

# ScaleMesh: A Scalable Dual-Radio Wireless Mesh Testbed

Sherif M. ElRakabawy, Simon Frohn, and Christoph Lindemann  
University of Leipzig  
Department of Computer Science  
Johannisgasse 26  
04103 Leipzig, Germany

**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 present a performance study in which we validate the feasibility of ScaleMesh for network emulation and protocol evaluation. We further conduct single-radio vs. dual-radio experiments in ScaleMesh, and show that dual-radio communication significantly improves network goodput. The median TCP goodput we observe in a typical random topology at 54 Mbit/s and dual-radio communication ranges between 1468 Kbit/s and 7448 Kbit/s, depending on the current network load.

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

## I. 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 [8] and Qualnet [15], the trend is increasingly moving towards deploying such networks in reality. Examples include MIT Roofnet [2], TFA-Rice ([3], [10]), and Freifunk [9], 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. Opposed to previous proposals such as [6], and [17], ScaleMesh comprises variable attenuators to variably adjust the transmission range and thus flexibly emulate large-scale networks. Such dynamic network scaling combined with dual-radio support allows emulating and evaluating a large variety of wireless mesh networks.

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. In a performance study, we validate the feasibility of ScaleMesh, and conduct single-radio vs. dual-radio experiments to give insight on the achievable performance in typical random network topologies.

The remainder of this paper is organized as follows. Section II summarizes related work on real deployments of wireless mesh networks as well as testbed prototypes. Section III describes the architecture and operation of the introduced miniaturized wireless mesh testbed, whereas in Section IV 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 performance study of the testbed is presented in Section V. Finally, concluding remarks are given.

## II. 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 omnidirectional antennas. Gambiroza et al. [10] simulated a multihop wireless backhaul network consisting of multiple Transit Access Points (TAPs), which are connected to the Internet through multiple entry points. Based on the findings in [10], 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.

Opposed to [2], [3], and [10], we introduce a scalable miniaturized mesh testbed rather than a large-scale mesh network. Using our testbed, networks similar to [2], [3], and

[10] can be emulated within a miniaturized experimentation area. Beyond [2], [3], and [10], our testbed further supports multiple radios rather than only a single radio.

De et al. [6] 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 [7], 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. Raychaudhuri et al. [17] 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 [16] on their experience in designing and deploying the UCSB MeshNet, a 30-node wireless mesh testbed which covers several floors inside a building. In [19], 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 [6], [7], [16], and [17], our testbed supports multiple radios. Opposed to [7], [17], and [19], 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 [6] significantly limits the spectrum of network topologies which can be considered due to the fixed transmission range associated with the attenuators.

### III. 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 in a 10m x 6m wireless mesh lab 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 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 to a few meters, 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 [11]. As driver for the wireless PCI cards, we employ the Linux 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 ([4], [14]) for multihop routing. This implementation of OLSR incorporates the Expected Transmission Count (ETX) metric [5] for selecting routes based on the current loss probability of the links. 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 I shows a detailed description of hardware and software components of the miniaturized testbed.

ScaleMesh supports dual-radio communication between mesh nodes by assigning a different channel to each of the two



Fig.1: ScaleMesh with a sample topology

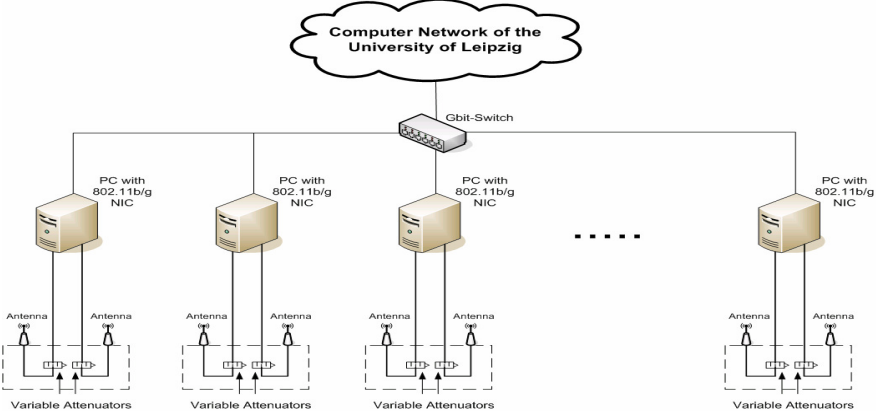


Fig. 2: Architecture of ScaleMesh

TABLE I: 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 2.1 dBi 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

wireless PCI cards of a node. Besides the different channels, each wireless card is assigned a different IP address. The routing protocol is responsible for routing the packets such that both channels are fully utilized in order to minimize inter-link interference. The IEEE 802.11b/g standard supports 11 different channels. According to the IEEE 802.11 specifications [12], channels 1, 6, and 11 are non-overlapping. Thus, in dual-radio experiments we assign channels 1 and 11 to the wireless NICs, respectively.

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.

#### IV. 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. In spite of the significant extra cost, 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 can only provide a fixed level of signal attenuation, limiting the spectrum of topologies that can be deployed.

Within this context, it's crucial to acquire certain topology-dependent 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 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  $\Omega_{sum}$ . 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  $\Omega_{sum}$  is to approximate the signal attenuation between nodes  $A$  and  $B$ . Following the Equivalent Isotropically Radiated Power (EIRP) [18] equation we get:

$$P_{out} = P_{tx} + G_{ant} \quad (1)$$

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

The signal attenuation 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:

$$L_{\langle A,B \rangle} = P_{out} - P_{rx} - \Omega_{sum} \quad (2)$$

where

$$\Omega_{sum} = \Omega_{cab} + \Omega_v \quad (3)$$

Here,  $\Omega_{cab}$  and  $\Omega_v$  describe the signal attenuation caused by the coaxial cable and the variable attenuator, respectively. In ScaleMesh,  $\Omega_{cab}$  is roughly 12.5 dB and  $\Omega_v$  is naturally variable.

The next step towards identifying the correlation between  $d^{non-scaled}$ ,  $d^{scaled}$ , and  $\Omega_{sum}$  is to derive the signal attenuation between nodes  $A$  and  $B$  as function of  $d_{\langle A,B \rangle}^{scaled}$ , 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 [18] is given by:

$$L = 20 \log_{10}(f_c) + 10p \log_{10}(d) \quad (4)$$

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. Due to findings from extensive measurements in ScaleMesh and following [18], we set  $p = 3$ .

Considering  $d_{<A,B>}^{scaled}$  as the distance between nodes  $A$  and  $B$ , we get according to Eq. 4:

$$L_{<A,B>} = 20 \log_{10}(f_c) + 10p \log_{10}(d_{<A,B>}^{scaled}) \quad (5)$$

Solving for  $d_{<A,B>}^{scaled}$  while considering Eq. 2 we get:

$$\begin{aligned} d_{<A,B>}^{scaled} &= 10^{\frac{L_{<A,B>} - 20 \log_{10}(f_c)}{10p}} \\ &= 10^{\frac{P_{out} - P_{rx} - \Omega_{sum} - 20 \log_{10}(f_c)}{10p}} \\ &= 10^{\frac{P_{out} - P_{rx} - 20 \log_{10}(f_c)}{10p}} \cdot 10^{-\frac{\Omega_{sum}}{10p}} \\ &= d_{<A,B>}^{non-scaled} \cdot 10^{-\frac{\Omega_{sum}}{10p}} \end{aligned} \quad (6)$$

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  $\Omega_{sum}$ . This delivers 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. Note that in case 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 V.A.

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

$$d_{<A,B>}^{non-scaled} = \frac{d_{<A,B>}^{scaled}}{10^{-\frac{\Omega_{sum}}{10p}}} \quad (7)$$

and

$$\Omega_{sum} = -\log\left(\frac{d_{<A,B>}^{scaled}}{d_{<A,B>}^{non-scaled}}\right) 10p \quad (8)$$

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.

## V. PERFORMANCE STUDY

We conduct an experimental performance study using ScaleMesh, in which we evaluate the influence of different key parameters on network performance. In all experiments, except for experiments 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. Dependent on the respective experiment, we generate TCP/UDP traffic using the Iperf bandwidth measurement tool for Linux [13]. Unless otherwise stated, RTS/CTS is disabled and the link-layer data rate is set to 54 Mbit/s.

Due to the increased number of IEEE 802.11 access points as well as other operating devices 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 conducted a 24-hour experiment to identify time slots with the least external interference. We found out 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.

### A. Experimental Cross-Validation

To validate the feasibility of ScaleMesh for emulating mesh networks, we conduct a cross-validation experiment. In this experiment, 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. For the non-scaled mesh topology, we deploy two real PC mesh nodes (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. 8, while considering a further parameter  $G_{diff}$ , which denotes the difference in antenna gain power between the 2.1 dBi 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. 6, while solving for  $\Omega_{sum}$ , we get:

$$\Omega_{sum} = -\log\left(\frac{d_{<A,B>}^{scaled}}{d_{<A,B>}^{non-scaled}}\right) 10p - G_{diff} \quad (9)$$

Inserting the given values, we get  $\Omega_{sum}=24.2$ dB. According to Eq. 3, for  $\Omega_{sum}=24.2$ dB and  $\Omega_{cab}=12.5$ dB we get 11.7dB for  $\Omega_v$ . Rounded up, we set the variable attenuators to provide an attenuation level  $\Omega_v$  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 3), as well as the end-to-end TCP goodput in Kbit/s (Figure 4), both for a duration of 5 minutes.




Fig. 3: Link quality vs. time for non-scaled and scaled topologies




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

Figure 3 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 4, 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. This validates the capability of ScaleMesh to reproduce results acquired from large-scale networks.

### B. Random Topologies

Random node topologies are typically found in community mesh networks such as [2] and [9], and are widely deployed in reality. To evaluate network performance in such topologies, we consider random placements of the testbed's 20 nodes. Unless otherwise stated, the 20 antenna-stations are distributed uniformly on a flat area of 2m x 3m such that full connectivity between each pair in the network over one or more hops is granted. In addition to the batch means method described earlier, and in order to achieve optimal results in terms of representativeness, we consider 20 replicates when 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 path length 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 5 shows that at a transmit power of -6 dBm, the transmission range of mesh nodes is at its minimum, such that only minimum connectivity to mostly one-hop and two-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 9 dBm 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. Figures 6 and 7 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 6 shows that the traffic load in the network has a significant impact on the goodput achieved. Specifically, while around 20% of all nodes achieve 4000 Kbit/s or more in the case of ten simultaneously active flows, around 40% of the nodes achieve such values in the case of one active flow at a time. 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 distribution in Figure 7 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 6, 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 vary the UDP transmission rate between each pair in the network and plot the corresponding goodput as box-and-whisker diagram in Figures 8 and 9 for single-radio communication and dual-radio communication, respectively. Due to space limitations, we omit the figures for the experiment with one active flow at a time. Recall that in the whisker diagrams, filled boxes visualize the distribution of 75% of all goodput values, whereas the black lines within the boxes represent the median. The vertical lines indicate the smallest and largest observations that are less than 1.5 times the interquartile range (IQR).

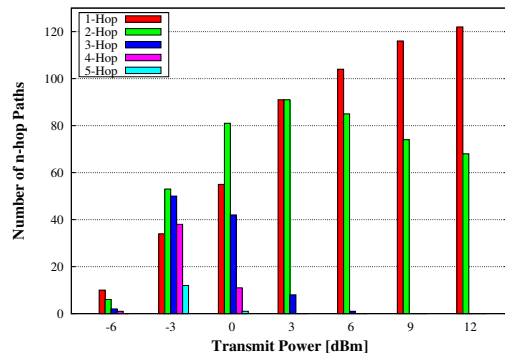


Fig. 5: Distribution of path lengths vs. outgoing transmission power of mesh nodes




Fig. 6: 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)

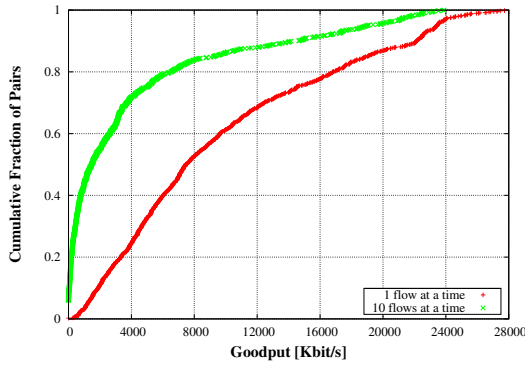
## VI. CONCLUSION AND FUTURE WORK

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 correlation allows approximating the attenuation level required for downscaling desired network topologies.

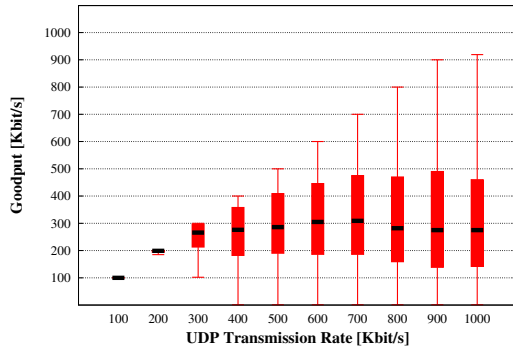
In a performance study, we validated the performance of ScaleMesh against a corresponding non-scaled topology. Further experiments with single and dual radios using TCP and UDP traffic showed that dual-radio communication yields a significant improvement in goodput compared to single-radio communication. In future work we will be further extending the number of nodes in ScaleMesh and emulating a subnet of the community mesh network Freifunk [9] to evaluate newly designed network protocols.

## REFERENCES

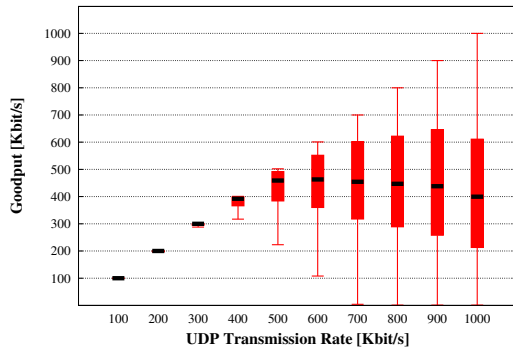
- [1] I. Akyildiz, X. Wang, and W. Wang, Wireless mesh networks: a survey, *Elsevier Computer Networks*, 47, 2005.
- [2] J. Bicket, D. Aguayo, S. Biswas, and R. Morris, Architecture and Evaluation of an Unplanned 802.11b Mesh Network, *Proc. ACM MOBICOM*, Cologne, Germany, 2005.
- [3] J. Camp, J. Robinson, C. Steger, and E. Knightly, Measurement Driven Deployment of a Two-Tier Urban Mesh Access Network, *Proc. ACM MobiSys*, Uppsala, Sweden, 2006.
- [4] T. Clausen and P. Jacquet, Optimized Link State Routing Protocol, *RFC 3626*, <http://www.ietf.org/rfc/rfc3626.txt>, October 2003.
- [5] 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.
- [6] P. De, A. Raniwala, R. Krishnan, K. Tatarvarthi, J. Modi, N. Syed, S. Sharma, and T. Chiueh, MiNT-m: An Autonomous Mobile Wireless Experimentation Platform, *Proc. ACM MobiSys*, Uppsala, Sweden, 2006.
- [7] J. Eriksson, S. Agarwal, P. Bahl, and J. Padhye, Feasibility Study of Mesh Networks for All-wireless Offices, *Proc. ACM MobiSys*, 2006, Uppsala, Sweden.
- [8] K. Fall and K. Varadhan (Ed.), The ns-2 Manual, Technical Report, The VINT Project, UC Berkeley, LBL, and Xerox PARC, 2007.
- [9] Freifunk Mesh Community, <http://start.freifunk.net/>.
- [10] V. Gambauro, B. Sadeghi, and E. Knightly, End-to-End Performance and Fairness in Multihop Wireless Backhaul Networks, *Proc. ACM MOBICOM*, Philadelphia, PA, 2004.
- [11] 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/>).
- [12] IEEE Standard for Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications, ISO/IEC 8802-11, August 1999.
- [13] Iperf, the TCP/UDP Bandwidth Measurement Tool, <http://dast.nlanr.net/Projects/Iperf/>.
- [14] OLSR.ORG Implementation for Linux, <http://www.olsr.org>.
- [15] The Qualnet Simulator, <http://www.scalable-networks.com/>.
- [16] 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.
- [17] D. Raychaudhuri, I. Sesar, M. Ott, S. Ganu, K. Ramachandran, H. Kremono, 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.
- [18] S. Saunders, Antennas and Propagation for Wireless Communication Systems, *Wiley & Sons*, May 2007.
- [19] 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.



**Fig. 7:** 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. 8:** 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. 9:** 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)

In Figure 8, 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 increases. 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.