T1 and T2 transmit with a power level appropriate for communication.

The transmission parameters are given in Table 3. The OFDMA signals occupy $B_{\text{PB}} = 18.8$ MHz in the 2.4 GHz passband. The artificial noise is colored and its bandwidth is varied, while the OFDMA parameters stay unaltered. The noise interferers transmit noise at center frequencies -5 MHz and +5 MHz in the baseband, respectively. The experiment consists of eleven series of transmissions. In each one, the noise bandwidth B_n of each interferer is increased in steps of 0.5 MHz. While in the first series the noise interferers are disabled, in the last series the total interference bandwidth reaches $2 \cdot 5 \text{ MHz} = 10 \text{ MHz}$. Each series contains 81 records of transmissions. Again, we perform the series in an interleaved order to avoid biased channel conditions. The experiment is conducted within 112 minutes.

Figure 7: Mean channel capacities as a function of B_n : Error bars indicate 5% and 95% quantiles.

7.1 Observations

The mean performance of the different strategies is depicted in Figures 6 and 7. Note that the total interference bandwidth is twice the noise interference bandwidth B_n as there are two interferers in disjoint frequency bands. Figure 6 shows that the error rates of OFDM increase rapidly between $B_n = 0$ MHz and $B_n = 0.5$ MHz, but not as steep for further increasing B_n . The error rates of $\text{fair}_{\text{dmax}}$ are only slightly below the respective error rates of OFDM, if at all. In contrast, the $\text{fair}_{\text{smin}}$, the $\text{fair}_{\text{dmin}}$ and the $\text{best}_{\text{dmax}}$ strategies are much less affected by interference. For higher values of B_n , these three strategies stay at a level around 1% SER. The SER of $\text{fair}_{\text{rand}}$ reaches a level of about 5% for increasing B_n . The mean CAP values fall by tendency with increasing B_n , as seen in Figure 7. The error bars indicate a CAP variability of up to $\pm 10\%$, which are due to variations of channel quality over time.

7.2 Evaluation of the Results

Out of the fair strategies, $\text{fair}_{\text{smin}}$ and $\text{fair}_{\text{dmin}}$ perform best. In terms of error rates, they achieve a reduction by an order of magnitude with respect to OFDM, which is significant. They also slightly outperform the other fair strategies in terms of CAP in this experiment. The $\text{best}_{\text{dmax}}$ strategy performs best with respect to error rates and CAP. Nevertheless, $\text{best}_{\text{dmax}}$ only marginally outperforms $\text{fair}_{\text{smin}}$ and $\text{fair}_{\text{dmin}}$ in terms of error rates, which indicates that these fair strategies avoid bad channel conditions very well.

7.2.1 Sensitivity to Interference

The steep increase of the OFDM error rates when switching from no interference to $B_n = 0.5$ MHz of interference indicates a strong sensitivity to narrowband interference. Even though the noise bandwidth of both interferers taken together is only $\frac{2 \cdot 0.5 \text{ MHz}}{18.824 \text{ MHz}} \approx 5.3\%$ of the total bandwidth, it has a serious impact on performance. This might be attributed to the fact that the spectra of OFDM subcarriers overlap with each other in the frequency domain, such that narrowband interference affects several subcarriers.

In contrast, $\text{best}_{\text{dmax}}$, $\text{fair}_{\text{smin}}$ and $\text{fair}_{\text{dmin}}$ perform very well under conditions of narrowband interference with respect to error rates. In addition, their SERs and BERs are quite unaffected by further increasing the interference bandwidth. This indicates that different communication links

are affected quite differently on different subchannels by interference, which can be exploited by dynamic subchannel allocation strategies that focus first on avoiding allocations resulting in poor channel conditions. Both $\text{fair}_{\text{smin}}$ and $\text{fair}_{\text{dmin}}$ slightly outperform the other fair strategies also in terms of CAP, which affirms that they are best suited among the fair strategies to sustain under conditions of interference. The $\text{best}_{\text{dmax}}$ strategy performs best out of all subchannel allocation strategies as it is not constrained to fairness.

The $\text{fair}_{\text{dmax}}$ strategy also prefers communication links with good channel conditions in the order of selection. However, as opposed to $\text{best}_{\text{dmax}}$, it is constrained to fairness and can finally be forced to allocate subchannels to links resulting in poor channel conditions. The $\text{fair}_{\text{rand}}$ strategy does not act in favor of link qualities and thus performs between the $\text{fair}_{\text{dmax}}$ strategy and the $\text{fair}_{\text{smin}}$ and $\text{fair}_{\text{dmin}}$ strategies.

In terms of capacity gain with respect to the OFDM baseline, $\text{best}_{\text{dmax}}$ performs about 25% to 30% better than OFDM. The fair strategies achieve about 10% to 15% more CAP than OFDM, while $\text{fair}_{\text{dmin}}$ performs best, followed by $\text{fair}_{\text{smin}}$, $\text{fair}_{\text{rand}}$ and $\text{fair}_{\text{dmax}}$.

7.2.2 Subchannel Avoidance Gain

Table 4 shows the SAGs of the different subchannel allocation strategies in relation to the OFDM baseline for different values of B_n . The $\text{best}_{\text{dmax}}$ strategy reaches a SAG of about 10% at $B_n = 0$ MHz. In the presence of narrowband interference, its SAG immediately increases to about 24%. For further increasing interference bandwidth B_n , it slowly declines to a level of about 18%. The $\text{fair}_{\text{smin}}$ and $\text{fair}_{\text{dmin}}$ strategies both experience a SAG of about 10% when the artificial noise interferers are disabled. In the presence of narrowband interference, their SAGs increase to about 24% and 19%, respectively. Their SAGs slowly decline with increasing B_n and finally reach a level of about 17% for $B_n = 5$ MHz. Thus, $\text{best}_{\text{dmax}}$, $\text{fair}_{\text{smin}}$, and $\text{fair}_{\text{dmin}}$ achieve their greatest SAGs in the presence of narrowband interference. The reason for this is that a communication link which suffers from interference on a few subchannels is more likely to find a subchannel with good channel conditions if there is a broader range of subchannels which is not affected by interference. Conversely, in the presence of broadband interference, these strategies inevitably also allocate subchannels with a little worse channel conditions, resulting in slightly more severe symbol errors.

The SAG of $\text{fair}_{\text{dmax}}$ drops from about 0% to about -15% when enabling narrowband interference. Since $\text{fair}_{\text{dmax}}$ is constrained to fairness, it finally also allocates subchannels to communication links that provide poor channel quality due to interference, resulting in more severe symbol errors. The $\text{fair}_{\text{rand}}$ strategy experiences SAGs around 0% for all levels of B_n as it chooses links randomly. In the long run, its mean SAG value is quite neutral, even though individual transmissions indeed result in positive or negative SAGs.

8. DISCUSSION

Our experiment in Section 7 shows that the error rates of $\text{fair}_{\text{smin}}$ and $\text{fair}_{\text{dmin}}$ are an order of magnitude smaller than the error rates of OFDM with interference. This finding is especially valuable for meshed networks since distant nodes can cause interference even when they are not in range. The reason for $\text{fair}_{\text{smin}}$ and $\text{fair}_{\text{dmin}}$ performing almost equally can be attributed to our setup of closely spaced nodes. Un-

Table 4: SAG($\cdot$, OFDM)

B_n [$\times 0.5$ MHz]	best _{dmax}	fair _{dmax}	fair _{rand}	fair _{smin}	fair _{dmin}
0	9.76%	-0.20%	4.00%	9.76%	9.76%
1	23.88%	-14.51%	1.18%	24.22%	19.15%
2	23.71%	-14.95%	0.77%	24.19%	18.80%
3	22.62%	-11.32%	-1.38%	22.93%	21.09%
4	23.42%	-9.55%	1.57%	23.18%	22.00%
5	22.42%	-9.00%	0.32%	22.34%	21.65%
6	22.12%	-8.16%	-0.57%	21.60%	21.33%
7	17.68%	-7.92%	0.73%	15.83%	15.42%
8	19.04%	-8.31%	2.32%	16.09%	15.09%
9	17.98%	-8.23%	2.01%	16.79%	16.52%
10	17.98%	-8.15%	3.51%	17.42%	17.53%

less fades are not extremely steep, fair_{dmin} makes decisions similar to those of fair_{smin}. However, since frequency selectivity increases with the channel delay spread, fair_{dmin} might outperform fair_{smin} in wider spaced topologies.

Channel estimates can be assumed to be valid only for a limited amount of time, namely the coherence time T_C . T_C decreases with increasing mobility of nodes and can theoretically be assumed to be infinite in immobile environments. Even though we merely focus on a network scenario with stationary nodes, our results apply for transmissions lasting no longer than T_C in mobile environments as well.

9. CONCLUSION

The scalability of wireless mesh networks does not only depend on the performance of routing algorithms, but also on the underlying PHY and MAC. We study the performance of dynamic OFDMA subchannel allocation strategies in a local network with multiple concurrent transmitters and receivers by means of the SDR platform WARP [1]. We propose five different dynamic subchannel allocation strategies that we compare to an OFDM strategy as a baseline. Four of these strategies are constrained to fairness with respect to the number of subchannels allocated per link, while the fifth always allocates a subchannel to its best link. We evaluate the strategies' ability to exploit subchannels providing excellent channel quality, and to avoid poor channel quality under conditions of interference. Our results show that strategies which focus on trying to avoid bad channel conditions consistently perform best. The overall channel capacity can be enhanced by 10% to 30%, while error rates can be reduced by an order of magnitude compared to OFDM.

Acknowledgments

This work was supported by the LOEWE Priority Program Cocoon (www.cocoon.tu-darmstadt.de), the DFG CRC 1053 MAKI (www.maki.tu-darmstadt.de) and LOEWE CASED (www.cased.de).

10. REFERENCES

- [1] Rice University WARP Project. <http://warp.rice.edu>.
- [2] IEEE Standard for Local and Metropolitan Area Networks - Part 16: Air Interface for Fixed

Broadband Wireless Access Systems. *IEEE Std. 802.16-2004*, 2004.

- [3] R. Chen, Q. Cai, K. Alecke, O. Lazar, and T. Kaiser. A Real-Time PRE-MIMO-LTE Software Radio Testbed. In *Proceedings of the European Signal Processing Conference (EUSIPCO)*, 2007.
- [4] H. Eslami, G. Patel, C. P. Sukumar, S. V. Tran, A. M. Eltawil, R. Rao, and C. Dick. Demonstration of Highly Programmable Downlink OFDMA (WiMax) Transceivers for SDR Systems. In *Proceedings of the 10th ACM International Symposium on Mobile Ad Hoc Networking and Computing (MobiHoc '09)*, 2009.
- [5] M. Fathi, H. Taheri, and M. Mehrjoo. Cross-Layer Joint Rate Control and Scheduling for OFDMA Wireless Mesh Networks. *IEEE Transactions on Vehicular Technology*, 59(8):3933–3941, Oct. 2010.
- [6] S. Hoteit, S. Secci, R. Langar, and G. Pujolle. A Nucleolus-based Approach for Resource Allocation in OFDMA Wireless Mesh Networks. *IEEE Transactions on Mobile Computing*, 2012.
- [7] S. Karande, Z. Wang, H. Sadjadpour, and J. Garcia-Luna-Aceves. Optimal Scaling of Multicommodity Flows in Wireless Ad Hoc Networks: Beyond The Gupta-Kumar Barrier. In *IEEE MASS 2008*, pages 102 – 113, Oct. 2008.
- [8] S. Katti, D. Katabi, H. Balakrishnan, and M. Medard. Symbol-level Network Coding for Wireless Mesh Networks. In *Proceedings of the ACM SIGCOMM 2008 Conference on Data Communication*, 2008.
- [9] R. Klose, A. Loch, and M. Hollick. A Rapid Prototyping Framework for Practical OFDMA Systems using Software Defined Radios. In *Proceedings of the 2013 International Conference on Sensing, Communication, and Networking (SECON '13)*, 2013.
- [10] K.-D. Lee and V. Leung. Fair Allocation of Subcarrier and Power in an OFDMA Wireless Mesh Network. *IEEE JSAC*, 2006.
- [11] H. Lichte, S. Valentin, F. Eitzen, M. Stege, C. Unger, and H. Karl. Integrating Multiuser Dynamic OFDMA into IEEE 802.11a and Prototyping it on a Real-Time Software-Defined Radio Testbed. In *Proceedings of the 3rd International Conference on Testbeds and Research Infrastructure for the Development of Networks and Communities (TridentCom '07)*, 2007.
- [12] S. Singh, M. Shahbazi, K. Pelechrinis, K. Sundaresan, S. V. Krishnamurthy, and S. Addepalli. A Case for Adaptive Sub-carrier Level Power Allocation in OFDMA Networks. In *Proceedings of the 13th ACM international symposium on Mobile Ad Hoc Networking and Computing (MobiHoc '12)*, 2012.
- [13] M. Veyseh, J. Garcia-Luna-Aceves, and H. Sadjadpour. OFDMA Based Multiparty Medium Access Control in Wireless Ad Hoc Networks. In *IEEE ICC 2009*, pages 1–6, June 2009.
- [14] M. Veyseh, J. Garcia-Luna-Aceves, and H. Sadjadpour. Adaptive Diversity Based Spectrum Allocation in Single-Radio Wireless Ad Hoc Networks. In *IEEE MASS 2010*, pages 232–241, Nov. 2010.
- [15] M. Veyseh, J. Garcia-Luna-Aceves, and H. Sadjadpour. Cross-Layer Channel Allocation Protocol for OFDMA Ad Hoc Networks. In *IEEE GLOBECOM 2010*, pages 1–6, Dec. 2010.