

A Scalable Dual-Radio Wireless Testbed for Emulating Mesh Networks

Sherif M. ElRakabawy, Simon Frohn, and Christoph Lindemann

Abstract

In this paper, we introduce and evaluate ScaleMesh, a scalable miniaturized dual-radio wireless mesh testbed based on IEEE 802.11b/g technology. ScaleMesh can emulate large-scale mesh networks within a miniaturized experimentation area by adaptively shrinking the transmission range of mesh nodes by means of variable signal attenuators. To this end, we derive a theoretical formula for approximating the attenuation level required for downscaling desired network topologies. We conduct a comprehensive performance study, in which we validate the feasibility of ScaleMesh for network emulation and protocol evaluation. Among others, we study the effect of channel selection, signal attenuation level, different topologies, and traffic load on network performance. We particularly focus on the performance of single-radio versus dual-radio communication, while investigating key parameters which can provide a substantial improvement in performance. We show that dual-radio communication improves network goodput by up to 100%, yet does not overcome TCP's fairness problems over IEEE 802.11.

Keywords:

Design and implementation of wireless mesh testbeds, IEEE 802.11 wireless networks, performance evaluation, real-world experiments

1 Introduction

In recent years, wireless mesh networks [1] have been within the focus of research in the networking community. Such networks are becoming increasingly attractive, since they can provide cost-efficient Internet access with minimal infrastructure expenditure. While most of the research in this area is still conducted using network simulators such as ns-2 [12] and Qualnet [27], the trend is increasingly moving towards deploying such networks in reality. Examples include MIT Roofnet [2], TFA-Rice ([3], [15]), and Freifunk [13], which have proven the feasibility of wireless mesh networks.

Within this context, wireless testbeds can contribute significantly to research by providing a real-world

platform for implementing and evaluating next-generation network protocols. Such testbeds possess crucial advantages with respect to network simulators. The latter often rely on optimistic assumptions compared to the real world and, thus, do not always deliver accurate results. Moreover, many physical measures in reality, such as the distance between nodes in a network, can be simply inquired in simulations, but are not available at nodes in reality due to the absence of global knowledge. Hence, working with testbeds improves the feasibility and reliability of newly designed protocols.

In this paper, we introduce and evaluate ScaleMesh, a 20-node scalable dual-radio wireless mesh testbed based on IEEE 802.11b/g technology. Using ScaleMesh, large-scale mesh networks can be emulated within a miniaturized experimentation area by using variable signal attenuators. By adaptively shrinking the transmission range of mesh nodes, large-scale networks can be downscaled on an area of a few square meters. Different network topologies can then be emulated by adjusting the positions of the testbed antenna-stations. Such a dynamic network scaling combined with dual-radio support allows emulating and evaluating a large variety of wireless mesh networks. On the other hand, ScaleMesh can help planning the deployment of large-scale mesh networks by building and evaluating a miniaturized version of the aspired network within the testbed area before transferring the acquired knowledge to reality.

Furthermore, we derive the correlation between three fundamental measures, which are crucial for downscaling large-scale networks using ScaleMesh. Specifically, for emulating particular mesh networks, it is desirable to identify the correlation between the inter-node distance in the network to be emulated, the inter-node distance in the downscaled version within the testbed area, and the attenuation level of the transmission signal. This correlation makes it possible to approximate the level of attenuation required to downscale certain mesh networks on a desired miniaturized area.

We conduct a comprehensive performance study, in which we evaluate a wide variety of parameters in the testbed. Among others, we study the effect of channel selection, signal attenuation level, and traffic load on network performance. We particularly focus on the

*The authors are with the University of Leipzig, Department of Computer Science, Johannisgasse 26, 04103, Leipzig, Germany.
E-mail: {sme, sf, cl}@ros.informatik.uni-leipzig.de.
Tel: +49 341 97 2222, Fax: +49 341 97 32289*

performance of single-radio versus dual-radio communication, while investigating key parameters which can provide a substantial improvement in performance. We also conduct ns-2 [12] simulations, and compare the acquired results to the corresponding experiments in the testbed.

The remainder of this paper is organized as follows. Section 2 summarizes related work on real deployments of wireless mesh networks as well as testbed prototypes. Section 3 describes the architecture and operation of the introduced miniaturized wireless mesh testbed, whereas in Section 4 we identify the correlation between the inter-node distance in the network to be emulated, the inter-node distance in the downscaled version on the testbed area, and the attenuation level of the transmission signal. A comprehensive performance study of the testbed is presented in Section 5. Finally, concluding remarks are given.

2 Related Work

Bicket et al. [2] evaluated a 37-node 802.11b community mesh network over an area of approximately four square kilometers in Cambridge, Massachusetts. The mesh network, denoted as MIT Roofnet, adopts off-the-shelf equipment, e.g. IEEE 802.11 wireless cards and standard omni-directional antennas. The authors evaluated multiple aspects of the architecture such as the effect of node density on connectivity and throughput as well as the characteristics of wireless links.

Gambiroza et al. [15] simulated a multihop wireless backhaul network consisting of multiple Transit Access Points (TAPs), which are connected to the Internet through multiple entry points. They studied TCP/UDP fairness, while considering different parameters such as the role of the link layer protocol, antenna technology, and traffic types.

Based on the findings in [15], Camp et al. [3] deployed a two-tier mesh network in Houston, Texas, that aims at providing Internet access over a wide area with minimal infrastructure. The deployed network comprises an access tier and a backhaul tier. The access tier connects mobile clients with mesh nodes, whereas the backhaul tier interconnects the mesh nodes and forwards traffic to and from the Internet. Using this network, the authors presented a measurement driven deployment strategy and a data driven model to study the impact of design and topology decisions on network-wide performance.

Opposed to [2], [3], and [15], we introduce a scalable miniaturized mesh testbed rather than a large-scale mesh network. Using our testbed, networks such as [2], [3], and [15] can be emulated within a miniaturized experimentation area. Beyond [2], [3], and [15], our testbed further supports multiple radios rather than only a

single radio.

In [21], Raniwala et al. proposed a dual-radio wireless mesh network comprising 9 PC nodes, each equipped with two IEEE 802.11a interfaces. The authors show that, by employing sophisticated channel assignment approaches, network throughput can be significantly improved.

De et al. [7] proposed a mobile 12-node experimentation testbed for multihop wireless networks. Each node in the testbed comprises a wireless computing device and a mobile robot. Fixed signal attenuators are used to limit the transmission range of the mobile nodes.

In [11], Eriksson et al. evaluated the feasibility of an all-wireless office mesh network consisting of 21 multi-radio mesh nodes. The authors captured user traffic on office PCs with wired ethernet connectivity and replayed them on the mesh network. A set of parameters, such as different routing metrics and hardware settings were evaluated.

Raychaudhuri et al. [20] proposed an open access research testbed called Orbit for evaluating next-generation wireless network protocols. The testbed consists of an indoor radio grid emulator for controlled experiments and an outdoor field trial software for end user evaluations.

Lundgren et al. reported in [18] on their experience in designing and deploying the UCSB MeshNet, a 30-node wireless mesh testbed which covers several floors inside a building. In [23], Vaidya et al. discussed preliminary ideas towards building a single-radio shielded testbed for a repeatable evaluation of wireless protocols. The authors proposed to shield the testbed using an electromagnetic chamber. Nodes in the testbed shall comprise laptops with fixed attenuators and shielding copper tapes.

Similar to [7], [11], [18], [20], and [21], our testbed supports multiple radios. Opposed to [11], [20], [21], and [23] our testbed comprises variable attenuators to variably adjust the transmission range and thus flexibly emulate large-scale networks. Deploying fixed signal attenuators such as in [7] significantly limits the spectrum of network topologies which can be considered due to the fixed transmission range associated with the attenuators.

3 ScaleMesh Architecture

ScaleMesh is a miniaturized mesh testbed, which is mostly built using off-the-shelf hardware and software components. The testbed, which is depicted in Figures 1 and 2, is built on a 2m x 3m flat table pool in a 10m x 6m lab with four 15cm thick light-gypsum walls, and comprises 20 wireless mesh nodes. Each node consists of a Siemens ESPRIMO P2510 PC with an Intel Celeron 3.2 GHz processor and two IEEE 802.11b/g Netgear WG311T wireless PCI network interface cards (NICs) with Atheros chipsets. Each wireless card is connected



Fig.1: ScaleMesh

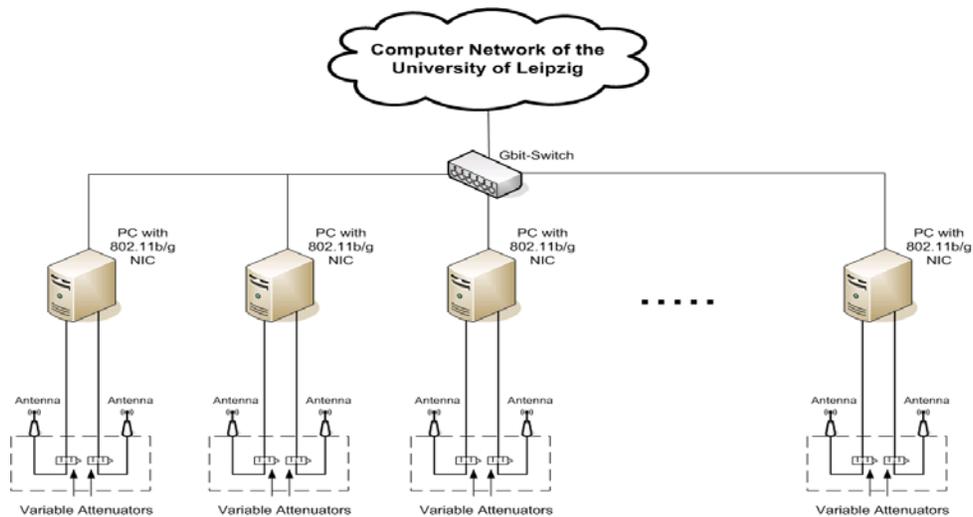


Fig. 2: Architecture of ScaleMesh

to a variable signal attenuator and a 2.1dBi low-gain antenna. Using the variable attenuators, the signal power of the wireless PCI cards can be adaptively shrunk in 1dB steps in order to limit the maximum transmission range of each node. Thus, large-scale wireless mesh networks can be scaled down on the testbed area, making quick topology and parameter modifications for efficient evaluation of network protocols possible. Adjusting the transmission power of the wireless NICs while omitting attenuators is not sufficient for an effective scaling, since the lowest adjustable transmission power of 0dBm (i.e. 1mW) still provides a transmission range of around 4m. Thus, networks of several hops would require an entire building floor to emulate. In such a scenario, a flexible adjustment of the mesh nodes for emulating certain topologies would be extremely hard, if not impossible.

Note that within this context, scalability refers to scaling the deployed environment of a network, not scaling its number of nodes.

The variable attenuators are connected to the wireless PCI cards through 50 Ohm, 7m long, highly shielded aircell5 coaxial cables, whereas the antennas are connected to the signal attenuators through a 50 Ohm, 3m long RG-174 coaxial cable. According to the technical specifications, both cables (i.e. from NIC to attenuator and from attenuator to antenna) add a total of 12.5dB signal attenuation.

Testbed nodes run a SuSE Linux 10.2 operating system with a custom-compiled kernel version 2.6.18 with the high-resolution timer subsystem patch [16]. As driver for the wireless PCI cards, we employ the Linux

TABLE 1: HARDWARE AND SOFTWARE COMPONENTS OF THE MINIATURIZED TESTBED

Hardware	
Component	Description
PC	Siemens ESPRIMO P2510 Celeron 3.2 GHz, 512 Mbytes RAM, 80 Gbytes HDD
Wireless NIC	Netgear IEEE 802.11b/g wireless PCI card WG311T with Atheros chipset
Variable attenuator	Broadwave 751-002-030 variable attenuator, attenuation range 0-30dB in 1 dB steps
Coaxial cable	7m aircell5 + 3m RG-174, 50 Ohm with SMA/RPSMA connectors
Antenna	Maldol mini 2.1dBi antenna with magnetic mount and 3m SMA cable
Software	
Component	Description
Operating System	SuSE Linux 10.2 with custom kernel version 2.6.18 with high resolution subsystem patch
Wireless NIC driver	Madwifi Linux kernel device driver for Atheros chipsets version 0.9.3.2
Multihop routing protocol	OLSR for Linux version 0.5.2 with ETX support

Madwifi kernel device driver version 0.9.3.2 for Atheros chipsets. All wireless cards operate in ad-hoc mode. Depending on the current scenario, we employ either static routing or the Optimized Link State Routing Protocol (OLSR) version 0.5.2 ([5], [26]) for multihop routing. This implementation of OLSR incorporates the Expected Transmission Count (ETX) metric [6] for selecting routes based on the current loss probability of the links. The ETX value of a link describes the expected number of transmissions (including retransmissions) required for sending a packet over that link. The ETX of a path is the sum of the ETX values of each link on that path. According to [6], ETX is given by

$$ETX = \frac{1}{D_f \cdot D_r} \quad (1)$$

where D_f denotes the measured probability that a data packet successfully arrives at the receiver, and D_r denotes the measured probability that the corresponding ACK packet is successfully received on the reverse path. Note that the more the ETX value converges to 1 the better is the quality of the link.

The idea of ETX has evolved because first generation multihop routing protocols had mainly used the hop count metric for determining best routes between source and destination. That is, the shortest route from source to destination was always chosen as the best one. However, as shown in [6], the shortest route is not always the best choice, since link interference and congestion along the route have a significant impact on its quality. In particular, a longer route may very well be more suitable in case it experiences less congestion and interference than a shorter one. Therefore, ETX aims at determining

routes based on their loss probabilities rather than on their hop count.

Each wireless node further possesses a Gigabit ethernet NIC, which is connected to the subnet of the University of Leipzig through a Gigabit switch. This allows a remote management of the wireless nodes from any wired host in the subnet. Hence, wireless experiments can be managed from a remote computer and traces can be copied and evaluated through the wired network. Table 1 shows a detailed description of hardware and software components of the miniaturized testbed.

ScaleMesh supports dual-radio communications by statically assigning a different channel to each of the two independent wireless PCI cards of a node. Channel assignment is performed according to the findings of section 5.2. Incorporating two independent wireless cards provides a better performance than using software-based channel switching techniques for single cards as proposed in [4]. Besides the different channels, each wireless card in a ScaleMesh node is assigned a different IP address. Which radio (i.e. wireless card) to use for a certain transmission is determined by the routing protocol OLSR. Specifically, OLSR uses the ETX values of both radios to determine which radio corresponds to the lowest packet loss probability. Such radio is then used for transmission, since it exhibits less packet drops and thus achieves more goodput. The IEEE 802.11b/g standard supports 11 different channels. According to the IEEE 802.11 specifications [24], channels 1, 6, and 11 are non-overlapping. However, in practice, non-overlapping channels strongly depend on the vendor of the corresponding network cards and may strongly vary. In particular, the authors in [14] showed through measurements in IEEE 802.11 that in practice, non-overlapping channels can interfere with each other due to the “near-far effect” of transmitting antennas. Thus, in Section 5.2 we study which channels show the least mutual interference, as well as how the mutual proximity of antennas can affect performance.

Mesh networks can be emulated using ScaleMesh by adjusting the positions of the antenna-stations according to the desired topology. An antenna-station is a joint magnetic board, on which every two antennas of each mesh node are mounted. Such antenna-stations define the logical structure of a mesh node. Since ScaleMesh is deployed in an indoor environment, the shadowing and fading characteristics of wireless signals correspond to the indoor propagation model [22], which takes into account reflections on walls and floors. For all-wireless office mesh networks as introduced in [11], these indoor shadowing characteristics are identical. For mesh networks operating in free space, different shadowing

characteristics apply. These different characteristics may well be considered using outdoor instead of indoor propagation models for downscaling mesh networks. While the signal-to-noise ratio in SaleMesh may not deliver one-to-one identical results as in a free space mesh network, the acquired results are representative due to the similar characteristics of the IEEE 802.11 wireless link (opposed to simulations). Furthermore, while a large-scale free space mesh networks has a fixed topology, nodes in ScaleMesh are variably adjustable, making it more convenient for evaluating network protocols.

4 Downscaling Mesh Networks

Using ScaleMesh, large-scale mesh networks can be scaled down to a small area of a few square meters. Such a feature makes it possible to emulate large-scale networks by adjusting the testbed antenna-stations to emulate arbitrary topologies. The key component which puts such a feature into practice is the variable signal attenuator which is connected to each wireless network card. We preferred variable signal attenuators over fixed ones, since variable attenuators allow variable transmission ranges, and thus more flexibility in deploying various network topologies. Fixed signal attenuators only provide a fixed level of attenuation, limiting the spectrum of topologies that can be deployed.

Within this context, it's crucial to acquire certain topology-dependant parameters in order to emulate particular mesh networks within the miniaturized area of the testbed. Specifically, for emulating particular mesh networks, it is desirable to identify the correlation between the inter-node distance in the network to be emulated, the inter-node distance in the downscaled version on the testbed area, and the attenuation level of the transmission signal. Such a correlation shall provide answers to questions such as: How much attenuation is required for scaling an inter-node distance of x meters down to a distance of y meters in the testbed area?

We denote the inter-node distance in the network to be emulated as $d^{non-scaled}$, the inter-node distance in the downscaled version on the testbed area as d^{scaled} , and the attenuation level of the transmission signal as Ω_{sum} . Consider a simple scenario, where two mesh nodes, A and B , communicate with each other over one hop. In this scenario, node A is the transmitter and node B is the receiver. The first step towards deriving a correlation between $d^{non-scaled}$, d^{scaled} , and Ω_{sum} is to approximate the signal attenuation between nodes A and B . Following the Equivalent Isotropically Radiated Power (EIRP) [22] equation the emitted transmission power at the antenna P_{out} is given by:

$$P_{out} = P_{tx} + G_{ant} \quad (2)$$

where P_{tx} denotes the transmission power of the wireless card at node A , and G_{ant} denotes the gain of the antenna.

The signal attenuation $L_{<A,B>}$ between nodes A and B is given by the difference between the received power P_{rx} at node B and the outgoing signal power from node A , minus the attenuation level of the signal:

$$L_{<A,B>} = P_{out} - P_{rx} - \Omega_{sum} \quad (3)$$

where

$$\Omega_{sum} = \Omega_{cab} + \Omega_v \quad (4)$$

Here, Ω_{cab} and Ω_v describe the signal attenuation caused by the coaxial cable and the attenuator. In ScaleMesh, Ω_{cab} is roughly 12.5dB and Ω_v is naturally variable.

The next step towards identifying the correlation between $d^{non-scaled}$, d^{scaled} , and Ω_{sum} is to derive the signal attenuation between nodes A and B as a function of d^{scaled} , which denotes the downscaled distance between A and B in meters. We refer to such distance as downscaled since, with respect to a normal mesh network without synthetic attenuation by variable attenuators or long cables, the distance between A and B is downscaled.

The general signal attenuation equation as described by the ITU-R indoor propagation model [22] is given by:

$$L = 20 \log_{10}(f_c) + 10p \log_{10}(d) \quad (5)$$

where f_c denotes the frequency of the transmitted signal, i.e. a channel in the 2.4 GHz band in our case, p denotes the path loss exponent, and d describes the distance between transmitter and receiver in meters. The path loss exponent p depends on the operating environment of the wireless nodes and ranges from 2 for propagation in free space up to 5 in dense indoor environments. Following the indoor propagation model in [17] and [22] and as validated in Section 5.1, we set $p = 3$ for our testbed.

Considering d^{scaled} as distance between nodes A and B , we get according to Eq. 5:

$$L_{<A,B>} = 20 \log_{10}(f_c) + 10p \log_{10}(d^{scaled}) \quad (6)$$

Substituting Eq. 3 into Eq. 6 and solving for d^{scaled} :

$$\begin{aligned} d^{scaled} &= 10^{\frac{P_{out} - P_{rx} - \Omega_{sum} - 20 \log_{10}(f_c)}{10p}} \\ &= 10^{\frac{P_{out} - P_{rx} - 20 \log_{10}(f_c)}{10p} - \frac{\Omega_{sum}}{10p}} \end{aligned} \quad (7)$$

In a non-scaled network Ω_{sum} equals 0, hence we get for $d^{non-scaled}$:

$$d^{non-scaled} = 10^{\frac{P_{out} - P_{rx} - 20 \log_{10}(f_c)}{10p}} \quad (8)$$

Substituting Eq. 8 into Eq. 7:

$$d^{scaled} = d^{non-scaled} 10^{-\frac{\Omega_{sum}}{10p}} \quad (9)$$

In other words, the scaled distance $d_{<A,B>}^{scaled}$ is the product of the non-scaled distance $d_{<A,B>}^{non-scaled}$ and a factor determined by the attenuation level Ω_{sum} . This delivers us the desired correlation between the inter-node distance in the network to be emulated, the inter-node distance in the downscaled version on the testbed area, and the attenuation level of the transmission signal. In case that the antenna gain of the node to be emulated differs from the testbed antenna, the difference in gain should be considered in the equation. Such a case is discussed in Section 5.1.

Solving for $d_{<A,B>}^{non-scaled}$ and Ω_{sum} , respectively, we get:

$$d_{<A,B>}^{non-scaled} = \frac{d_{<A,B>}^{scaled}}{10^{-\frac{\Omega_{sum}}{10p}}} \quad (10)$$

and

$$\Omega_{sum} = -\log\left(\frac{d_{<A,B>}^{scaled}}{d_{<A,B>}^{non-scaled}}\right)10p \quad (11)$$

Note that other signal propagation models (e.g. free space) may well be considered along with the ITU-R indoor propagation model, depending on the physical environment of the deployed mesh network.

5 Performance Study

We conduct a comprehensive performance study using ScaleMesh, in which we validate ScaleMesh's feasibility and evaluate the influence of different key parameters on network performance. In all experiments, except for ones showing transient behavior, we conduct steady-state experiments starting with an initially idle system. In each run, we activate TCP/UDP connections until 55,000 packets are successfully transmitted, and split the output of the experiment in 11 batches, each 5,000 packets in size. The first batch is discarded as initial transient. The considered performance measures are derived from the remaining 10 batches with 95% confidence intervals by the batch means method. The default TCP/UDP packet size is set to 1460 bytes, unless otherwise stated. Dependant on the respective experiment, we generate TCP/UDP traffic using the Iperf bandwidth measurement tool for Linux [25]. Unless otherwise stated, RTS/CTS is disabled, the inter-antenna distance is set to 10cm, and the link layer data rate is fixed to 54 Mbit/s (with rate adaptation disabled). This eliminates undesired effects that may be caused by the rate adaptation algorithm, which can influence the fairness of the results when evaluating and comparing certain performance aspects (i.e. single-radio vs. dual-radio). Moreover, prior work such as [19] showed that the rate adaptation functionality of 802.11 can influence the throughput of other hosts that share the same radio channel. That is, a host with a lower bit rate can pull down bit rates of other hosts in the vicinity, degrading their performance.

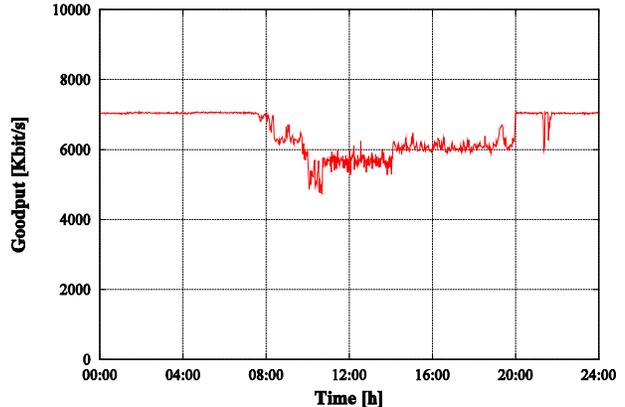


Fig.3: Effect of external interference on TCP goodput over a 24h period

Due to the increased number of IEEE 802.11 access points as well as other devices operating in the ISM 2.4 GHz band, external interference within the testbed's environment (i.e. in nearby offices) may affect running experiments. In order to eliminate such external interference, we conduct a 24-hour experiment to identify time slots with the least external interference. Figure 3 shows the TCP goodput between the rightmost and leftmost nodes of a 4-hop chain over a 24h interval. We see that during the core working time between 8am and 8pm, the measured goodput is influenced by external interference, especially due to students who access the web wirelessly through their IEEE 802.11 equipped laptops. Therefore, experiments in this paper are conducted in the time with the least external interference, between 8pm and 8am.

5.1 Experimental Cross-Validation

To validate the feasibility of ScaleMesh for emulating mesh networks, we conduct cross-validation experiments. In these experiments, we compare the results acquired from a non-scaled mesh network with the results acquired from the corresponding downscaled version of the network in ScaleMesh.

In the first experiment, we deploy two real PC mesh nodes for the non-scaled mesh topology (i.e. not mesh nodes from the testbed). Similar to the testbed nodes, both PC mesh nodes are equipped with Netgear WG311T NICs and run SuSE Linux 10.2. Unlike testbed nodes, the two deployed PC mesh nodes are not connected to any cables or attenuators. Each wireless NIC is attached to the standard 5dBi antenna which is jointly shipped with the cards. Both PC mesh nodes are placed 4m apart while the transmission power of the wireless NICs is set to 18dBm.

The next step is to emulate this one-hop topology in the testbed by using two testbed nodes with 0.5m inter-

node distance. The required attenuation level to scale the 4m distance down to 0.5m can be determined using Eq. 11, while considering a further parameter G_{diff} , which denotes the difference in antenna gain power between the 2.1dBi antennas in the testbed and the 5dBi antennas of the real mesh nodes. This parameter has to be considered in order to reflect the reduced signal amplification of the testbed antennas compared to the PC mesh antennas. By considering G_{diff} as an additional attenuation factor in Eq. 7, while solving for Ω_{sum} , we get:

$$\Omega_{sum} = -\log\left(\frac{d_{<A,B>}^{scaled}}{d_{<A,B>}^{non-scaled}}\right)10p - G_{diff} \quad (12)$$

Inserting the given values, we get $\Omega_{sum}=24.2$ dB. According to Eq. 4, for $\Omega_{sum}=24.2$ dB and $\Omega_{cab}=12.5$ dB we get 11.7dB for Ω_v . Rounded up, we set the variable attenuators to provide an attenuation level Ω_v of 12dB.

As a next step, we compare measures acquired from the non-scaled one-hop topology to the corresponding downscaled version in the testbed. As measures of interest we determine the quality of the wireless link in dBm (Figure 4), as well as the end-to-end TCP goodput in Kbit/s

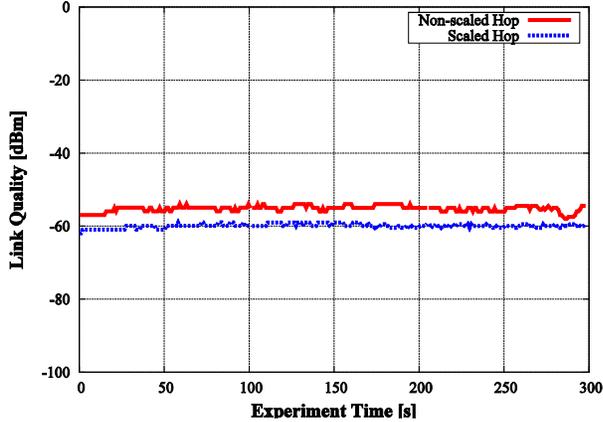


Fig. 4: Link quality vs. time for non-scaled and scaled topologies

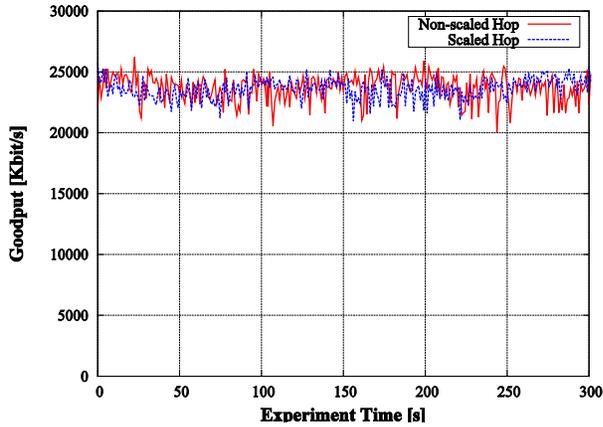


Fig 5: TCP goodput vs. time for non-scaled and scaled topologies



Fig. 6: 8-hop chain topology with a single flow

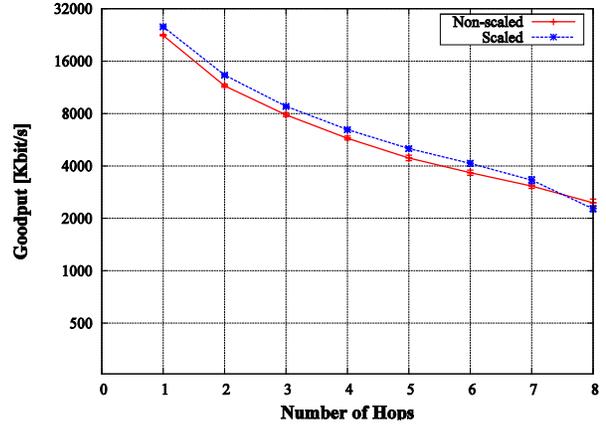


Fig. 7: TCP goodput vs. number of hops for non-scaled and scaled chain

(Figure 5), both for a duration of 5 minutes. Figure 4 shows that the dBm values for the scaled and non-scaled topologies lie relatively close to each other with a deviation of around 8%. In Figure 5, we observe that such deviation has no impact on the end-to-end TCP goodput between the nodes, since it exhibits similar values for both scaled and non-scaled topologies.

In a second validation experiment we set up a multihop, equally spaced chain of nodes as illustrated in Figure 6. Consistent with the previous experiment, we build a non-scaled version of the chain using real PC mesh nodes, as well as a corresponding scaled version of the chain within the testbed area. We utilize a single TCP flow from the leftmost node to the rightmost node of the chain and measure the achieved goodput for increasing number of hops. Figure 7 shows the results both for the non-scaled as well as for the scaled version of the chain. It's obvious that the goodput values of both scaled and non-scaled chain are nearly identical through all number of hops. This validates the capability of ScaleMesh to reproduce results acquired from non-scaled indoor networks.

In a third validation experiment, we set up a parallel chains topology as depicted in Figure 23, however, with only one hop on each chain (due to the space limitation within our lab). Consistent with the previous experiments, we build a non-scaled version of the chains using real PC mesh nodes, as well as a corresponding scaled version of the chains within the testbed area. We utilize one TCP flow on each chain and measure the aggregate goodput of both flows for the non-scaled as well as for the scaled topology. Furthermore, we consider dual-radio communication by assigning a different channel for each chain. We choose channels 1 and 11, as we show in the next section that these

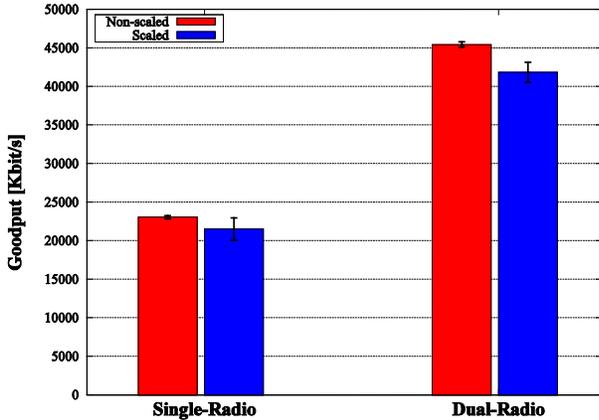


Fig. 8: Aggregate TCP goodput on two 1-hop parallel chains for non-scaled and scaled topologies

channels are indeed non-overlapping. Figure 8 shows the results of this experiment. As we see, both for single-radio as well as for dual-radio communication, the goodput of the non-scaled vs. scaled topology is almost identical with a deviation of at most 6%.

5.2 Identifying Non-Overlapping Channels

Naturally, multi-radio communication using IEEE 802.11 technology can only be beneficial if the adopted radio channels are, to a large extent, non-overlapping. According to the IEEE 802.11 specifications [24], there exists 3 non-overlapping channels out of 11 within the spectrum of 2.412 GHz to 2.462 GHz. Supposedly, channels 1, 6 and 11 are non-overlapping, which means that they do not mutually interfere. However, as mentioned in Section 3, prior work [14] showed that this assumption does not always hold in practice and can also depend on the vendor of the wireless cards.

In order to investigate this assumption, we conduct a set of experiments for identifying non-overlapping channels. Therefore, we consider a 2-hop topology comprising 3 mesh nodes in a chain, as depicted in Figure 9. The inter-node distance is 1m, whereas the attenuation level of the variable attenuators, apart from the cable attenuation, is set to 4dB. At each node, different NICs are used for the receipt and transmission of packets, respectively. Each NIC is statically assigned

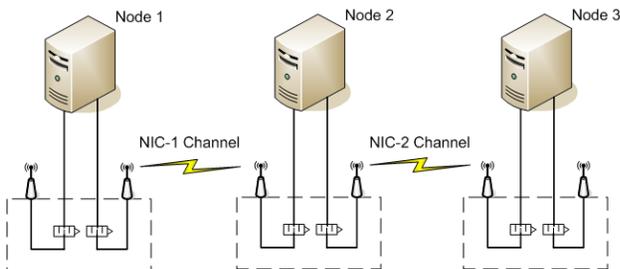


Fig. 9: 2-hop mesh topology

a channel within the available 11 channels of IEEE 802.11b/g.

Using Iperf [25], we utilize a TCP connection from node 1 as a source to node 3 as a destination. We vary the channels assigned to the NICs and plot the goodput accordingly. Naturally, the achieved goodput is correlated to the level of interference between the channels of the first and the second NIC. The less mutual interference there exists, the more goodput is achieved.

We further investigate the impact of two key parameters on the mutual interference between the wireless NICs. Particularly, we examine the impact of different link layer data rates, i.e. 2 Mbit/s, 11 Mbit/s, and 54 Mbit/s as supported by IEEE 802.11b/g. Furthermore, we study the impact of the spatial distance between the two antennas of a mesh node (i.e. inter-antenna distance), since previous work [8] has reported that such distance may have impact on the interference level between supposedly non-overlapping channels. Figures 10 to 12 show the results of the experiments with 1cm inter-antenna distance, whereas Figures 13 to 15 compare the results in Figures 10 to 12 with the results corresponding to an inter-antenna distance of 100cm.

In Figure 10, we observe that NIC channels which are separated by at least 5 channels (i.e. 1, 6, and 11) provide the highest goodput. This is consistent with the IEEE specifications [24], which specify the channel combinations indicated by the grey regions as non-overlapping.

As shown in Figure 12, varying the link layer data rate has an impact on the mutual interference between the channels. Specifically, opposed to the cases with 2 Mbit/s and 11 Mbit/s, at 54 Mbit/s, the rule of the 3 non-overlapping channels is not unambiguous. In the figure, we notice a larger variance in goodput and several performance gaps. The reason for such performance gaps may be attributed to the higher sensitivity to frequency synchronization problems as well as to possible external interference due to the long duration of these experiments.

From Figures 10 to 12 we conclude that the two non-overlapping channels which suite best for the dual NICs are 1 and 11, respectively. Therefore, throughout the rest of this paper, dual-radio experiments are conducted using these two non-overlapping channels.

Considering the spatial distance between the antennas, in Figures 13 and 14, we see that the difference in goodput between 1cm and 100cm inter-antenna distance is almost negligible at 2 Mbit/s and 11 Mbit/s. However,

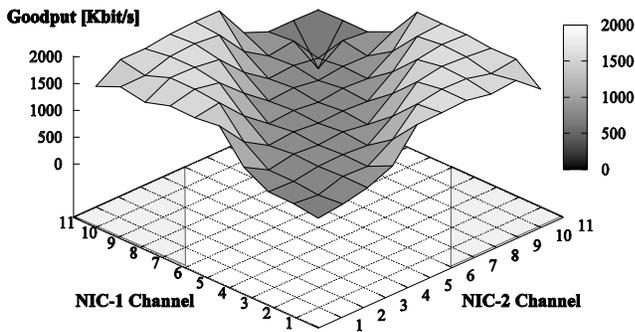


Fig. 10: TCP goodput vs. varying channel combinations at 2 Mbit/s and 1cm inter-antenna distance (Dual NIC channels which are separated by at least 5 channels provide highest goodput. Non overlapping channels according to IEEE specifications are indicated by the grey regions)

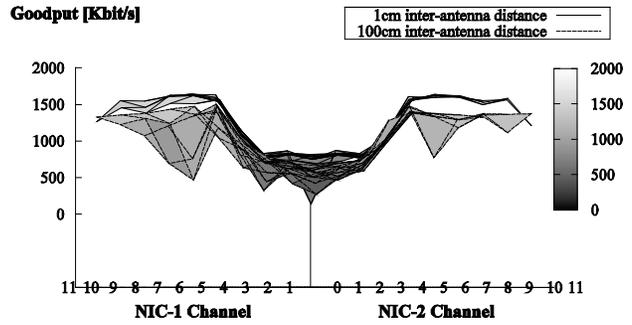


Fig. 13: TCP goodput at 2 Mbit/s for 1cm inter-antenna distance vs. 100cm inter-antenna distance (No improvement for 100cm inter-antenna distance)

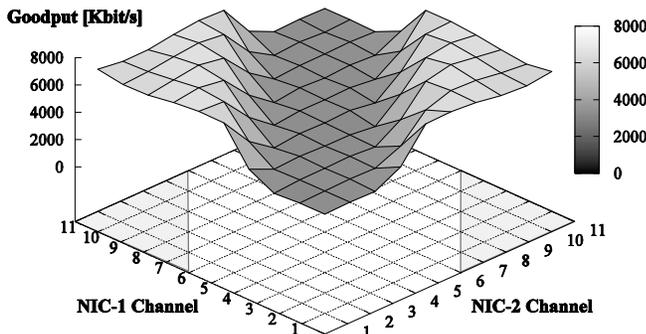


Fig. 11: TCP goodput vs. varying channel combinations at 11 Mbit/s and 1cm inter-antenna distance (Dual NIC channels which are separated by least 5 channels provide highest goodput. Non overlapping channels according to IEEE specifications are indicated by the grey regions)

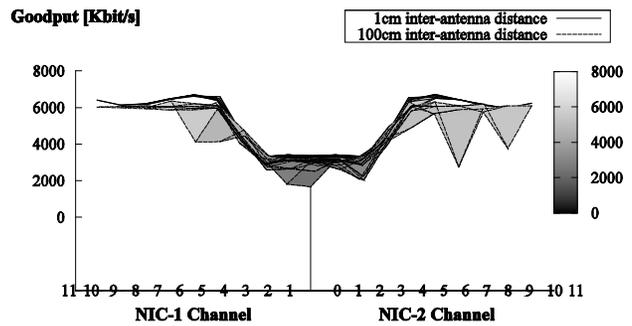


Fig. 14: TCP goodput at 11 Mbit/s for 1cm inter-antenna distance vs. 100cm inter-antenna distance (No improvement for 100cm inter-antenna distance)

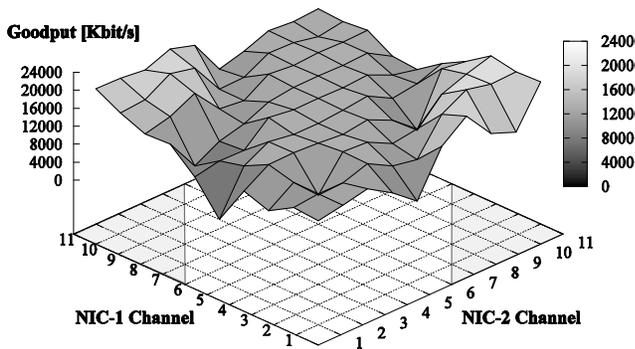


Fig. 12: TCP goodput vs. varying channel combinations at 54 Mbit/s and 1cm inter-antenna distance (Higher variance in goodput due to increased sensitivity, Non overlapping channels according to IEEE specifications are indicated by the grey regions)

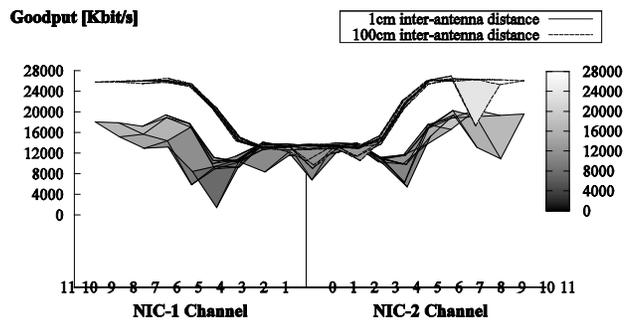


Fig. 15: TCP goodput at 54 Mbit/s for 1cm inter-antenna distance vs. 100cm inter-antenna distance (Significant improvement for 100cm inter-antenna distance)

at 54 Mbit/s, we observe a significant performance improvement of about 45% in case the antennas are separated by 100cm. This is consistent to the findings of Figures 10 to 12, which prove that at high data rates (i.e. 54 Mbit/s), the sensitivity towards external interference increases. Thus, separating the antennas of the wireless NICs decreases the mutual interference between the channels, yielding improved goodput.

5.3 Chain Topology

In this set of experiments, we consider an equally spaced chain comprising $h+1$ nodes (h hops) with a single flow, as depicted in Figure 6. Packets traverse along the chain from the leftmost node (i.e., the source) to the rightmost node (i.e., the destination). Nodes in the chain are positioned such that only direct neighbors can communicate with each other over one hop.

First, we study the impact of signal attenuation on TCP goodput. Therefore, we consider the 2-hop chain as depicted in Figure 9, where we vary the attenuation level of the attenuators. A TCP connection is utilized from node 1 to node 3 at 54 Mbit/s, and the achieved goodput is plotted for varying attenuation level. Figure 16 shows the results for single- and dual-radio communication. The setup for dual-radio communication is consistent with the description in section 3.

The figure shows that using dual radios, the goodput remains relatively constant up to a value of 4dB. As attenuation keeps increasing, the goodput decreases gradually until reaching zero at 12dB. Using a single radio, the goodput stays constant up to 7dB before decreasing sharply. This shows that dual-radio communication is more sensitive to attenuation than single-radio communication. On the other hand, dual-radio communication yields up to 70% more goodput with respect to single-radio communication.

In a further experiment, we evaluate the TCP goodput for varying chain length. In order to get intuition on the

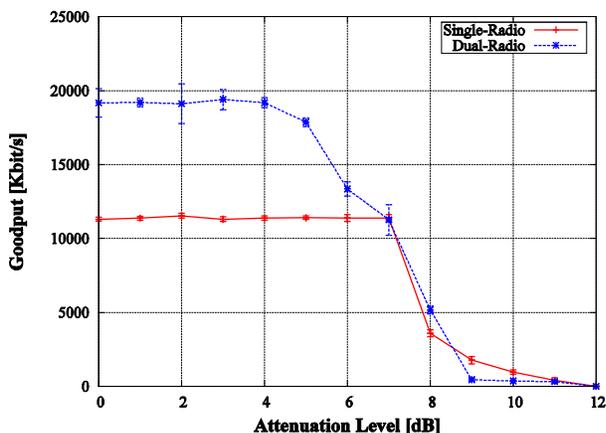


Fig. 16: TCP goodput vs. attenuation level

reliability of network simulation with respect to real-world measurements, we conduct this experiment using both testbed measurements as well as ns-2 [12] simulation. Therefore, we extended ns-2 to support dual-radio communication, since the default ns-2 implementation only supports single-radio communication. The inter-node distances in ns-2 correspond to the non-scaled distances of the considered chain topology. We also evaluate the impact of the RTS/CTS handshake on the overall performance. Figure 17 shows the results without RTS/CTS handshake, whereas in Figure 18, RTS/CTS is activated.

Figure 17 shows that, both in ScaleMesh and ns-2, dual-radio communication achieves up to 100% more goodput than single-radio communication. The ns-2 results approximate testbed results quite well, though the absolute values are naturally not identical. That is, ns-2 provides more optimistic quantitative results compared with the testbed results, mainly due to the idealized modeling assumptions in ns-2.

As observed in Figure 18, activating RTS/CTS reduces goodput significantly for all considered variants. Such severe reduction, by up to factor 4, emphasizes that the RTS/CTS mechanism does not automatically yield a performance improvement in arbitrary topologies, but rather depends on the node density and traffic patterns of a given scenario.

In a third experiment, we deploy a chain of 5 hops, vary the TCP packet size from 100 bytes up to 1500 bytes, and plot the corresponding goodput accordingly. Figure 19 shows that goodput increases with increasing packet size, since more payload is transmitted at constant link layer overhead. The absolute deviation between ns-2 results and testbed results lie in an interval of about 1%-10%. Consistent with the results of the previous experiment, Figure 20 shows that RTS/CTS causes a significant decrease in the goodput achieved by all variants.

To evaluate the achievable goodput against varying traffic load, we conduct an experiment using UDP traffic. Consistent with the previous experiment, we consider a 5-hop chain and a UDP connection between the leftmost node and the rightmost node. We vary the UDP transmission rate to increase the traffic load, and plot the UDP goodput for 2 Mbit/s, 11 Mbit/s, and 54 Mbit/s accordingly. Figures 21 and 22 show the results of this experiment.

Figure 21 shows how the goodput increases gradually with increasing UDP transmission rate up to the point where the wireless channel is fully utilized. It is obvious how a significant performance gain is achieved when choosing high data rates. In particular, around 6000 Kbit/s are achieved at 54 Mbit/s compared to 1200 Kbit/s

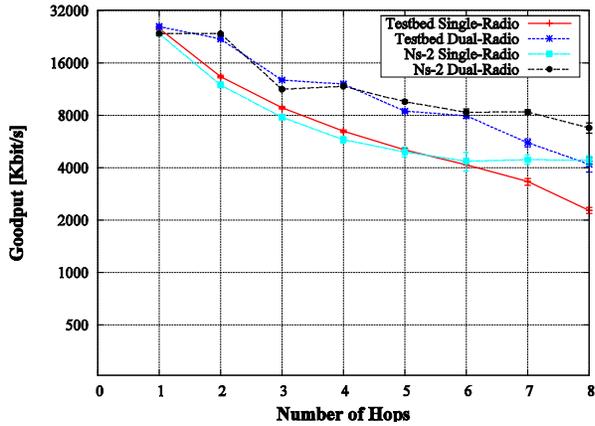


Fig. 17: TCP goodput vs. number of hops without RTS/CTS

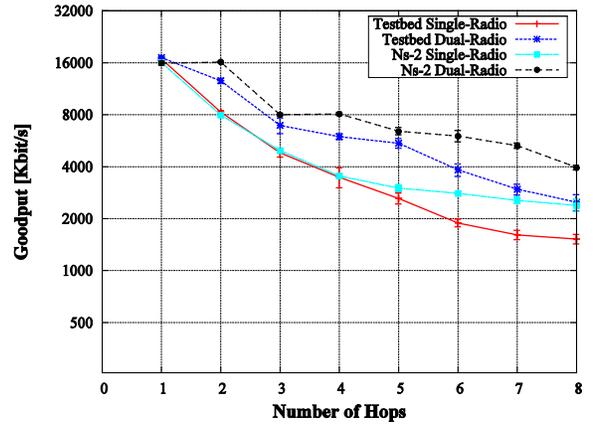


Fig. 18: TCP goodput vs. number of hops with RTS/CTS

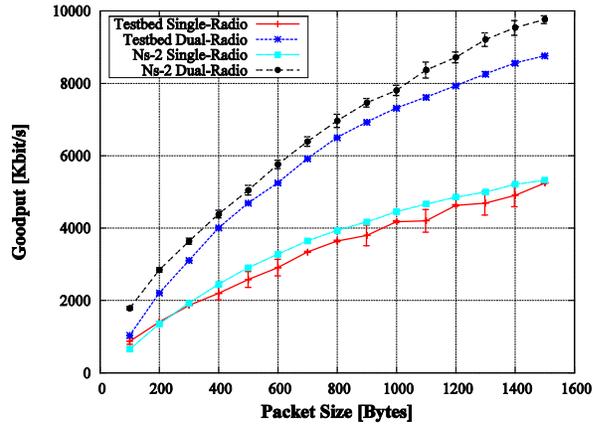


Fig. 19: TCP goodput vs. TCP packet size without RTS/CTS

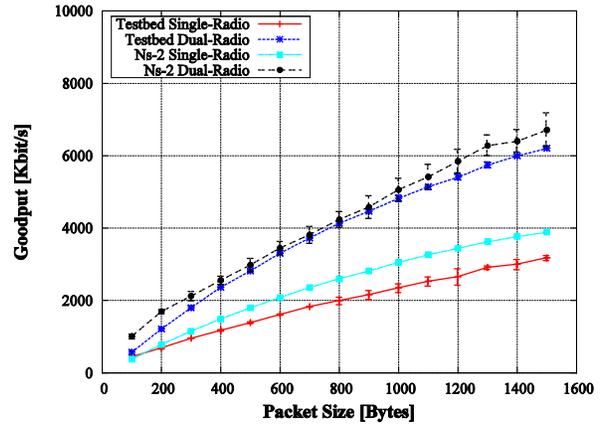


Fig. 20: TCP goodput vs. TCP packet size with RTS/CTS

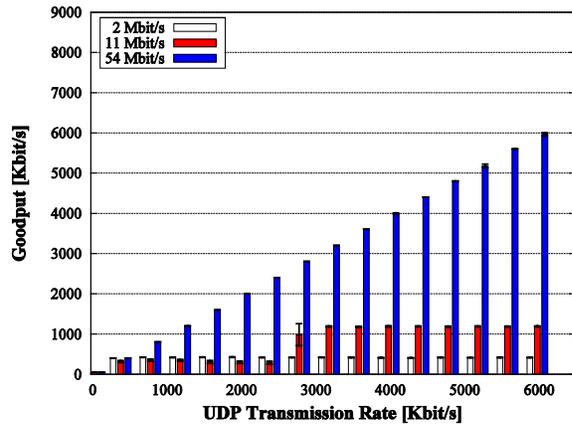


Fig. 21: UDP goodput vs. UDP transmission rate for different bandwidths and single-radio communication

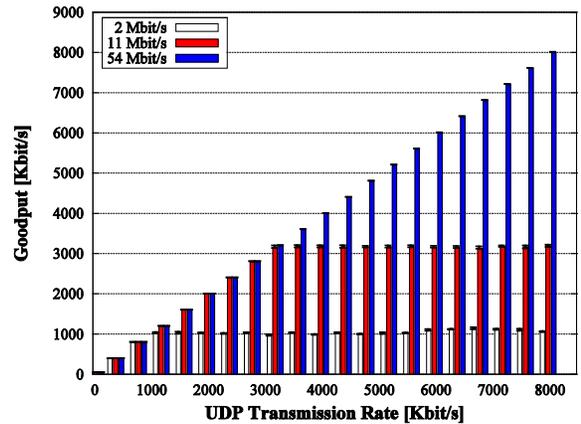


Fig. 22: UDP goodput vs. UDP transmission rate for different bandwidths and dual-radio communication

at 11 Mbit/s and 400 Kbit/s at 2 Mbit/s. As shown in Figure 22, and consistent with the previous results, using multi-radio communication yields a substantial performance improvement across all data rates.

5.4 Parallel Chains Topology

In a further set of experiments, we consider a topology of two 4-hop symmetric parallel chains, as depicted in Figure 23. The chains lie beyond each other's transmission range, but within each other's interference range. The ratio of transmission range to interference range is roughly 1 to 2.2. The transmission range was estimated using Eq. 6 (while solving for $d_{<A,B>}^{scaled}$), while the interference range was obtained using empirical measurements of the goodput as well as the signal-to-noise ratio. Specifically, similar to the approach in Figure 3, we measured the mutual influence among both chains on goodput (as well as signal-to-noise ratio) for increasing inter-chain distance. Thus, the range within which both chains mutually interfere could be determined. We experiment with both UDP and TCP flows, while evaluating the goodput achieved on each chain as well as the aggregate goodput, i.e. the sum of the goodput achieved by both chains.

In the first set of experiments, we run one UDP flow on each chain, while varying the UDP transmission rate. Figures 24 and 25 show the results of the experiments, both for single-radio as well as dual-radio communication. The setup for dual-radio communication is consistent with the description in section 3. Consistent with the previous results, we observe that dual-radio communication yields over 70% more aggregate goodput than single-radio communication.

Aside from the aggregate goodput, we notice that both for single-radio and dual-radio communication, the two UDP flows achieve similar goodput at all rates. This optimal fairness is mainly due to the absence of congestion control mechanisms in UDP. That is, when the UDP transmission rate exceeds the available bandwidth, packets are dropped at link layer, either due to buffer overflow or to unsuccessful transmissions. Since a dropped packet is followed directly by a new UDP packet, the channel on both chains stays fully utilized. In particular, opposed to TCP, UDP neither needs to wait for retransmission timeouts or acknowledgments nor maintains a window which can differ for both flows. Thus they achieve similar goodput, however, at cost of increased packet drops and low transport reliability.

To evaluate the level of fairness of TCP, we also consider a TCP flow on each chain instead of a UDP flow. Figure 26 shows the goodput achieved by each flow as well as the aggregate goodput of both flows for single-radio and dual-radio communication, respectively.

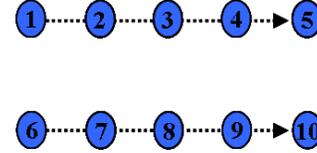


Fig. 23: Parallel chains topology

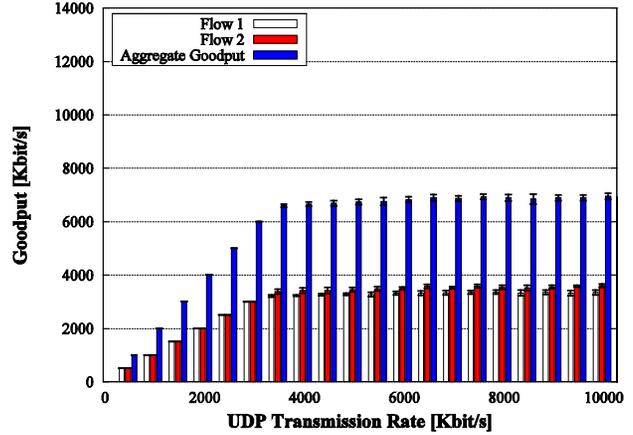


Fig. 24: Single and aggregate UDP goodput for varying UDP transmission rates and single-radio communication

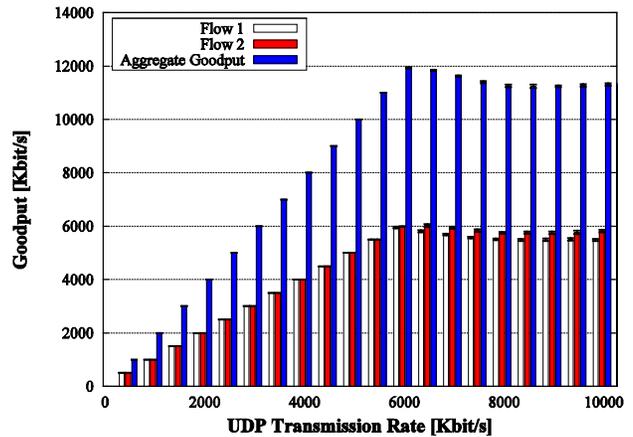


Fig. 25: Single and aggregate UDP goodput for varying UDP transmission rates and dual-radio communication

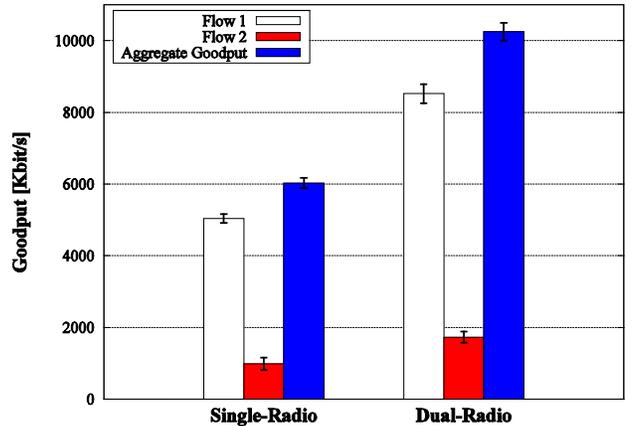


Fig. 26: Single and aggregate TCP goodput

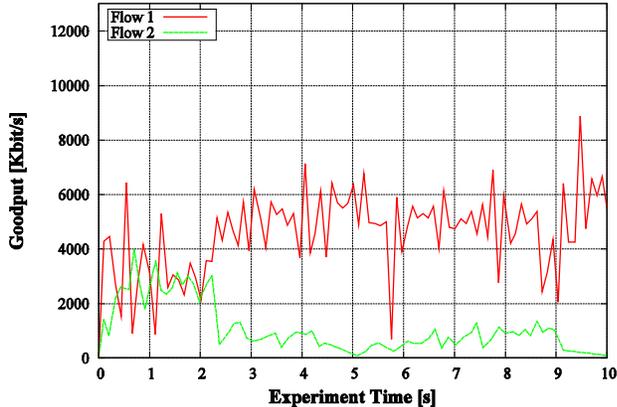


Fig. 27: TCP goodput vs. time for single-radio communication

Opposed to the case with UDP, we see that both for single-radio and dual-radio communication, flow 1 achieves significantly more goodput than flow 2. Due to the aggressive congestion control mechanism of standard TCP in combination with the well known fairness deficiencies of IEEE 802.11, the available bandwidth is not distributed fairly among both flows. These results give new insight with respect to the fairness of TCP in dual-radio mesh networks, showing that dual-radio communication does not automatically solve the fairness problem of standard TCP over IEEE 802.11.

Figures 27 and 28 show the transient goodput of TCP sampled over the initial ten seconds of the experiment, both for single-radio and dual-radio communication. Consistent with the previous findings as well as previous work on TCP in IEEE 802.11 multihop networks ([9], [10]), flow 1 acquires most of the available bandwidth over time, resulting in a severe starvation of flow 2. That is, since IEEE 802.11 favors aggressive flows over less aggressive flows, the flow which succeeds to acquire the channel first, i.e. flow 1, succeeds to take control of the channel. Unfortunately, deploying dual-radio communication does not overcome this problem.

5.5 Random Topologies

Random node topologies are found in community mesh networks such as [2] and [13], and are widely deployed in reality. To evaluate network performance in such topologies, we consider random placements of the testbed's 20 antenna-stations. Unless otherwise stated, the 20 antenna-stations are distributed uniformly on a flat area of 2m x 3m such that connectivity between each pair in the network over one or more hops is granted. Similar to Section 5.1, the signal attenuation is set to provide a transmission range of 0.5m. The setup for dual-radio communication is consistent with the description in section 3. In addition to the batch means method and in order to achieve optimal results in terms of

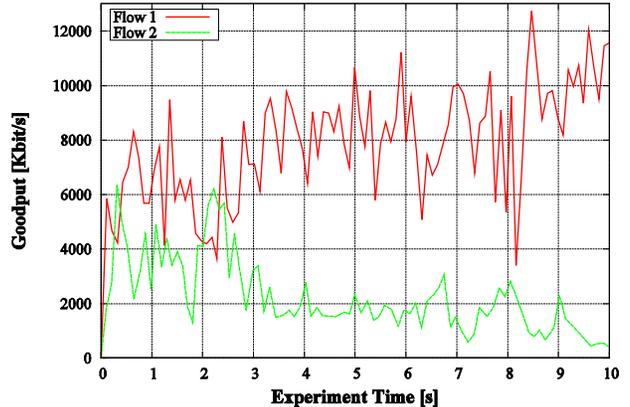


Fig. 28: TCP goodput vs. time for dual-radio communication

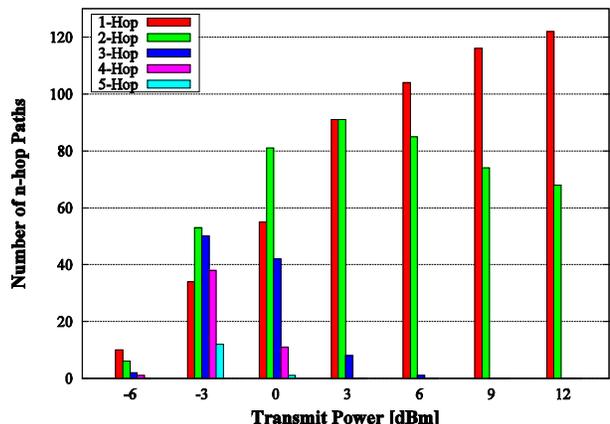


Fig. 29: Distribution of path lengths vs. transmit power of mesh nodes

representativeness, we consider 20 replicates for deriving performance measures. Each replicate corresponds to a different random placement of the nodes.

First, we investigate the correlation between the transmit power of the wireless network cards and the length of the paths between mesh nodes. Therefore, we utilize a UDP connection between each pair in the network and determine the path length in terms of number of hops for varying outgoing transmit power P_{out} . Figure 29 shows that at a transmit power of -6dBm, the transmission range of mesh nodes is at its minimum, such that only minimum connectivity to mostly 1-hop and 2-hop neighbors is granted. As transmit power increases, the number of short paths also increases, since more nodes can be reached over less number of hops. At 9dBm and above, all nodes can be reached over 2 hops or less.

In a further set of experiments, we measure the TCP goodput between each pair in the network, while varying the number of simultaneously active flows in order to vary network load. In the first experiment, only one TCP flow is active at a time, whereas in the second experiment, ten TCP flows are simultaneously active.

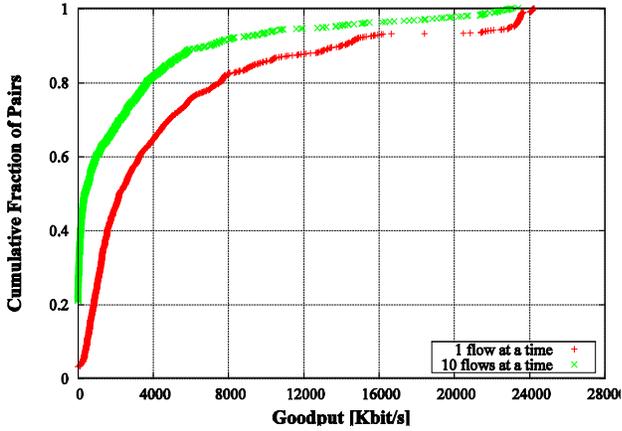


Fig. 30: Cumulative distribution function (CDF) of TCP goodput between each pair in the network for single-radio communication (Median goodput: 1 flow: **2202 Kbit/s**, 10 flows: **376 Kbit/s**)

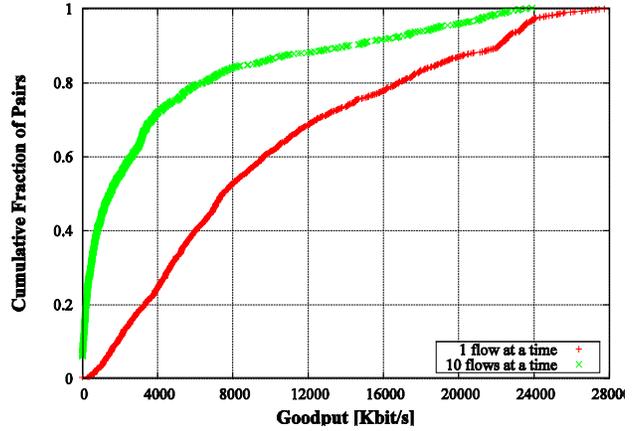


Fig. 31: Cumulative distribution function (CDF) of TCP goodput between each pair in the network for dual-radio communication (Median goodput: 1 flow: **7448 Kbit/s**, 10 flows: **1468 Kbit/s**)

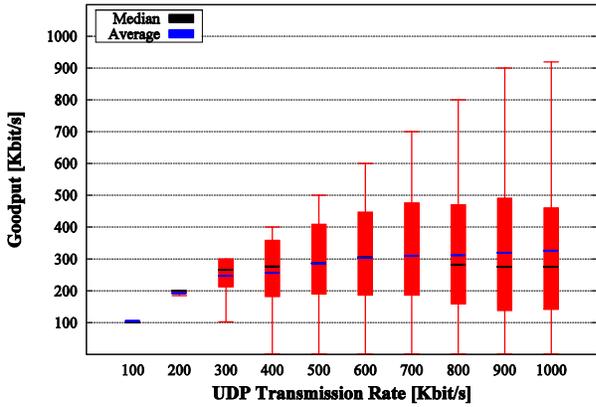


Fig. 32: Box-and-whisker diagram of UDP goodput between each pair in the network for single-radio communication and different transmission rates (10 flows at a time)

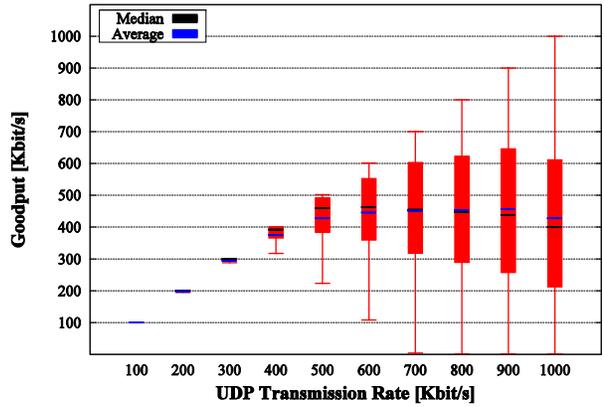


Fig. 33: Box-and-whisker diagram of UDP goodput between each pair in the network for dual-radio communication and different transmission rates (10 flows at a time)

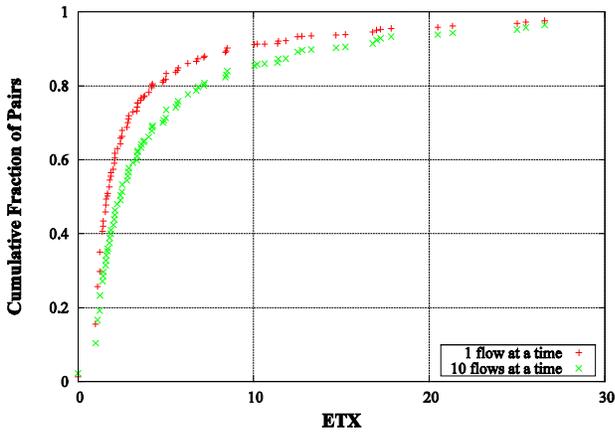


Fig. 34: Cumulative distribution function (CDF) of ETX values between each pair in the network for single-radio TCP communication (Median ETX: 1 flow: **1.678**, 10 flows: **2.419**)

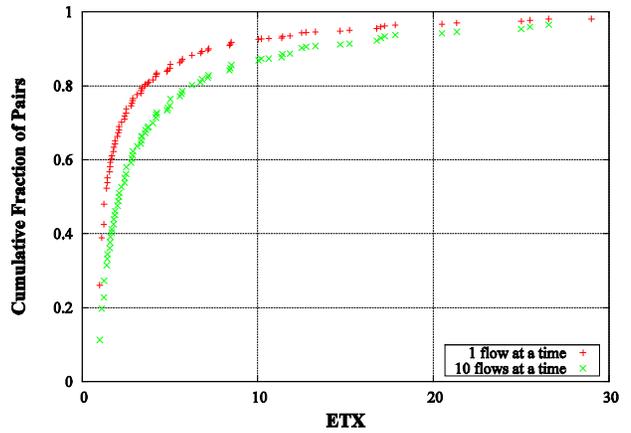


Fig. 35: Cumulative distribution function (CDF) of ETX values between each pair in the network for dual-radio TCP communication (Median ETX: 1 flow: **1.392**, 10 flows: **2.111**)

Figures 30 and 31 show the cumulative distribution function (CDF) of the TCP goodput between each pair in the network for single-radio and dual-radio communication, respectively. Figure 30 shows that the traffic load in the network has a significant impact on the goodput achieved. Specifically, the median goodput for the case of one flow at a time is 2202 Kbit/s versus 376 Kbit/s for 10 flows at a time. The CDF in Figure 31 shows that deploying dual-radio communication, mesh nodes achieve higher goodput values than the case with single-radio communication. Consistent with the findings in Figure 30, active flows achieve significantly more goodput at moderate traffic load than at high traffic load. Specifically, the median goodput for the case of one flow at a time is 7448 Kbit/s versus 1468 Kbit/s for 10 flows at a time.

To get further insight on the effect of varying traffic load on network performance, we re-conduct the previous experiment using UDP instead of TCP traffic. Thereby, we successively vary the UDP transmission rate between each pair in the network and plot the corresponding goodput as box-and-whisker diagram in Figures 32 and 33 for single-radio communication and dual-radio communication, respectively. Recall that in the whisker diagrams, filled boxes visualize the distribution of 75% of all goodput values, whereas the black and blue lines within the boxes represent the median and the mean, respectively. The vertical lines indicate the smallest and largest observations that are less than 1.5 times the interquartile range (IQR).

In Figure 32, we observe that for a UDP transmission rate below 300 Kbit/s, the load is moderate such that optimum goodput is achieved. As the UDP transmission rate increases, the variance also increases, and the median even decreases for transmission rates higher than 700 Kbit/s, indicating an increasing load in the network. That is, at a transmission rate of 700 Kbit/s, mesh nodes achieve more goodput than at higher transmission rates. Opposed to the case with single-radio communication, for dual-radio communication, the variance at 400 Kbit/s and below is almost negligible. The load is so moderate such that almost all UDP connections achieve optimum goodput. Starting from 500 Kbit/s, the variance grows. While the median at a transmission rate of 500 Kbit/s is nearly-optimal with around 460 Kbit/s, it decreases down to 400 Kbit/s at 1000 Kbit/s, indicating an increasing load in the network.

As mentioned in Section 3, we employ OLSR as routing protocol with ETX [6] support. To get intuition on the distribution of the ETX values of the corresponding links, we plot the CDF of ETX values for all links in the network. Figures 34 and 35 show the results. Note that, as given in Eq. 1, a low ETX value

indicates a better link quality than a high ETX value.

Consistent with the previous results, the figures show that the ETX distribution exhibits lower values at lower traffic load. Figure 35 implies that dual-radio communication does substantially improve link quality, however, not to the extent of the goodput improvements shown in Figures 30 and 31. In particular, according to Eq. 1, the probability for a successful packet delivery equals $1/ETX$. For 1 flow, this is 0.6 for the ETX median 1.678 (single-radio) and 0.72 for the ETX median 1.392 (dual-radio). This is an improvement of 0.12. For 10 flows, we get a probability of 0.41 for the ETX median 2.419 (single-radio) and 0.47 for the ETX median 2.111 (dual-radio). This is an improvement of 0.07. This shows that the improvement in ETX values does not proportionally correlate with the improvement in goodput.

6 Conclusion

We introduced ScaleMesh, a scalable miniaturized dual-radio wireless mesh testbed based on IEEE 802.11b/g technology. ScaleMesh emulates large-scale mesh networks on a miniaturized experimentation area by deploying variable signal attenuators that shrink the transmission range of wireless nodes. We formally derived the correlation between the inter-node distance in the network to be emulated, the inter-node distance in the downscaled version within the testbed experimentation area, and the attenuation level of the transmission signal. Such a correlation allows approximating the attenuation level required for downscaling desired network topologies.

In a comprehensive performance study, we evaluated various aspects of ScaleMesh. After identifying non-overlapping channels in the 2.4 GHz spectrum, we conducted various sets of experiments over different network topologies and traffic patterns using TCP and UDP. The results showed that dual-radio communication yields a significant improvement in goodput, up to 100% more than single-radio communication. Yet, we showed that TCP fairness deficiencies over IEEE 802.11 are not overcome using dual radios.

Ns-2 [12] simulations, which we conducted along with testbed experiments, showed that qualitatively, the ns-2 output approximates ScaleMesh's output within a deviation of roughly 10%. Although absolute values do not match due to the high level of abstraction in ns-2, simulations can indeed help gain initial results on the performance of newly designed protocols.

7 References

- [1] I. Akyildiz, X. Wang, and W. Wang, Wireless Mesh Networks: A Survey, *Elsevier Computer Networks*, 47, 2005.
- [2] J. Bicket, D. Aguayo, S. Biswas, and R. Morris, Architecture and Evaluation of an Unplanned 802.11b Mesh Network, *Proc. ACM MOBICOM*, Cologne, Germany, 2005.
- [3] J. Camp, J. Robinson, C. Steger, and E. Knightly, Measurement Driven Deployment of a Two-Tier Urban Mesh Access Network, *Proc. ACM MobiSys*, Uppsala, Sweden, 2006.
- [4] C. Chereddi, P. Kyasanur, and N. Vaidya, Design and Implementation of a Multi-Channel Multi-Interface Network, *Proc. ACM REALMAN*, Florence, Italy, 2006.
- [5] T. Clausen and P. Jacquet, Optimized Link State Routing Protocol, *RFC 3626*, <http://www.ietf.org/rfc/rfc3626.txt>, October 2003.
- [6] D. De Couto, D. Aguayo, J. Bicket, and R. Morris, A High-Throughput Path Metric for Multi-Hop Wireless Routing, *Proc. ACM MOBICOM*, San Diego, CA, 2003.
- [7] P. De, A. Raniwala, R. Krishnan, K. Tatavarthi, J. Modi, N. Syed, S. Sharma, and T. Chiueh, MiNT-m: An Autonomous Mobile Wireless Experimentation Platform, *Proc. ACM MobiSys*, Uppsala, Sweden, 2006.
- [8] R. Draves, J. Padhye, and B. Zill, Routing in Multi-Radio, Multi-Hop Wireless Mesh Networks, *Proc. ACM MOBICOM*, Philadelphia, PA, 2004.
- [9] S. ElRakabawy, A. Klemm, and C. Lindemann, TCP with Adaptive Pacing for Multihop Wireless Networks, *Proc. ACM MobiHoc*, Urbana-Champaign, IL, 2005.
- [10] S. ElRakabawy, A. Klemm, and C. Lindemann, TCP with Gateway Adaptive Pacing for Multihop Wireless Networks with Internet Connectivity, *Computer Networks*, 52, 2008.
- [11] J. Eriksson, S. Agarwal, P. Bahl, and J. Padhye, Feasibility Study of Mesh Networks for All-wireless Offices, *Proc. ACM MobiSys*, 2006, Uppsala, Sweden.
- [12] K. Fall and K. Varadhan (Ed.), The ns-2 Manual, Technical Report, The VINT Project, UC Berkeley, LBL, and Xerox PARC, 2007.
- [13] Freifunk Mesh Community, <http://start.freifunk.net/>.
- [14] P. Fuxjager, D. Valerio, and F. Ricciato, The Myth of Non-Overlapping Channels: Interference Measurements in IEEE 802.11, *Proc. IEEE/IFIP WONS*, Obergurgl, Austria, 2007.
- [15] V. Gambiroza, B. Sadeghi, and E. Knightly, End-to-End Performance and Fairness in Multihop Wireless Backhaul Networks, *Proc. ACM MOBICOM*, Philadelphia, PA, 2004.
- [16] T. Gleixner and D. Niehaus, Hrtimers and Beyond: Transforming the Linux Time Subsystems, *Proc. 8th OLS Linux Symposium*, Ottawa, Canada, 2006 (Source code available at <http://www.tglx.de/projects/ktimers/>).
- [17] M. Lihan, T. Tsuchiya, and K. Koyanagi, Orientation-Aware Indoor Localization Path Loss Prediction Model for Wireless Sensor Networks, *Lecture Notes in Computer Science*, Vol. 5186, 2008.
- [18] H. Lundgren, K. Ramachandran, E. Belding-Royer, K. Almeroth, M. Benny, A. Hewatt, A. Touma, and A. Jardosh, Experiences from the Design, Deployment, and Usage of the UCSB MeshNet testbed, *IEEE Wireless Communications*, 13(2), 2006.
- [19] A. Munaretto, M. Fonseca, K. Al Agha, and G. Pujolle, Fair Time Sharing protocol: A solution for IEEE 802.11b hot spots, *Lecture Notes in Computer Science*, Vol. 3124, 2004.
- [20] D. Raychaudhuri, I. Seskar, M. Ott, S. Ganu, K. Ramachandran, H. Kremo, R. Siracusa, H. Liu, and M. Singh, Overview of the ORBIT Radio Grid Testbed for Evaluation of Next-Generation Wireless Network Protocols, *Proc. IEEE WCNC*, New Orleans, LA, 2005.
- [21] A. Raniwala, T. Chiueh, Architecture and algorithms for an IEEE 802.11-based multi-channel wireless mesh network, *Proc. IEEE INFOCOM*, Miami, FL, 2005.
- [22] S. Saunders, Antennas and Propagation for Wireless Communication Systems, *Wiley & Sons*, May 2007.
- [23] N. Vaidya, J. Bernhard, V. Veeravalli, P. Kumar, and R. Iyer, Illinois Wireless Wind Tunnel: A Testbed for Experimental Evaluation of Wireless Networks, *Proc. ACM E-WIND Workshop*, Philadelphia, PA, 2005.
- [24] IEEE Standard for Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications, ISO/IEC 8802-11, August 1999.
- [25] Iperf, the TCP/UDP Bandwidth Measurement Tool, <http://dast.nlanr.net/Projects/Iperf/>.
- [26] OLSR.ORG Implementation for Linux, <http://www.olsr.org>.
- [27] The Qualnet Simulator, <http://www.scalable-networks.com/>.