15th IEEE Internation Symposium on a World of Wireless, Mobile and Multimedia Networks MWoWM, June 2014, Sydney, Australia

Corridor-based Routing: Opening Doors to PHY-Layer Advances for Wireless Multihop Networks

Adrian Loch and Matthias Hollick
Technische Universität Darmstadt
Secure Mobile Networking Lab
Mornewegstr. 32, 64293 Darmstadt, Germany
{adrian.loch, matthias.hollick}@seemoo.tu-darmstadt.de

Alexander Kuehne and Anja Klein
Technische Universität Darmstadt
Communications Engineering Lab
Merckstr. 25, 64283 Darmstadt, Germany
{a.kuehne, a.klein}@nt.tu-darmstadt.de

***Abstract*—As of today, the performance of routing mechanisms in Wireless Multihop Networks (WMNs) is still limited by the lower layers. While recent cross-layer approaches take advantage of the characteristics of the medium, they are often based on traditional physical layers such as OFDM. State-of-the-art techniques used in one-hop scenarios, such as OFDMA or MIMO, pose a significant challenge in practical multihop networks, since typically Channel State Information (CSI) at the transmitter is required. Due to its volatile nature, disseminating timely CSI in the network is often infeasible. We propose Corridor-based Routing, which systematically and flexibly supports state-of-the-art physical layer techniques in WMNs. Instead of routing packets from node to node, we forward them along fully-connected groups of nodes. As a result (1) CSI only needs to be exchanged locally to enable cooperation of nodes in a group, and (2) groups can adaptively choose the best physical layer technique based on CSI. As a proof-of-concept, we implement the core mechanisms of Corridor-based Routing on software-defined radios and evaluate it in a multihop testbed using OFDMA as an exemplary physical layer technique. We show that our routing paradigm is feasible in practice, providing up to 2x throughput gain compared to routing agnostic to the lower layers.**

I. Introduction

While state-of-the-art physical layer techniques such as, for instance, Multiple-input Multiple-output (MIMO) or Orthogonal Frequency-Division Multiple Access (OFDMA) are widely studied for one-hop wireless communications, the multi-hop case has barely seen advances at the lower layers. Recent cross-layer approaches [1], [2] exploit the characteristics of the medium but typically build on traditional physical layers such as Orthogonal Frequency-Division Multiplexing (OFDM). Enabling advanced physical layers in Wireless Multihop Networks (WMNs) leads to a whole new class of cross-layer optimizations. Still, WMNs pose a fundamental problem to such state-of-the-art techniques, since often Channel State Information (CSI) is required at neighboring nodes. Disseminating CSI over multiple hops is often prohibitive—by the time it has been exchanged, CSI is typically outdated, since channel conditions change at a timescale orders of magnitude smaller than the forwarding of packets. Additionally, adaptation becomes crucial, as certain physical layer techniques might only be feasible under certain channel conditions. That is, the network must be able to support multiple mechanisms at the physical layer to keep up throughput for flows traversing areas with varying channel quality. This highlights the dimensions of the problem—supporting advanced physical layers in WMNs and allowing for transitions among them is a hard challenge.

There is a dearth of WMN routing mechanisms supporting advanced physical layers. Existing approaches are typically complex custom-built solutions which only support one specific physical layer, e.g., [3]. Moreover, they often focus on theoretical aspects, whereas the key challenge in such distributed scenarios frequently lies with the practical issues. A routing mechanism that systematically and flexibly supports in practice state-of-the-art physical layer techniques for WMNs as well as the resulting cross-layer interactions is missing.

In this paper we (a) propose a WMN routing paradigm which addresses the aforementioned challenges and, as a proof-of-concept, (b) we show its performance in practice on a software-defined radio testbed for the exemplary case of an OFDMA physical layer. The key to many current physical layer techniques is the exploitation of spatial diversity and the cooperation among nodes. Still, typical WMN routing mechanisms follow single paths from node to node, as inspired by wired networks. We propose to *widen* paths to *span multiple nodes per hop* in order to provide the foundation for spatial diversity. We call such a widened path a *corridor* and each widened hop a *stage*. Figure 1 shows an example. Essentially, each hop is formed by m transmitters and n receivers, which is a topology on top of which many state-of-the-art physical layer techniques fit. The routing layer builds corridors as a base structure to support such techniques. CSI only needs to be disseminated within each corridor stage. We refer to this routing paradigm as *Corridor-based Routing*.

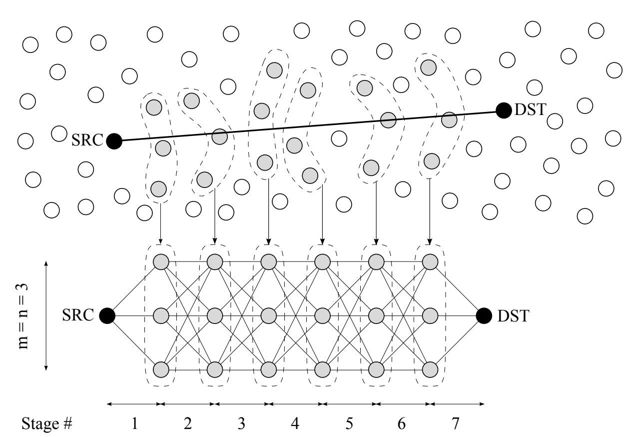

Fig. 1. Corridor example. For clarity we use a structured representation as shown in the lower part, but it can be directly derived from an actual network.

We choose OFDMA as a selected technique to show the benefits of corridors in practice. As opposed to traditional routing, at each hop there is not just one link available for forwarding data, but $m \times n$ links. Hence, OFDMA can be used to allocate each subcarrier to the link which provides best quality in terms of signal-to-noise ratio (SNR), which in turn translates into lower bit error rates (BERs) and ultimately higher throughput. We use OFDMA as an example, but we introduce the general notion of Corridor-based Routing for any physical layer technique. Our contributions are as follows:

1) We propose a routing paradigm to exploit state-of-the-art physical layer techniques in WMNs.
2) As a proof-of-concept, we show how the protocol can support OFDMA in a wireless multihop network.
3) We implement multihop OFDMA on a software-defined radio testbed and analyze the gains in practice.

In Section II we present motivating examples which provide an insight of the advantages of using corridors. We survey related work in Section III and then introduce Corridor-based Routing in Section IV. In Section V we show how corridors can be used for OFDMA. Section VI gives an overview of our implementation and in Section VII, we discuss our results. Finally, Section VIII concludes the paper.

II. Motivating features

OFDMA. While a corridor may look similar to a group of multiple paths, it enables optimizations which are infeasible with multi-path approaches. For example, when using OFDMA on top of such a corridor, data is not just forwarded along the best link of a stage, but is sent *over all links at the same time*. Nodes decide locally on how to allocate subcarriers to links to maximize performance. At each stage, data switches subcarriers according to local link conditions and hence *dynamically splits and joins* at nodes as it is forwarded through the corridor.

Combining physical layer techniques. Existing cross-layer approaches typically require all nodes of a network to use the same scheme to forward packets. On the contrary, corridors enable the network to use different physical layer techniques *at each stage*. Hence, each stage can locally choose the mechanism which best suits the current channel conditions, thus providing high adaptability and flexibility. A stage only needs to coordinate with its neighboring stages in order to guarantee that the used scheme is compatible to the mechanisms in use at its neighbors. For example, a node should not receive more data at stage i than what it can forward at stage $i+1$.

Timescale adaptation. The wireless medium is often considered to be random and change continuously. However, some variations may affect performance at significantly different *timescales* than others. For example, in an indoor setting, a person walking by produces quick *short-term* link quality variations. Still, the overall *long-term* trend of link characteristics due to walls typically dominates the communication quality. Corridors allow for appropriate adaptation to both types of variations. Short-term CSI changes are captured by local stage measurements and may trigger a stage to switch the physical layer technique in use. This allows a quick reaction to such a change. In contrast, when long-term variations are registered, the corridor itself is adapted to avoid the affected area, similarly to how paths are switched in multi-path approaches.

III. Related work

Research in WMNs has moved from dealing mainly with the network layer to the need of considering also lower layers [4], [5], [6], [7]. While **cross-layer** approaches exploit the characteristics of the wireless medium, they typically assume physical layers similar to the 802.11 standards. For instance, opportunistic routing [4] achieves gain by selecting the node which receives a packet best as a forwarder. Random network coding [6] also builds on the broadcast nature of wireless to opportunistically combine packets according to the data overheard by neighboring nodes. Partial packet recovery [7] reduces overhead by allowing nodes to retransmit only missing data instead of full packets in case of corrupted frames. A powerful synergy of these techniques is harvested in protocols such as MIXIT [1]. Still, the underlying physical layer is OFDM in all cases. Our work stands apart from such cross-layer schemes since we aim at allowing by design for virtually any physical layer technique. For the specific case of **OFDMA**, there exist approaches which explore how to combine it with WMNs. On the one hand, existing work analyzes subchannel allocation algorithms [8], [9] to find the optimal resource allocation in a WMN. On the other hand, some authors deal with building complete systems to exploit OFDMA in WMNs [10], [3], but do not take routing into account. OFDMA itself has been widely studied for cellular networks [11], but we (a) focus on the multihop case and (b) consider corridors for generic techniques beyond OFDMA. We do not try to improve the already largely studied domain of resource allocation, but only use OFDMA as an exemplary mechanism to illustrate the benefits of Corridor-based Routing. This includes **concurrency at the physical layer**, i.e., allowing simultaneous transmissions of different nodes. Concurrency is the underlying technique to many state-of-the-art lower layer techniques and has proven [12], [13] to overcome the well known limitations described in [14]—namely, that throughput in WMNs does not scale with network size. Enabling physical layer techniques in WMNs to circumvent this limitation motivates our research on Corridor-based Routing.

In previous work, we present the notion of **corridors** [15], [16] and analyze in theory whether they can be beneficially combined with OFDMA. In particular, in [15] we optimize the network throughput by means of resource and power allocation in a given corridor with up to five stages, whereas in [16] we devise a scheme to select the nodes of a random network topology that shall participate in the corridor. Still, we do not investigate in neither of both cases the operation of Corridor-based Routing itself, but assume a routing protocol that establishes, maintains and coordinates the corridor. Our contribution in this work is designing the architecture of such a protocol. A concept similar to corridors was introduced by Gui et al. [17], but they consider traditional multipaths within the corridor instead of OFDMA. Such closely grouped multipaths often have to deal with self-interference among paths, which is not the case for OFDMA. In earlier work [18], the same authors analyze the outage performance of OFDMA in a topology analog to a corridor, but as opposed to our previous work, they do not find optimal solutions for subchannel and power allocation. While the aforementioned papers deal with corridor-like structures, they do not consider routing protocols designed to operate a WMN based on corridors. We believe this is the first work in this direction.

IV. Corridor-based Routing

In this section, we present our first contribution, i.e., the design and operation of a protocol for Corridor-based Routing. Figure 2 gives an overview of our scheme.

A. Corridor construction

In a first step, the corridor itself must be built, similarly to the establishment of a routing path in traditional approaches. The goals are to (a) find a suitable corridor placement in the network, (b) decide which nodes are part of the corridor, and (c) organize nodes in stages. This can be achieved by extending well-known topological or geographical routing mechanisms. In the **topological** case, routing protocols like AODV or DSR can be used to find a hop-by-hop path to the destination. After the initial setup, each hop is widened by adding neighboring nodes. The corridor width is given as the number of nodes each hop shall add. Protocols using corridor-like alternative paths at each hop are well known (e.g., [19], [20]), but do not exploit yet the interaction with the lower-layers. In the **geographical** case, trajectory-based forwarding [21] can be used to find a path along an arbitrary curve from source to destination. The curve is extended to a band whose width is given in geographic units. All nodes falling into that band are part of the corridor. As opposed to the topological case, the width in terms of nodes might vary, since it depends on node density. For both cases, the routing metric determines the corridor shape, e.g., for geographical routing the curve may avoid regions which only provide substandard service.

B. Stage maintenance

As a result of corridor construction, nodes are aware of (a) the corridor they are part of and (b) the stage they are involved into. The nodes of a stage periodically probe each other to locally maintain the control information required for data forwarding. Maintenance ranges from periodical "Hello" messages required in basic stage mechanisms, to full CSI measurements required in advanced physical layer techniques. The forwarding mechanism used at a certain stage is adapted according to the results of stage maintenance. Adaptation might include adjusting the parameters of the current forwarding mechanism, switching to a different mechanism or changing the nodes involved in the current stage. In this paper, we perform stage maintenance measurements to adjust the parameters of the OFDMA forwarding mechanism we implement as a proof-of-concept for Corridor-based Routing. Specifically, we periodically measure CSI to determine the subcarrier allocation at each stage (c.f. Section V).

C. Stage mechanisms

In order to forward data, each stage can choose any conceivable mechanism which can exploit its $m \times n$ structure. Stage mechanisms are defined by the control information they require (e.g., full/partial CSIT), the data assumed to be available at the transmitting nodes (e.g., all transmitters have all data or not) and the way in which data is delivered at the receiving nodes (e.g., all data is delivered to all receivers or data is distributed among receivers). For example, in the simplest case one of the m transmitters sends all data to one of the n receivers. Such a stage mechanism requires no CSIT, data only needs to be available at one transmitter, and is only received at one receiver. Note that this example only serves to illustrate the parameters of a stage mechanism and would provide no advantage compared to traditional routing. More elaborated mechanisms could be based on, e.g., Interference Alignment, distributed MIMO or Analog Network Coding.

D. Stage coordination

Depending on the chosen stage mechanism, **intra-stage** coordination might be required. For example, when using a mechanism based on OFDMA, stage nodes need to share CSI and decide which node shall send data on which subcarrier. Such coordination incurs overhead and might be needed each time stage maintenance reports a change in CSI. Still, it is only required locally within a stage. In Section VI, we present more details on how intra-stage coordination can be done efficiently. Moreover, **inter-stage** coordination might be needed. This type of coordination is less frequent than intra-stage coordination and is more lightweight in terms of exchanged control data, e.g., no CSI needs to be shared. It is used to ensure that neighboring stages employ stage mechanisms which are compatible to each other, e.g., one stage does not deliver more data to a node than the amount this node can forward in the next stage. In this paper we address this issue by using a mechanism which delivers the same amount of data to each node. Still, we need inter-stage coordination to track how data switches subcarriers, as mentioned in Section II.

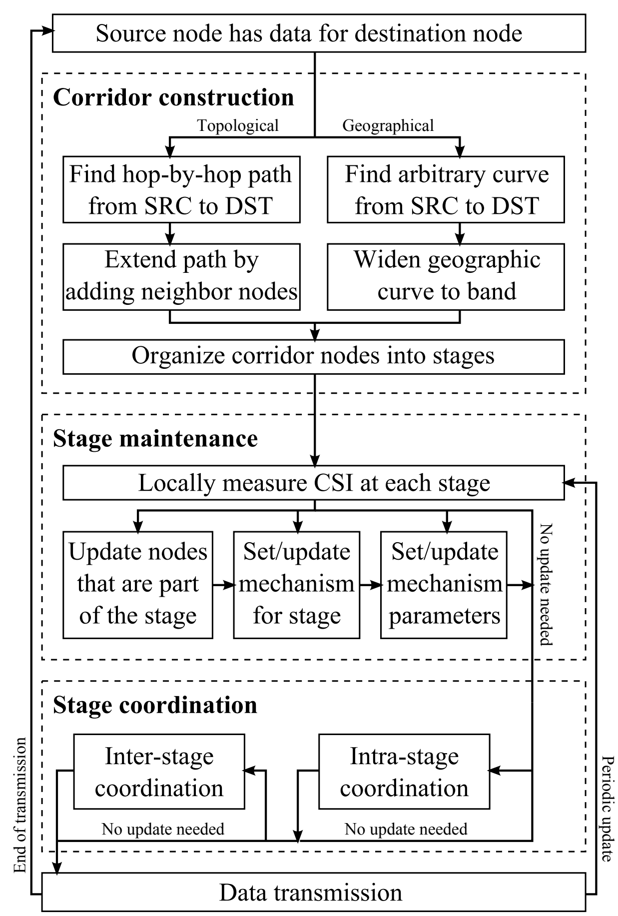

Fig. 2. Overview on Corridor-based Routing.

E. Data transmission

Finally, data is sent using the chosen stage mechanisms. The election of these mechanisms might depend on traffic type. If data is generated in small bursts and latency requirements impede buffering at stages, the stage coordination effort per burst may only be worth for mechanisms with small overhead. In contrast, if data can be buffered, the overhead might be negligible and thus more elaborated mechanisms can be used.

V. OFDMA FOR CORRIDORS

Next, we describe a stage mechanism for Corridor-based Routing based on OFDMA to improve performance in WMNs.

A. Scenario

Approach. OFDMA allows multiple nodes to transmit simultaneously without interference by assigning disjoint sets of subcarriers to each node. That is, instead of a single transmitter which uses all subcarriers to send data to a single receiver, there are m transmitters that *share* the available subcarriers to send data to n receivers. The gain stems from the assignment of subcarriers to links which experience good channel conditions, since this causes transmission errors to become less likely. As a result, the BER in a stage using OFDMA is lower than when using OFDM. Similarly to a traditional forwarding scheme, our OFDMA stage mechanism uses error correction codes and, if needed, retransmissions to obtain correct packets at each stage. In other words, errors do *not* accumulate as data flows through the corridor, but are corrected. While error correction is thus done per subcarrier instead of per frame, overhead is identical since the amount of redundancy of typical error correction codes is proportional to the amount of transmitted data.

Gains. Since OFDMA reduces the BER, it allows stages to operate at a higher modulation scheme compared to OFDM for a certain error correction capability. That is, if the BER of OFDMA for a certain modulation is similar to the BER of OFDM for a lower modulation, OFDMA can operate at the higher modulation and thus forward data significantly faster than an OFDM stage. For instance, if OFDM can only use BPSK while OFDMA can employ 4-QAM, the OFDMA stage only requires *half the time* to relay data. We use such an adaptive modulation scheme for our OFDMA stage mechanism—while all nodes within a stage use the same modulation on all subcarriers, modulations can be different at each stage. Additionally, channel quality may change throughout the corridor. Thus, if the first stage allowed for a high modulation order, subsequent stages with worse channel conditions and thus constrained to lower modulations may need to split data into multiple transmissions, incurring additional overhead. Here again, OFDMA achieves gains w.r.t. OFDM, as it typically needs less transmissions due to the use of higher modulations.

B. Forwarding scheme

We assume a stage with the same number of transmitters and senders, i.e., $m = n$. Our stage mechanism forwards data homogeneously, that is, it requires each of the m transmitters to have $1/m$ of the data and delivers $1/m$ of the data to each receiver. Alternatively, data could be distributed heterogeneously to nodes, but we choose the homogeneous variant because it ensures that in each stage all nodes need the same number of transmissions to forward data, thus simplifying scheduling and reducing coordination overhead. The goal of our mechanism is to distribute the available N_c subcarriers to the m^2 links of the stage. Note that subcarriers are allocated not just to transmitters but to *links*, since each outgoing link of a transmitter may have different quality. For example, while a certain subcarrier on the link from transmitter t_1 to receiver r_1 might be good, the same subcarrier from the same sender to r_2 might be very poor. Each transmitter and receiver has m outgoing and m incoming links, respectively. Our mechanism assigns to each link a $1/m^2$ fraction of the available subcarriers, i.e., each link gets the same share of resources. Thus, each transmitter gets a total of $m \times N_c/m^2 = N_c/m$ subcarriers and uses N_c/m^2 of them for each receiver. Conversely, each receiver receives N_c/m^2 subcarriers from each transmitter and thus gets overall the aforementioned $1/m$ fraction of the data. In other words, each transmitter sends to each receiver the same amount of data. Note that at each stage, data is shifted between subcarriers and thus becomes disordered. For example, data transmitted using subcarrier s_1 in stage i may be shifted to subcarrier s_2 in stage $i+1$ because the node transmitting the data in stage $i+1$ does not get subcarrier s_1. Thus, each stage must tell the next stage the current order of data to allow the destination to reorder it correctly. This control data is the inter-stage coordination overhead caused by our mechanism.

C. Subcarrier allocation

To allocate N_c/m^2 subcarriers to each link, we use an allocation strategy proposed in [22] and identified as best out of multiple allocation approaches. The strategy allocates the same number of subcarriers to each link, as required by our OFDMA scheme. We briefly summarize its operation. Essentially, the strategy follows an iterative approach. In each iteration, the strategy assigns one subcarrier to one link. The process is repeated until all subcarriers are allocated. Each iteration is divided into two steps. In a first step, the strategy chooses a *link* out of the m^2 available ones, excluding the links which already have their full share of subcarriers. More precisely, the strategy chooses the link which has the subcarrier with the *lowest* SNR. In a second step, the strategy assigns to the chosen link the subcarrier on which it experiences its *highest* SNR. The rationale behind this somewhat counter-intuitive approach is that the strategy tries to ensure that links with very bad SNR on some subcarriers at least get the N_c/m^2 subcarriers on which they perform best. Hence, the criteria used in the first step to choose a link defines which links are prioritized in the first iterations of the algorithm and thus can choose out of more still available subcarriers. If links with high SNR are prioritized, it might happen that in the last iterations the strategy is forced to assign subcarriers with very bad channel conditions to the links with low SNR. In [22], this approach clearly outperforms strategies which prioritize links with high SNR, specially when interference is present.

D. Operation

Our OFDMA scheme is a stage mechanism and thus does not need to deal with the construction of the corridor. Still, it requires stage maintenance and coordination in order to allocate subcarriers according to the aforementioned strategy. Figure 3 shows the frame format of our OFDMA mechanism, which can be divided into three main parts, namely (a) CSI measurement, (b) CSI and subcarrier distribution sharing, and (c) data transmission. We assume channels to be reciprocal, i.e., the channels are equal in both directions. Our measurements confirm this for our testbed environment. Hence, for the CSI measurement in (a) the *receivers* send pilot symbols to the *transmitters* one at a time. After this first step, each transmitter knows its m outgoing links, but does not know how the remaining $m^2 - m$ links are. The transmission order is defined during corridor construction, but can be adapted during stage maintenance.

The second step (b) deals with the intra- and inter-stage coordination. First, each transmitter shares its CSI with all other transmitters. In order to do this efficiently, we use a codebook approach, similarly to other technologies using OFDMA, such as LTE. Basically, a codebook is a list of quantized CSI values shared by all nodes of the stage. Hence, instead of encoding and sending a full CSI value, transmitters only need to share the index of a similar value of the codebook, which results in much less overhead. In Section VI, we discuss multiple codebook sizes and choose a suitable one for our purposes. In particular, we share CSI in terms of SNR on each subcarrier, since we use the SNR as a metric for allocation. Once all transmitters know the SNR on each subcarrier of all m^2 links, each transmitter can independently determine the subcarrier allocation according to the aforementioned strategy. This is the **intra-stage** coordination required by our OFDMA mechanism. Regarding the **inter-stage** coordination, the previous stage needs to tell the next stage how subcarriers are ordered, as discussed in Section V-B. To that end, after sharing the SNR values, the last transmitter broadcasts the current order to all receivers. Concretely, it sends on each subcarrier a sequence of bits representing the index number of the subcarrier on which the subsequent data on that subcarrier was originally located. For example, if subcarrier s_2 in the current stage contains the data which originally was sent over subcarrier s_1, the data on s_2 would be preceded by a binary representation of the index number s_1. The number of required bits is directly related to the number of subcarriers, i.e., it is $\lceil \log_2 N_c \rceil$, where $\lceil . \rceil$ denotes the ceiling function.

Finally, in (c) all transmitters send at the same time using OFDMA. We count as overhead all transmitted symbols which are not part of actual data, including pilot symbols. We send pilots prior to each transmission even if nodes might already know the channel to the transmitter from previous transmissions, since each transmission might be affected by a different phase offset in the frequency domain due to different symbol time offsets (STO) in the preamble detection.

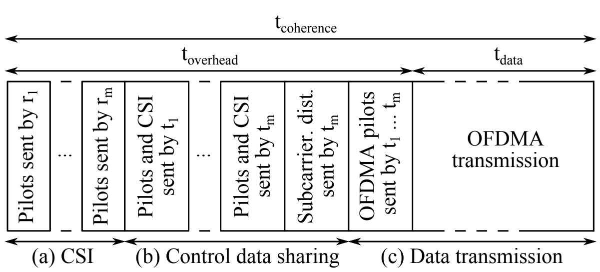

Fig. 3. OFDMA frame structure.

VI. Implementation

A. Testbed setup

Hardware platform. We use the Wireless Open-Access Research Platform (WARP), which is an FPGA-based Software Defined Radio (SDR) developed at Rice University [23]. It enables experiments in settings similar to an 802.11 network, but with full control regarding the lower layers. We use WARPLab, which is a framework that allows for rapid prototyping based on Matlab. First, we calculate in Matlab the samples to be transmitted. These samples are then transferred via Ethernet to the sending WARP board, which transmits them over the wireless medium. The receiving WARP board samples the signal and sends it back to Matlab. Note that in between frame parts (a), (b), and (c), depicted in Figure 3, data is processed in Matlab, i.e., while not fully real-time due to delays for transferring signals to and from Matlab, we do *not* process data offline, but online and *interactively*. This approach only relocates processing from the WARP board to Matlab.

Experiment setup. We carry out our experiments on ten WARP boards. Due to the limited testbed size, we cannot realize reasonable corridor construction, as we need all nodes to be part of the corridor to have a meaningful hop count. For simulative insights on corridor construction, we refer the reader to our results in [16]. Here, we focus on stage maintenance and coordination. We consider corridors with stage widths $m = [2, 4]$ but without start and end stages (c.f. Figure 1), as both are variants of a generic stage with fewer links. Figure 4 shows both setups. In the $m = 2$ case, each stage is formed by two nodes. Hence, with 10 boards the corridor would be limited to four stages. To look also into longer corridors, we take advantage of the two radios available on each board. We use each radio as if it were an individual node and thus only need one board per additional stage. All data sent and received is treated independently. In the $m = 4$ case, we follow a similar approach, using two boards per additional stage. We use nine boards for the corridor and one as an artificial interferer (c.f. Section VII-E). We place nodes in a regular pattern and in one room to better understand the results. We achieve typical indoor SNRs (20 to 30 dB) by using low transmit gains.

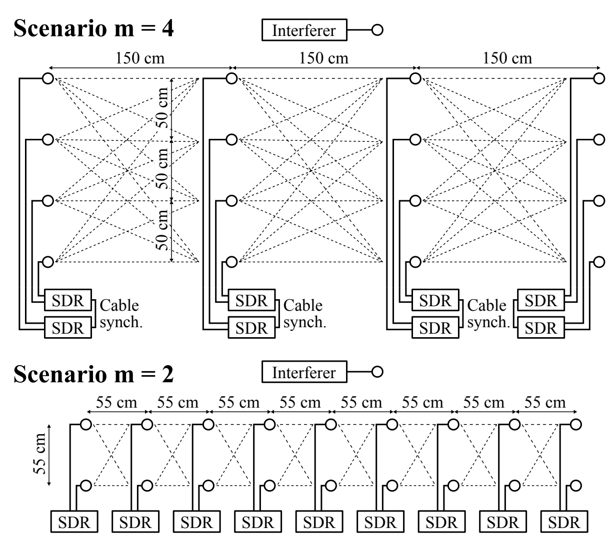

Fig. 4. Setup. The interferer is placed far enough to affect all stages similarly.

Synchronization. OFDMA requires transmitters to be synchronized. We achieve this either by having both transmitters on the same board for $m = 2$, or via wired synchronization for $m = 4$. While there exist mechanisms to achieve synchronization, we abstract from this issue to focus on the performance of the OFDMA stage mechanism. Additionally, receivers need to be synchronized to transmitters in order to (a) find out when to start receiving data and (b) avoid carrier-frequency offset (CFO). For (a), we prefix the signal with a preamble as in 802.11a. Receivers can determine the start by correlating the incoming signal with the preamble. To address (b), we use a pilot-aided technique to determine and compensate CFO.

B. Practical considerations

Coherence time. Our testbed is static and thus channels are stable, i.e., the coherence time is long. We exploit this to overcome the delays for transferring signals to/from Matlab. While the delays are in the order of *milliseconds*, in our testbed channels remain constant over *minutes*. Thus, CSI obtained in frame part (a) (Figure 3) is still up-to-date in part (c).

Acceptable BER. We use adaptive modulation at each stage according to CSI (c.f. Section V-A). Similarly to a real-world system, we aim at choosing the highest modulation which still achieves a certain acceptable BER. Typically, the acceptable BER is the error correction capability of the channel code in use. We do not restrict our results to certain codes, but analyze performance for a range of acceptable BER values. That is, we do not implement specific error correction schemes, but obtain how gains using *any* code with an error correction capability within the aforementioned range would be.

Throughput. The delays incurred by WARPLab prevent us from measuring throughput directly, since they would strongly affect the result. Moreover, the large coherence times in our testbed would lead to CSI measurements at much larger intervals than in a real-world scenario. To circumvent these limitations, we obtain throughput by extrapolating our measurements. We consider an indoor scenario and assume a realistic coherence time for such a setting, namely, $\mathrm{t_{coherence}} \approx 45$ ms [24]. For calculations, we assume that the first stage of the corridor sends a frame (Figure 3) as long as the coherence time before transmission continues in the next stage, but in practice the frame can be divided into multiple smaller packets. In subsequent stages, the transmission may take less time if the stage can use a higher modulation order than the first stage. Conversely, it also may take longer if only a lower modulation order is supported. In that case, multiple frames are needed, since each frame must not be longer than the coherence time. We compute the end-to-end corridor throughput “thp” as the transported data divided by the sum of all stage transmission durations, i.e., $\mathrm{thp} = \mathrm{tx_{data}}/\mathrm{sum_{time}}$. We calculate how much data is transmitted in the first stage based on CSI, acceptable BER and coherence time as follows.

$$\mathrm{tx_{data}} = \mathrm{bits_{S1}} \frac{\mathrm{t_{coherence}} - \mathrm{t_{overhead}}}{\mathrm{t_{measure}}} \quad (1)$$

where “$\mathrm{bits_{S1}}$” is the number of bits sent in the first stage using the highest modulation order possible according to the acceptable BER during the measurement duration $\mathrm{t_{measure}}$. The transmit buffers of WARPLab are not large enough to send data during the whole coherence time, i.e., $\mathrm{t_{measure}} < \mathrm{t_{coherence}}$. Thus, we extrapolate our measurement as shown in Equation 1 by calculating how many times $\mathrm{t_{measure}}$ fits into $\mathrm{t_{coherence}}$, after subtracting the time required for overhead $\mathrm{t_{overhead}}$. Next, we compute the time required to transport data through the corridor. Essentially, at each stage we calculate the number of transmissions f_m required according to the highest modulation order the stage supports. f_m may be smaller than one if the current modulation order is larger than in the first stage. Moreover, for each transmission in the stage we add the time required for overhead. To account for the number of times the overhead needs to be added at a stage, we define $f_o = \max[1, f_m]$. We compute the sum of all stage transmission times as follows.

$$\mathrm{sum_{time}} = \sum_{\forall \mathrm{stages}} f_o \cdot \mathrm{t_{overhead}} + f_m \cdot (\mathrm{t_{coherence}} - \mathrm{t_{overhead}}) \quad (2)$$

Gain control. Gain control in OFDMA is challenging. The complexity is due to the overlapping at the receiver of multiple signals *in time*, which can be decomposed *after* quantization using the Fourier Fast Transform (FFT). Still, *before* quantization the receiver can only operate on the sum. If one signal arrives with more power than others, the smaller ones suffer from higher quantization noise. The key problem is that the receiver can only adjust the largest signal to the input range of the quantizer, while all others are sampled with less accuracy. The impact of this issue can be observed in Figure 5, which depicts the BER for 64-QAM and $m = 2$ at subsequent stages for our OFDMA mechanism as well as for an OFDM baseline, which forwards data in a traditional hop-by-hop manner (c.f. Section VII-A). Additionally, we show the performance of a *sequential* OFDMA variant, which allows nodes to send in sequence on the subcarriers allocated to them. Hence, signals do not overlap in time and thus the gain control issue is circumvented. Note that this is *not* how we envision the scheme to operate once deployed, but allows us to illustrate the impact of gain control. As shown in Figure 5, the impact of gain misadjustments varies for each stage, since it is directly dependent on the physical environment surrounding the stage. We observe that simultaneous OFDMA performs worse than sequential OFDMA, specially at stage 5. To solve this problem, transmitter gains must be adjusted to ensure that all signals are received with similar power at the receivers. This is intrinsic to OFDMA and orthogonal to the gains achievable by subcarrier allocation. Hence, we do not tackle it in our implementation. To obtain the actual subcarrier allocation gain, for the bulk of our experiments we show the sequential OFDMA results and compute throughput as if transmissions were simultaneous.

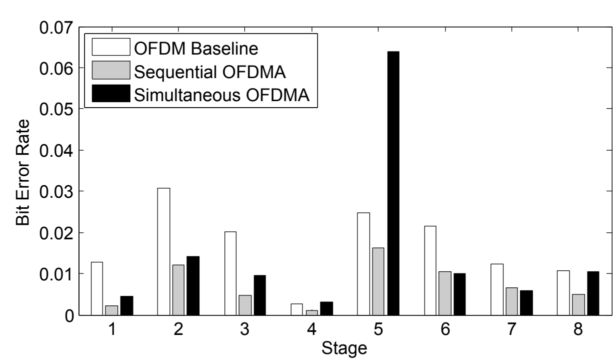

Fig. 5. Gain misadjustments. BER for 64-QAM at each corridor stage.

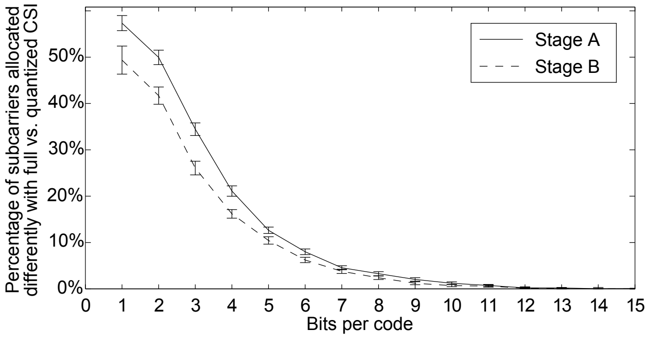

Fig. 6. Impact of quantized CSI on subcarrier allocation for two stages A/B.

TABLE I. BER FOR 64-QAM WITH INCREASING CODEBOOK SIZE.

Codebook size	16 codes	64 codes	512 codes	8192 codes
OFDMA BER	0.0179	0.0170	0.0146	0.0139

Feedback. We use quantized CSI feedback and account for the resulting overhead. To find a suitable codebook size, we measure how many bits per CSI value are needed to obtain the same subcarrier allocation than with full CSI. Figure 6 depicts the result for two stages of a $m = 2$ corridor. Both behave similarly, which means that the required codebook size does not depend on the specific channels of a stage. We choose a default value of 13 bits per code—8192 codebook entries—to achieve virtually identical allocations compared to full CSI. However, Table I shows that the BER achieved with much smaller codebooks is similar. While allocations using smaller codebooks are different than with full CSI, subcarriers allocated differently have similar performance. Hence, larger codebooks only provide marginal improvement. Still, we stick to 13 bits per code to show that overhead is reasonable even in that case. For smaller codebooks, the gained extra bits can be used to e.g. protect feedback against interference.

VII. EVALUATION

In this section we evaluate our OFDMA stage mechanism for Corridor-based Routing. Table II shows the parameters of our experiments and Table III gives an overview of our results.

A. Baseline mechanism

We compare OFDMA with traditional hop-by-hop forwarding based on OFDM. We use the same scenario than OFDMA, but choose randomly a node of each stage as a forwarder. Hence, data is relayed from node to node instead of from stage to stage. While pilot symbols are still needed to equalize the channel at the receiver, coordination and feedback are not required, which translates into smaller overhead.

B. OFDMA corridor gain

Our first experiment aims at showing the throughput gain in an OFDMA corridor. We expect the gain to stem from OFDMA being able to choose higher modulations than OFDM for a certain error correction capability. In Figure 7 we depict the time required to transport data through our $m = 4$ corridor for any acceptable BER up to 10%. We "normalize" the result dividing the time by the amount of sent data to highlight that higher modulations transport each single bit faster—this is

TABLE II. EXPERIMENT PARAMETERS

FFT size	256	**Subcarrier spacing**	256.25 kHz
Usable subcarriers	112	**Passband bandwidth**	17.5 MHz
Symbol duration	6.4 μs	**Pilot symbols**	3
Guard space	0.8 μs	**Data symbols**	10

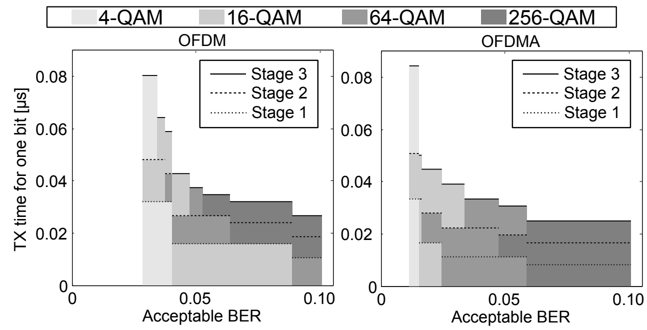

Fig. 7. Normalized transmission times through the corridor for $m = 4$.

not apparent in the non-normalized results, since with higher modulations we send more data. The gray tones indicate the modulation used in each stage. For example, for OFDM and 10% acceptable BER, the first stage uses 64-QAM, while the second and the third use 256-QAM. The transmission time decreases as expected with increasing acceptable BER, since the more errors can be corrected, the higher modulations are possible and thus the faster data is transported. For any acceptable BER, OFDMA can use higher modulations and thus requires less time than OFDM to transport data, which directly translates into throughput gain. For small acceptable BER values, both OFDM and OFDMA cannot operate, since the BER is too high even for the lowest modulation scheme in at least one of the stages. This is shown by the left white areas in Figure 7. However, note that OFDMA can already operate at about 0.012 acceptable BER, while OFDM requires at least 0.029 error correction capability to become feasible. For our $m = 2$ corridor results are equivalent, but we do not reproduce them here due to space constraints.

In Figure 8 we show the throughput gains resulting from our observations in Figure 7. The $m = 4$ corridor *doubles* throughput for certain acceptable BER values and the $m = 2$ corridor achieves up to 1.4x gain. We conclude that gains increase with corridor width—the wider the corridor, the more links there are and thus the higher is the probability that OFDMA can allocate links with good channel conditions to each subcarrier. The curve for $m = 4$ starts at 0.029 acceptable BER since OFDM cannot operate for lower values and thus no gain can be computed. That is, gain would be *infinite* in the range from 0.012 to 0.029 acceptable BER. For $m = 2$, we depict the throughput gain at stage 3—to make it comparable to $m = 4$—and at stage 8. Both curves behave similarly, which means that longer corridors do not strongly affect gains.

C. Per stage BERs

Since gains stem from adaptive modulation being able to choose higher modulations when using OFDMA due to lower BERs, we next analyze BERs for multiple modulations at each stage of our $m = 4$ corridor. Figure 9 shows our results. As expected, OFDMA generally achieves lower BERs for a certain modulation scheme than our OFDM baseline. Still, the

TABLE III. EXPERIMENT OVERVIEW

Experiment	Results	§
OFDMA corridor gain	We achieve 1.4x throughput gain for a corridor of width $m = 2$ and 2x gain for $m = 4$, i.e., gains increase with corridor width.	VII-B
Per stage BERs	OFDMA reduces the BER differently depending on the SNR requirements of the modulation and the physical environment of the stage.	VII-C
Low SNRs	At lower SNRs, we achieve up to 2.3x throughput gain and OFDMA can operate at lower error correction capabilities than OFDM.	VII-D
Narrowband interference	Subcarrier allocation allows OFDMA to avoid narrowband interference. As a result, its BER degrades less compared to OFDM.	VII-E

behavior depends on (a) the modulation scheme and (b) the stage. Regarding (a), we observe that the BER improvement for 16-QAM is limited or even non-existent. The reason for this is that in our testbed the SNR requirements for 16-QAM are low enough for OFDM to operate without incurring large BERs, i.e., channels are good enough for OFDM to work correctly and thus OFDMA does not provide significant improvements. Results for 256-QAM are similar but for the opposite reason—in this case SNR requirements are very high and thus OFDMA can also only provide marginal improvement. The largest gains are achieved for modulation schemes which are at neither of both extremes regarding SNR requirements, which in this case is 64-QAM. As to our aforementioned observation (b), the specific stage also influences the achieved improvement. For instance, in Figure 9, BER improvements at stage 3 become marginal compared to the previous stages. The reason lies with the specific channels at that stage—if channel conditions are strongly impaired for all subcarriers, subcarrier selection becomes less effective. Moreover, our allocation mechanism assigns the same number of subcarriers to each link in a stage (c.f. Section V-C) and thus is forced to assign resources to a link even if it has bad quality, which results in high BERs.

D. Low SNRs

In the following, we investigate the behavior of our OFDMA stage mechanism for low SNRs. However, the close positioning of antennas in our testbed impedes us to lower SNRs below ≈20 dB by adjusting transmit gains. Hence, we use the interferer depicted in Figure 4 to generate a small amount of white noise on all subcarriers and thus artificially achieve a lower SNR. Figure 10 depicts our results for three transmit gains at the interferer. As expected, for larger noise values—that is, lower SNRs—the acceptable BER required to achieve throughput gains increases. Still, the gains themselves are similar, reaching up to 2.3x, and follow virtually the same pattern, i.e., in Figure 10 the curve is just shifted to the right. Note that the acceptable BER at which each curve starts is the minimum BER at which *OFDM* can operate, but OFDMA works already at significantly lower BERs represented by the vertical dashed lines. Hence, while the 2.3x throughput gain only becomes possible at large BERs, a key advantage of OFDMA at low SNRs is that it can operate at significantly lower BERs than OFDM, potentially enabling communication at all if the channel code in use cannot cope with large BERs.

E. Narrowband interference

Finally, our last experiment deals with narrowband interference. In contrast to Section VII-D, we now use the interferer to generate noise only on a certain fraction of subcarriers. Our goal is to analyze whether subcarrier selection allows stages to *avoid* subcarriers affected by narrowband noise and thus achieve a better robustness against external interference, which is likely to occur in ISM bands. We set low transmit gains at the interferer to avoid jamming completely the affected subcarriers. As a result, nodes of a stage closer to the interferer are more affected by noise than nodes further away. In each stage, we expect our OFDMA mechanism to assign the subcarriers impacted by noise to the links which are less affected and thus degrade less for increasing noise bandwidth than OFDM. In Figure 11, we focus on the BER for 64-QAM at the first stage of our $m = 4$ corridor to investigate this effect. We observe that the behavior matches our expectations—the curves fitted on top of the results highlight that the OFDMA BER degrades at a lower rate than the the OFDM BER, hence providing a better robustness against external interference.

F. Discussion

Our results show that OFDMA for Corridor-based Routing is **feasible in practice** and improves throughput significantly. The **overhead** is not critical, since we achieve large gains while accounting for stage maintenance and coordination including virtually full CSI feedback. Moreover, we observe that coarse CSI feedback provides similar improvements. That is, gains do not stem from a fine-granular classification of subcarriers according to their channel conditions, but just from coarsely identifying subcarriers with very bad channel conditions. Further, corridors using OFDMA **enable communication** at SNR

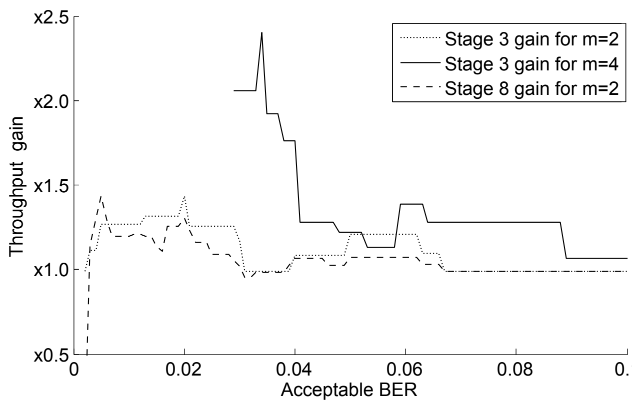

Fig. 8. Throughput gain for corridors of widths $m = 2$ and $m = 4$.

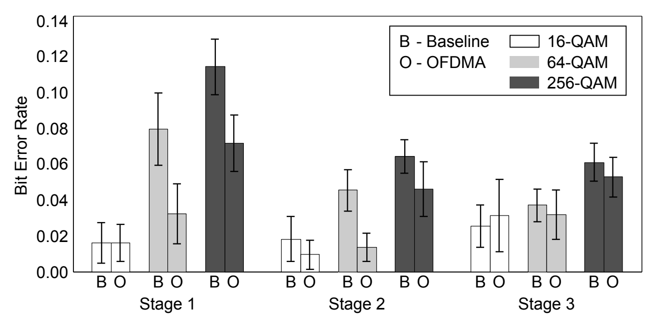

Fig. 9. BERs at each stage of our $m = 4$ corridor for multiple modulations.

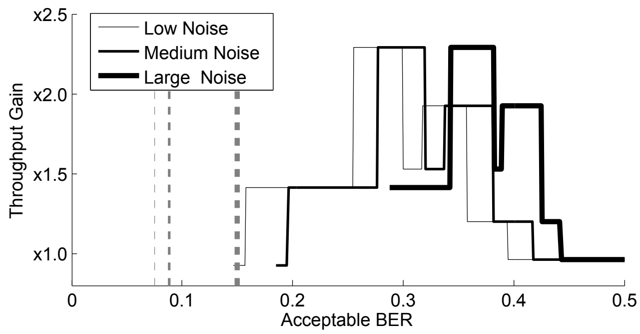

Fig. 10. Throughput gain with increasing noise for $m = 4$. The vertical dashed lines indicate the acceptable BER from which on OFDMA is operable.

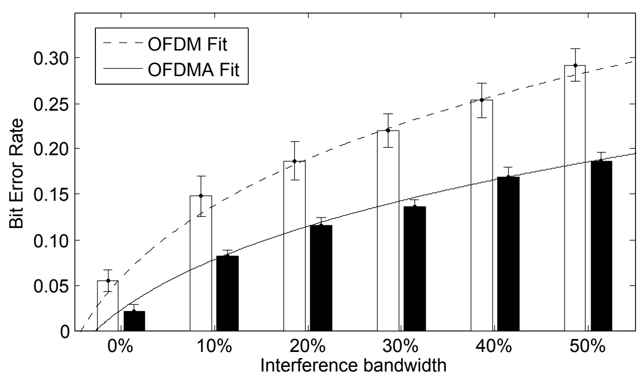

Fig. 11. BER for 64-QAM and increasing noise bandwidth (stage 1; $m = 4$).

values at which OFDM cannot operate at all, which shows its robustness. The **corridor shape** influences performance of our OFDMA stage mechanism regarding width but not length—gains become larger for wider corridors but are similar for increased hop counts. Still, we expect a trade-off w.r.t. width, since from a certain width on additional links per stage only improve diversity slightly. Finally, additional transmissions in a stage due to **adaptive modulation** choosing low modulations are not critical, since the resulting overhead is greatly compensated by other stages being able to choose high modulations.

VIII. Conclusion

We present a routing paradigm for Wireless Multihop Networks that supports state-of-the-art physical layer techniques. It widens traditional hop-by-hop paths in order to span a group of nodes at each hop. We call such a path a corridor and two subsequent groups a stage. A stage mechanism is a technique which forwards data from one group to the next by exploiting spatial diversity. As a proof-of-concept, we design and implement a stage mechanism based on OFDMA. This mechanism achieves performance gains in terms of throughput by allocating subcarriers to the links which provide good channel conditions. We implement it on software-defined radios and evaluate it in a multihop testbed. We achieve up to 2x throughput gain compared to traditional hop-by-hop forwarding. We observe that gains increase with stage width and that OFDMA can efficiently avoid subcarriers affected by interference. Future work includes designing stage mechanisms based on further physical layers and combining different stage mechanisms as well as corridor widths in one corridor.

References

[1] S. Katti, D. Katabi, H. Balakrishnan, and M. Medard, "Symbol-Level Network Coding for Wireless Mesh Networks," in *Proc. of SIGCOMM '08*, 2008.

[2] K. Jamieson and H. Balakrishnan, "PPR: Partial Packet Recovery for Wireless Networks," in *Proc. of SIGCOMM '07*, 2007.

[3] M. Veyseh, J. Garcia-Luna-Aceves, and H. Sadjadpour, "Adaptive Diversity Based Spectrum Allocation in Single-Radio Wireless Ad Hoc Networks," in *Proc. of MASS '10*, 2010.

[4] S. Biswas and R. Morris, "ExOR: Opportunistic Multi-Hop Routing for Wireless Networks," in *Proc. of SIGCOMM '05*.

[5] S. Gollakota and D. Katabi, "ZigZag Decoding: Combating Hidden Terminals in Wireless Networks," in *Proc. of SIGCOMM '08*.

[6] S. Katti, H. Rahul, W. Hu, D. Katabi, M. Medard, and J. Crowcroft, "XORs in the Air: Practical Wireless Network Coding," *IEEE/ACM Transactions on Networking*, 2008.

[7] G. R. Woo, P. Kheradpour, D. Shen, and D. Katabi, "Beyond the Bits: Cooperative Packet Recovery Using Physical Layer Information," in *Proc. of MobiCom '07*, 2007.

[8] S. Hoteit, S. Secci, R. Langar, and G. Pujolle, "A Nucleolus-based Approach for Resource Allocation in OFDMA Wireless Mesh Networks," *Transactions on Mobile Computing*, 2012.

[9] K.-D. Lee and V. Leung, "Fair Allocation of Subcarrier and Power in an OFDMA Wireless Mesh Network," *IEEE Journal on Selected Areas in Communications*, 2006.

[10] M. Veyseh, J. Garcia-Luna-Aceves, and H. Sadjadpour, "OFDMA Based Multiparty Medium Access Control in Wireless Ad Hoc Networks," in *Proc. of ICC '09*, 2009.

[11] A. Richard, A. Dadlani, and K. Kim, "Multicast Scheduling and Resource Allocation Algorithms for OFDMA-Based Systems: A Survey," *IEEE Comm. Surveys & Tutorials*, 2012.

[12] J. Garcia-Luna-Aceves, H. Sadjadpour, and Z. Wang, "Challenges: Towards Truly Scalable Ad Hoc Networks," in *Proc. of MobiCom '07*, 2007.

[13] 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 *Proc. of MASS '08*, 2008.

[14] P. Gupta and P. Kumar, "The Capacity of Wireless Networks," *IEEE Transactions on Information Theory*, 2000.

[15] A. Kuehne, A. Klein, A. Loch, and M. Hollick, "Corridor-based Routing Using Opportunistic Forwarding in OFDMA Multihop Networks," in *Proc. of PIMRC '12*, 2012.

[16] A. Kuehne, A. Loch, M. Hollick, and A. Klein, "Node Selection for Corridor-based Routing in OFDMA Multihop Networks," in *Proc. of WCNC '13*, 2013.

[17] B. Gui, L. Dai, and L. Cimini, "Routing Strategies in Multihop Cooperative Networks," *IEEE Transactions on Wireless Communications*, 2009.

[18] L. Dai, B. Gui, and L. Cimini, "Selective Relaying in OFDM Multihop Cooperative Networks," in *Proc. of WCNC '07*, 2007.

[19] S.-J. Lee, E. M. Belding-Royer, and C. E. Perkins, "Scalability Study of the Ad Hoc On-demand Distance Vector Routing Protocol," *International Journal of Network Management*, vol. 13, no. 2, 2003.

[20] S. J. Lee, W. Su, and M. Gerla, "On-demand Multicast Routing Protocol in Multihop Wireless Mobile Networks," *Mobile Networks and Applications*, vol. 7, no. 6, 2002.

[21] D. Niculescu and B. Nath, "Trajectory Based Forwarding and its Applications," in *Proc. of MobiCom '03*, 2003.

[22] R. Klose, A. Loch, and M. Hollick, "Evaluating Dynamic OFDMA Subchannel Allocation for Wireless Mesh Networks on SDRs," in *Proc. of the 2nd Workshop of Software Radio Implementation Forum (SRIF 2013)*, 2013.

[23] "Rice University WARP Project." [Online]. Available: http://warp.rice.edu

[24] A. Loch, T. Nitsche, A. Kuehne, M. Hollick, J. Widmer, and A. Klein, "Practical challenges of IA in frequency," TU Darmstadt, SEEMOO, Tech. Rep., 2013. [Online]. Available: http://www.seemoo.de/ia