

Validation of a Miniaturized Wireless Network Testbed ^{*}

Yang Su
Computer Science Department
ETH Zurich
8092 Zurich Switzerland
ysu@inf.ethz.ch

Thomas Gross
Computer Science Department
ETH Zurich
8092 Zurich Switzerland
trg@inf.ethz.ch

ABSTRACT

Experimentation with wireless network testbeds is preferred to simulations for performing high fidelity testing. However, realistic experimentation with full-scale networks is difficult, time-consuming, and expensive. To side-step some of these problems, we have built IvyNet, a miniaturized wireless mesh network testbed, where a high degree of attenuation (up to 60dB) is applied to the wireless radios for multi-hop network testing. In this paper, we present a validation study of the miniaturized wireless networks via experiments in the IvyNet testbed. Our results show that IvyNet is able to facilitate convenient experimentation over multi-hop networks within a small space; at the same time, it preserves reasonable fidelity compared to real-world full-scale networks. Because IvyNet is based upon commodity hardware, the principal observations from the IvyNet experiments also apply to other miniaturized wireless networks.

Categories and Subject Descriptors

C2.1 [Network Architecture and Design]: Wireless communication

General Terms

Experimentation, Measurement, Performance

Keywords

Wireless Networks, IvyNet, Experiments, Miniaturized Testbed

1. INTRODUCTION

Wireless multi-hop networks have become increasingly popular over the past several years. Most of the time, re-

^{*}The work presented in this paper was partially supported by the National Competence Center in Research on Mobile Information and Communication Systems (NCCR-MICS), a center supported by the Swiss National Science Foundation under grant number 5005-67322

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

WinTech '08, September 19, 2008, San Francisco, California, USA.
Copyright 2008 ACM 978-1-60558-187-3/08/09 ...\$5.00.

searchers rely on wireless simulators for evaluating new protocols/applications for wireless networks. However, simulators are often based upon idealized and simplified radio propagation and physical environment models that fail to capture the complex channel characteristics of a real system faithfully. Additionally, most of the simulators do not take into account the software/hardware constraints from real systems. The fidelity of the simulation results remain a concern [10], [15]. There is an ever increasing need in the research community for research testbeds that allow users to conveniently perform controllable experiments using real-world wireless devices. Researchers have built several full-scale wireless network testbeds [5, 2, 9]. Because the long transmission range of the wireless radios, a full-scale multi-hop testbed normally requires a large space to operate, and the deployment and management of such testbeds are expensive.

To reduce the complexity of developing full-scale wireless testbeds, researchers have built miniaturized testbeds [8, 17, 7, 1] to accommodate hundreds of nodes into a smaller space while maintaining link characteristics through power control. They reduce the transmission power via software control and/or RF attenuator hardware, and reception power via augmented environmental noise and/or attenuator hardware. Although the miniaturized techniques are effective in reducing the communication distance among wireless devices, their fidelity in a large-scale network is still not fully understood.

We have built IvyNet [1], a scaled-down wireless mesh network testbed, where heavy attenuations (up to 60dB) are applied to the wireless radios for convenient multi-hop network testing. In this paper, we present a validation study of the IvyNet testbed. Specifically, we try to answer the questions that we often get regarding the fidelity of the miniaturized testbed:

1. Do network applications and protocols achieve comparable performance over links in the miniaturized network as in a real-world full-scale network?
2. In the miniaturized network, do experiments suffer more/less from interference than they do in a full-scale network?
3. For a network setup in the miniaturized testbed, is it possible to construct a corresponding setup in a real-world full-scale network and preserve the same radio-level characteristics?

Our results show that the heavily attenuated radios enable convenient multi-hop experiments within a limited space.



Figure 1: IvyNet Testbed.

At the same time, the attenuated wireless links preserve a spectrum of signal qualities that is large enough to reproduce most of the conditions that network applications would experience in the real-world networks. In addition, the wireless devices in the testbed show strong capture effects, and they strictly follow the SINR model when resolving the interference. Consequently experiments in the miniaturized network experience similar interference effects as in the full-scale network although the actual interference level might be different. In the end, for a single link in the miniaturized testbed, it is possible to map it into a network at another spatial scale. Whereas the miniaturized testbed operates in a relatively clean environment, and it has a very small path loss factor that cannot be found in typical large-scale deployments. As a result, there is a non-linear mapping between the networks in the miniaturized testbed and the networks in full-scale deployment. For networks with complex topologies, it would be difficult to find a direct mapping that preserves both the radio characteristics and the spatial relations among wireless devices. Hence the miniaturized testbed should be used with caution to evaluate protocols for which absolute spatial structure matters (e.g. routing protocols).

The rest of the paper is organized as follows. Section 2 discusses related work. Section 3 describes the overall architecture of the IvyNet testbed. Section 4 presents the validation of the miniaturized testbed. Section 5 discusses the network scaling problem. Section 6 concludes the paper.

2. RELATED WORK

Many full-scale wireless mesh network testbeds [5, 2, 9] have been built in recent years. The roofnet testbed [5] and the Purdue wireless mesh network testbed [2] both consist of large number of 802.11 wireless devices spread over a large campus area. Magnet [9] is a high-performance multi-channel wireless mesh network testbed deployed in the city of Berlin. On the one hand, all the full-scale testbeds provide a realistic platform for studying the behavior of wireless mesh networks. On the other hand, it is expensive to deploy and manage large number of geographically distributed wireless nodes. Furthermore, effective monitoring and convenient experiment control over large full-scale wireless networks remain challenging.

To avoid some of the difficulties with the full-scale wireless networks and to support realistic wireless experiments, researchers have built several miniaturized wireless testbeds [17, 7, 8]. The ORBIT project [17] implements a laboratory-

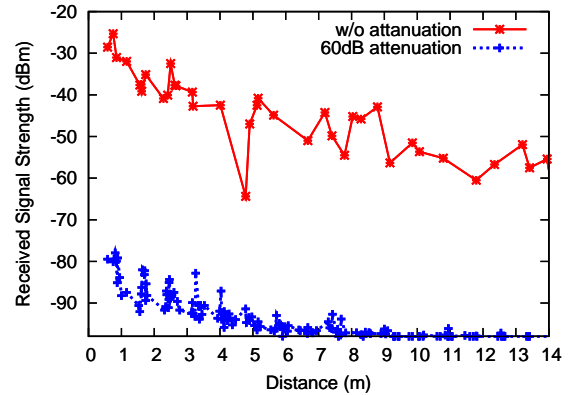


Figure 2: Received signal strength.

based wireless network emulator that uses a two-dimensional grid of 802.11 radio nodes that can be dynamically interconnected to form specified topologies with reproducible wireless channel models. The communication range of the wireless radios is reduced by creating a higher noise level at the receiver. Kansei [8] is a hybrid sensor network testbed that consists of heterogeneous testbeds with different wireless technologies. It also has an indoor grid of wireless nodes with 802.11 radio interfaces for wireless mesh network experiments. The MiNT testbed [7] is a miniaturized 802.11 testbed for mobile wireless networks with an integrated simulation and emulation framework. Both Kansei and MiNT testbed insert RF attenuators at both the transmitters and receivers to reduce the wireless communication range. Although various validation tests have been carried out by the developers of the miniaturized testbeds [17, 7, 8], the fidelity of the miniaturization techniques in a large-scale network is still not fully understood.

Recent studies have addressed the questions of whether one can design a performance-preserving network downscaling technique for the case of wireless networks [12, 13]. The authors of [12] show that it is possible to deploy the Kansei testbed [8] in a smaller area and yield the same performance, as if the same testbed was deployed in a larger area. This is achieved by appropriately reducing the transmission power of the wireless nodes and at the same time adjusting the separation distances between the nodes. Papadopoulos et al. [13] further investigate the scaling problem for arbitrary mobile ad hoc networks operating in outdoor environments in the presence of both large- and small-scale fading effects. They [13] also investigate time-downscaling for the wireless network, a technique that can be used to significantly expedite the time required to perform experiments with testbeds. In our work, we investigate the wireless network scaling problem with more realistic constraints, i.e., networks at different scale might have a different large-scale path loss exponent.

3. IVYNET TESTBED

In this paper, we conduct experiments using the IvyNet testbed [1]. IvyNet is a miniaturized wireless mesh network testbed deployed at ETH Zurich. In this section, we present an overview of the IvyNet testbed.

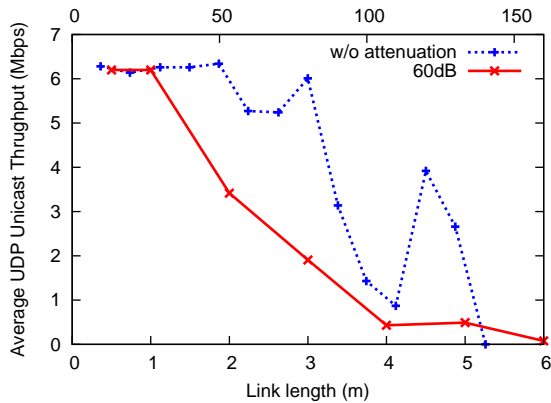


Figure 3: UDP throughput (11Mbps bitrate).

Hardware Architecture.

The IvyNet testbed consists of a collection of Soekris 4826 nodes. Each Soekris 4826 node is a single-board linux-based computer [3]. It uses a Geode SC1100UFH-266 CPU. It has 128 MB RAM and 64 MB on-board flash storage. There is a Compaq ZFM Micro USB 1.1 controller on the board for external persistent storage. Soekris 4826 has one fast Ethernet controller and 2 mini-PCI slots for extensions. We equip each node with one EnGenius NL-5354MP plus ARIES2 Mini PCI wireless network card. It uses the Atheros 5213 chip-set and supports 802.11 a/b/g modes. A low-gain 5dBi antenna is connected through a fixed attenuator of 30dB to the wireless card. We have 96 Soekris 4826 nodes, organized in a 4x24 grid, with an inter-node separation of 0.58 meter in the Y dimension and between 0.75 to 1.03 meters in the X dimension. As seen in Figure 1, all the nodes are put on IKEA shelves, and they are connected to control servers via the wired Ethernet. All the control and data traffics are communicated via the wired Ethernet to avoid interference with experimental traffics. Several columns of nodes are equipped with a second antenna. By setting the antenna diversity, we can adjust the attenuation level on these nodes by 30dB purely via software control.

Software Architecture.

To reduce the management overhead and get around the storage limitations of the Soekris board, we keep all the IvyNet wireless nodes from keeping their persistent state. Their memory and local disks keep only the volatile soft state. The hard state of the wireless nodes, such as system images and system configurations, is kept at the control server. To retain local disk modifications, wireless nodes must save such modifications on the remote persistent storage of the control server. In our current implementation, wireless nodes boot disk-less and stateless over the network and mount the root file system via NFS. The NFS export for the root file system is set to be read-only. Since all wireless nodes share the same system image on the file server, any changes we make to the system image on the file server are immediately visible to all the wireless nodes. Since most OS distributions run with writable media, a read-only root file system is not convenient for them. We solve this problem using the unionfs [22] file system. During system boot, a memory file system is created on each wireless node to store the temporal changes to their root file system. We use




Figure 4: Broadcast delivery ratio (1Mbps bitrate).

unionfs to combine the read-only root file system mounted via NFS and the memory file system into a read-writable root file system. Applications can write to the root file system, but all the changes are actually stored temporarily in the memory file system and will be lost once the node powers down. In addition, each wireless node has a private NFS writable directory on the file server; this directory serves as the private persistent storage for the wireless nodes.

Throughout the experiments in this paper, we configure the wireless network in IEEE 802.11b ad hoc mode. All the experiments are carried out either late at night or during weekends when there is the least external interference and its impact on the experimental results is negligible.

4. FIDELITY OF THE MINIATURIZED RADIO

In this section, we study the fidelity of the miniaturized network. The experimental results verify that the miniaturization technique effectively shrinks the physical space used by the network. At the same time, the miniaturized links preserve the same link-layer characteristics as the unattenuated links, and they do not alter the behavior of upper-layer protocols even in the presence of internal interferences.

4.1 Single Link

We first look at the behavior of a single wireless link. Figure 2 shows the received signal strength (RSS) distribution in IvyNet. Each point in the figure indicates the average RSS of all the links with the same length. For comparison, we also show the average RSS readings when we turn off attenuations on IvyNet nodes. In the experiments, we configure each node in turn to broadcast 180 100-byte packets. Meanwhile, all the nodes record the received signal strength. In this way, we collect the received signal strength profile for every link in the network. Although the Atheros chipset supports dynamic noise level adjustment, the observed noise level during the experiments is constant at $-98dBm$. From Figure 2 we can observe that with 60dB attenuation, the received signal strength is heavily reduced in IvyNet, and the signal strength degrades from $-79dBm$ to the noise level $-98dBm$ within 6 meters. Such heavily reduced transmission power leads to a much smaller radio communication distance that enables convenient multi-hop network setup in IvyNet. At the same time, it also brings up the question:




Figure 5: Link quality distribution.

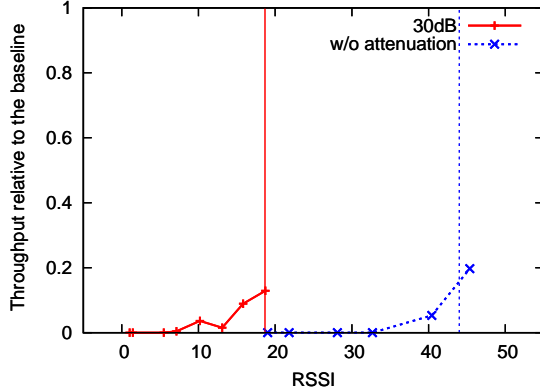


Figure 6: Strong interference.

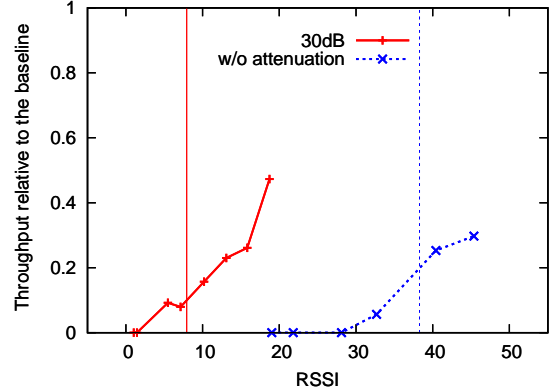


Figure 7: Intermediate interference.

is it still possible to preserve reasonable link-layer performance such that the upper-layer protocols show the same behavior over the miniaturized links as over unattenuated links? We use two metrics to evaluate the link-layer performance: link-layer throughput and link-layer packet delivery ratio.

Figure 3 shows the results for link-layer throughput. We measured the UDP throughput over the miniaturized links with different lengths in the IvyNet warehouse. In the experiments, the wireless cards are set to use the fixed 11Mbps bitrate with a transmission power of $19dBm$. Iperf is used to generate the UDP traffic. We also repeat the experiments for unattenuated links in an office building. Figure 3 depicts that both links show very similar link throughput distribution, although at different spatial scale. Both of them achieve the highest throughput when the link is short, e.g., shorter than 1 meter for miniaturized links. Over very long links, i.e., links longer than 4 meters for miniaturized links, only a few packets get through. In most of the cases UDP achieves intermediate throughput over links with medium length.

We also measure the packet delivery ratio at the link layer. In Figure 4, we present results for a single link. The transmitter sends broadcast packets with a rate smaller than the physical bit-rate, i.e., 1Mbps in the experiments, such that no packets are dropped due to queue overflow. At the receiver, we record the number of successfully received packets. The tests for the attenuated links are carried out in an office building. For all the tests, we set the transmission power to $19dBm$ and use packet size of 1000 Bytes. From Figure 4 we can observe the same as for the link-layer throughput: the miniaturized links preserve similar performance as the unat-

tenuated links at a different spatial scale. Additionally we measure the packet delivery ratio for all the links in IvyNet. Figure 5(a) summarizes the results. Each bar in the figure represents one bidirectional link. The end point of a bar represents the delivery ratio of a link measured at one direction. As comparison, Figure 5(b) shows the link delivery ratio distribution for Roofnet [5]. We can observe that in both IvyNet and Roofnet, there are many asymmetric links and also substantial amount of links with intermediate quality. Furthermore, IvyNet contains a much larger set of links with diverse quality. For most of the links from Roofnet, we can always find links from IvyNet that have a similar link quality.

4.2 Interference

In this section, we study the fidelity of the miniaturized radios in the presence of interference.

External interference.

In the miniaturized network, the radio signals are attenuated at both the transmitter and the receiver. When there are interferences from external interference sources, e.g., microwave ovens, cordless phones, and interfering signals from other 802.11 networks, the interference signals are attenuated only at the receivers, and the attenuation of the data signals is more than that of the interferences. Consequently, miniaturized testbeds are more vulnerable to external interferences. To reduce the impact of external interferences, we deploy IvyNet at a relatively isolated environment that is away from the campus WLAN networks and encourage users to repeat their experiments multiple times over different time period of a day.




Figure 8: Weak interference.

Internal Interference.

The “Hidden Terminal” problem is a major cause of internal interference in large wireless networks. We carried out experiments to study the effect of hidden terminal interference in the miniaturized network. The transmitter T sends broadcast traffic at the full rate to the receiver R . The hidden terminal I generates interference traffics by sending broadcast packets as fast as it can. We move T away from R . For each distance, we measure the broadcast throughput and the average RSSI (Received Signal Strength Indicator) at R . We repeat the measurements for four times. Firstly, we turn off the hidden terminal I . The measurement results from this round serve as the baseline reference. Then, we turn on I and place it at different distances away from R to create three different interference levels: strong interference, intermediate interference, and weak interference. We repeat the measurements for each of the three interference levels. We summarize the results with lines marked by $30dB$ in Figure 6, 7, 8. We then remove the attenuator at the receiver R and hence increase both the signal strength and the interference level by $30dB$. We repeat the experiments and summarized the results in Figure 6, 7, 8 with lines marked by *w/o attenuation*. The vertical lines in the figures represent the interference level. We can observe that the IvyNet wireless NIC hardware exhibits a strong capture effect [21, 19], i.e. the ability to decode one transmission when multiple transmissions collide. When the signal from the transmitter T is much stronger than the interference, as shown in Figure 8, the signal from T is always correctly received by R , even when the actual interference level is high, i.e., when there is no attenuation at the receiver R . In contrast, when the signal from the transmitter T is weaker than the interference, as shown in Figure 6, most of the packets from T are corrupted. Furthermore, for all the interference levels, the curves in the same figure have very similar shape. This observation indicates that the IvyNet NIC hardware resolves interference following strictly the SINR model, which predicts that a packet can be successfully decoded if the signal to interference plus noise ratio exceeds a certain threshold [18]. Consequently, when we scale down the power of a network, if the data signal and the interference are scaled down by the same factor, the scaled-down version of the network would experience similar interference effects as the original network, although the actual interference level at the two networks might be quite different. We did not ver-

ify our observation in the presence of multi-way interference because in practice multi-way interference has limited impact on network performance [6].

4.3 Case Study: TCP Performance over Multi-hop Path

Figure 9 shows TCP Reno performance over a 5-hop network [20]. In the 5-hop network, only adjacent stations can directly communicate with each other at 11 Mbps. Each data point in Figure 9 shows the average throughput from three runs of the same experiment. With TCP CWND clamping, we set TCP’s maximum congestion window limit. The RTS/CTS handshake is designed to coordinate the network transmissions when there are hidden terminals. By comparing the results with and without RTS/CTS, we can infer the significance of hidden terminals in our experiments. SampleRate [4] is used to select the bit-rate when the link-layer bit-rate adaptation is enabled.

Let us first consider the results without RTS/CTS. As we increase the CWND limit, we observe that if we enable the link-layer bit-rate adaptation, TCP achieves improved throughput until the network becomes overloaded. Further increasing the CWND limit hurts the TCP throughput. Whereas with fixed bit-rates, the TCP throughput drops only slightly before it stabilizes.

We can also observe the effect of RTS/CTS handshake from Figure 9. When the network operates with saturated contention level (i.e., CWND limit > 10 MSS) over adaptive bit-rate links, using RTS/CTS significantly improves the TCP throughput. The effect of RTS/CTS indicates that the hidden terminal problem is the major cause of packet losses in the experiments with adaptive bit-rate links when the TCP CWND limit is large. However, when TCP runs over fixed bit-rate links, the network operates with moderate contention. Fewer packets are lost due to hidden terminal problems and the benefit of using RTS/CTS does not compensate for its overhead. Hence using RTS/CTS over fixed bit-rate links is always detrimental for TCP in the experiments. Our results about the TCP performance over adaptive bit-rate multi-hop networks confirm the recent experiments over full-scale multi-hop wireless mesh networks [11].

5. SCALING THE MINIATURIZED TESTBED

5.1 Network Scaling in a Nutshell

The goal of miniaturized testbeds is to provide a platform for convenient and efficient testing of wireless protocols developed for real-world full-scale networks. Often we need to predict the real-world performance of a network protocol based on its behavior in the testbed. As the first step to conduct such a prediction, we need to know, for a wireless network, if it is possible to construct another network at a different power level that preserves the same performance and spatial structure. This is known as the network scaling problem [12, 13].

Scaling Links: Radio propagation models capture the behavior of wireless radios and estimate the received signal level as a function of the spatial relation between transmitter and receiver pairs. Large-scale propagation models characterize the received signal power over large transmitter-receiver separation distances. They indicate that the average received signal power decreases logarithmically with distance [16]:

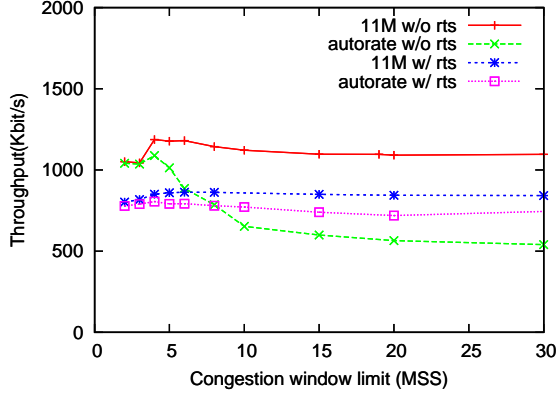


Figure 9: TCP performance over 5-hop path.

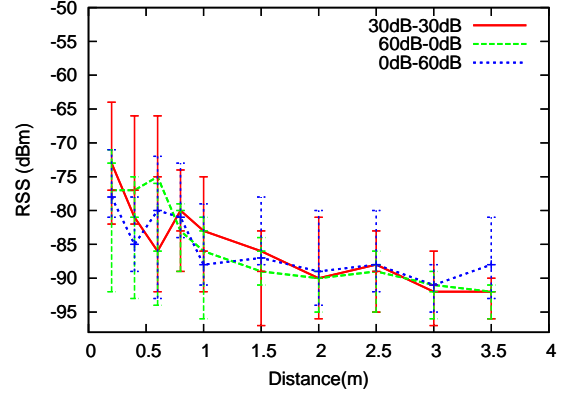


Figure 10: Effects of RF attenuation.

$$\overline{P_r}(d)_{dBm} = P_r(d_0)_{dBm} - 10n \log\left(\frac{d}{d_0}\right) \quad (1)$$

where $\overline{P_r}(d)$ is the average received signal power (in dBm) at distance d . n is called path loss exponent. $P_r(d_0)$ is the signal power (also in dBm) at the close-in reference distance d_0 . $P_r(d_0)$ is often computed from the transmission power P_t with the free space propagation model [16]:

$$P_r(d) = \frac{P_t G_t G_r \lambda^2}{(4\pi)^2 d^2 L} \quad (2)$$

Where λ is the wavelength. G_t and G_r describes the antenna gains. L is the system loss factor not related to propagation. When we take into account the variation of the surrounding environmental clutter, the received signal power at a certain distance from a transmitter is actually a random variable [16]:

$$P_r(d)_{dBm} = P_r(d_0)_{dBm} - 10n \log\left(\frac{d}{d_0}\right) + X_\sigma \quad (3)$$

X_σ describes the shadowing effects. It is a zero-mean Gaussian distributed random variable (in dB) with standard deviation σ (also in dB). Both n and X_σ are environment-specific and usually computed by field measurements. In this paper, we make the following assumptions:

1. Networks at different scales operate in similar environments that have the same shadowing effects and hence the same random variable X_σ .
2. All the stations in a network use the same transmission power.

To formulate the network scaling problem, we consider two networks N and N' . N has transmission power P_t and path loss exponent n . N' has transmission power αP_t and path loss exponent n' . Under the large-scale propagation model, we have the following theorems to map links from two networks:

THEOREM 1. *If N and N' have the same path loss exponent, given a link l in N with length d , we can construct a link l' with length d' in N' that preserves the same performance as the link l by scaling the length of l by a constant factor of $\sqrt[n]{\alpha}$, [12, 13]:*

$$d' = \sqrt[n]{\alpha} d \quad (4)$$

THEOREM 2. *If N and N' have different path loss exponents, there is a logarithmically linear mapping between link l and l' :*

$$\frac{d'}{d_0} = \sqrt[n']{\alpha} \left(\frac{d}{d_0}\right)^{\frac{n}{n'}} \quad (5)$$

PROOF. For both networks, we use the same close-in reference distance d_0 . From (2), we can get the received signal strength at the close-in reference distance in N' :

$$P_r'(d_0) = \alpha P_r(d_0)$$

Because X_σ is the same for N and N' , we only need to consider the average receiving signal strength (1). The received signal strength over the link in N' is given by:

$$\overline{P_r}'(d)_{dBm} = P_r(d_0)_{dBm} + 10 \log(\alpha) - 10n' \log\left(\frac{d'}{d_0}\right)$$

To achieve the same received signal strength over links from the two networks, we need:

$$10 \log(\alpha) - 10n' \log\left(\frac{d'}{d_0}\right) = -10n \log\left(\frac{d}{d_0}\right)$$

As the result, we get:

$$\frac{d'}{d_0} = \sqrt[n']{\alpha} \left(\frac{d}{d_0}\right)^{\frac{n}{n'}}$$

□

While large-scale propagation models capture the average path loss between a transmitter and receiver pair on large spatial scale, on a short time (or small spatial) scale the path loss between the transmitter-receiver pair may vary substantially due to small-scale fading [16]. If we consider only static networks, where nodes in the network are stationary, and the networks at different power levels have the same fading factor, e.g., the K -factor in the Ricean fading model, that depends on the structure of the environment, we can still apply Theorem 1 and Theorem 2 to scale links between different networks even in the presence of the small-scale fading effect. On the other hand, the IEEE 802.11 MAC




Figure 11: RSS profile (w/o attenuation, 19dBm txpower).




Figure 13: RSS profile (60dB attenuation, 1dBm txpower).




Figure 12: RSS profile (60dB attenuation, 19dBm txpower).

Setup	d_0 (m)	P_{d0} (dBm)	n
Without attenuation	0.58	-25.60	2.40
60dB attenuation 19dBm	0.58	-79.82	1.55
60dB attenuation 1dBm	0.58	-82.73	1.46

Table 1: Path loss parameters for networks at different power level.

protocol becomes seriously inefficient when the propagation delay increases above some threshold [14]. This fact limits the range of links that can be scaled faithfully. But for most cases in practice, the impact of the propagation delay difference is likely to remain insignificant.

Scaling Networks: Above we discussed issues related to how to scale links from one power level to another. We can easily extend this discussion to address the network scaling problem:

COROLLARY 1. *With the linear-mapping relation in Theorem 1, we can easily scale a network to another transmission power level by simply scaling all the links with the same constant factor [12, 13].*

COROLLARY 2. *In contrast, with the nonlinear relation in Theorem 2, it is not always feasible to construct a network with another transmission power level that preserves both the same radio performance and the same spatial structure as the original network.*

5.2 Scaling the IvyNet

IvyNet achieves reduced transmission range by introducing 30dB attenuation at both the transmitter and the receiver. Figure 10 shows the measured received signal strength for a transmitter-receiver pair when we add 30dB attenua-

tion at both the transmitter and the receiver, 60dB attenuation all at the transmitter, and 60dB attenuation all at the receiver. Not surprisingly, the three configurations yield similar results, no matter where the attenuation is inserted on the radio propagation path.

As the first step to map IvyNet to a real-world network, we conduct experiments to characterize the relationship between received signal strength and transmitter-receiver distance within IvyNet. We let each node in turn broadcast 180 100-byte packets. At the same time, all the nodes record the received signal strength (RSS). This way, we collect the received signal strength profile for every link in the network. We repeat the measurements for three different transmission power levels: maximum transmission power (iwconfig reading: 19dBm), maximum transmission power with 60dB attenuation, minimum transmission power (iwconfig reading: 1dBm) with 60dB attenuation. The results are summarized in Figure 11, Figure 12, and Figure 13.

We conducted linear regression tests to fit the RSS data to the large-scale path loss model (1). For the miniaturized network, we consider only links with a length less than 3 meters (7.14 distance units in Figure 12, 13). Table 1 summarizes the estimated parameters. We can observe that even with the same surrounding environment, networks with different transmission power levels experience quite different path-loss characteristics. A network with smaller transmission power tends to experience less path loss than a network with higher transmission power. For example, for the miniaturized network with maximum transmission power, the network has a path loss exponent of 1.55. When the transmission power is increased by 60dB, the path loss exponent is in turn increased significantly to 2.4. In addition, the path loss exponent from the miniaturized network is even smaller than the free space path loss exponent, and this fact indicates an ideal indoor environment with direct line-in-sight

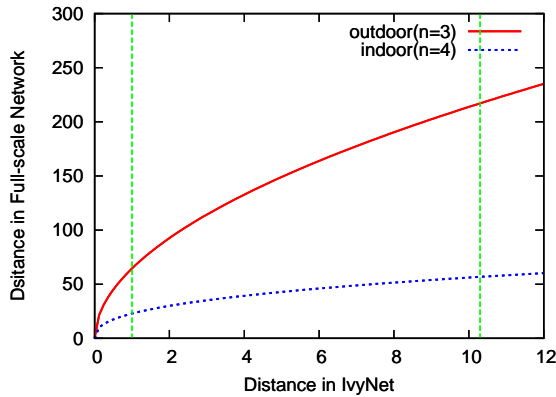


Figure 14: Distance mapping. The distances are normalized to d_0 (0.58m). Two vertical bars mark the shortest link (0.58m) and long link (6m) in IvyNet

paths. On the other hand, real-world large-scale networks are less likely to operate in such an ideal environment. Consequently, to scale IvyNet links into full-scale networks, we must consider the case when two networks have different path loss exponents, as shown in Theorem 2. By following Theorem 2, Figure 14 shows how we can scale IvyNet links into typical indoor and outdoor networks with the assumption that those networks have the same shadowing effects and fading factors as in IvyNet. As an example, the shortest link in IvyNet with a length of 0.58m achieves a similar performance as 23m links in an indoor environment, or 65m links in an outdoor urban environment.

In terms of scaling the whole network, Corollary 2 shows that we cannot directly map a network from one power level to another while preserving both signal-level characteristics and the spatial structure. This implies that for protocols that are influenced by the network spatial structure, e.g. routing protocols, we cannot always predict their performance over full-scale networks purely based on the testing over miniaturized wireless testbeds.

6. CONCLUSION

We present a validation study of IvyNet, a miniaturized wireless testbeds. Our results show that the miniaturized testbed is able to facilitate convenient experimentation over multi-hop networks within a small space, and at the same time preserves reasonable fidelity compared to the real-world full-scale networks for the majority, although exceptional cases exist when the miniaturized testbed should be used with caution.

Only commodity hardware has been used in the construction of IvyNet. Therefore we expect that the principle observations from IvyNet also apply to other miniaturized wireless networks.

7. REFERENCES

- [1] IvyNet wiki: <https://trac.id.ethz.ch/projects/ivyNet>.
- [2] Purdue wireless mesh network: <https://engineering.purdue.edu/mesh>.
- [3] Soekris engineering: <http://www.soekris.com/>.

- [4] J. Bicket. Bit-rate selection in wireless networks. Master's thesis, MIT, 2005.
- [5] J. Bicket, D. Aguayo, S. Biswas, and R. Morris. Architecture and evaluation of an unplanned 802.11b mesh network. In *Proc. MOBICOM*, 2005.
- [6] S. M. Das, D. Koutsonikolas, Y. C. Hu, and D. Peroulis. Characterizing multi-way interference in wireless mesh networks. In *WinTECH*, pages 57–64, 2006.
- [7] P. De, A. Raniwala, S. Sharma, and T. cker Chiueh. Mint: A miniaturised network testbed for mobile wireless research. In *Proc. INFOCOM*, 2005.
- [8] E. Ertin, A. Arora, and R. Ramnath. Kansei: A testbed for sensing at scale. In *Proc. IPSN*, 2006.
- [9] R. Karrer, I. Matyasovszki, A. Botta, and A. Pescape. Magnets: experiences from deploying a joint research-operational next-generation wireless access network testbed. In *TridentCom*, 2007.
- [10] D. Kotz, C. Newport, R. S. Gray, J. Liu, Y. Yuan, and C. Elliott. Experimental evaluation of wireless simulation assumptions. In *Proc. MSWiM*, 2004.
- [11] D. Koutsonikolas, J. Dyaberi, P. Garimella, S. Fahmy, and Y. C. Hu. On tcp throughput and window size in a multihop wireless network testbed. In *WinTECH*, 2007.
- [12] V. Naik, E. Ertin, H. Zhang, and A. Arora. Wireless testbed bonsai. In *Proc. WinMee*, 2006.
- [13] F. Papadopoulos and K. Psounis. Predicting the performance of mobile ad hoc networks using scaled-down replicas. In *Proc. ICC*, 2007.
- [14] R. Patra, S. Nedeveschi, S. Surana, A. Sheth, L. Subramanian, and E. Brewer. Wildnet: Design and implementation of high performance wifi based long distance networks. In *Proc. NSDI*, 2007.
- [15] K. Pawlikowski, H. Jeong, and J. Lee. On credibility of simulation studies of telecommunication networks. In *IEEE Communications Magazine*, 2002.
- [16] T. S. Rappaport. *Wireless Communications: Principles and Practice*. Prentice Hall, 2001.
- [17] D. Raychaudhuri and I. Seskar. Overview of the orbit radio grid testbed for evaluation of next-generation wireless network protocols. In *Proc. WCNC*, 2005.
- [18] C. Reis, R. Mahajan, M. Rodrig, D. Wetherall, and J. Zahorjan. Measurement-based models of delivery and interference in static wireless networks. In *SIGCOMM*, 2006.
- [19] S. Seshan, X. Liu, and P. Steenkiste. Interference-Aware Transmission Power Control for Dense Wireless Networks. In *Proceedings of the Annual Conference of ITA*, 2007.
- [20] Y. Su, P. Steenkiste, and T. Gross. Performance of tcp in multi-hop access networks. In *Proc. IWQOS*, 2008.
- [21] K. Whitehouse, A. Woo, F. Jiang, J. Polastre, and D. Culler. Exploiting the capture effect for collision detection and recovery. In *Proc. EmNets*, 2005.
- [22] C. P. Wright, J. Dave, P. Gupta, H. Krishnan, E. Zadok, and M. N. Zubair. Versatility and unix semantics in a fan-out unification file system. Technical report, 2004.