

“RF in the Jungle”: Effect of Environment Assumptions on Wireless Experiment Repeatability

Ryan Burchfield, Ehsan Nourbakhsh, Jeff Dix, Kunal Sahu, S. Venkatesan, Ravi Prakash
Department of Computer Science, University of Texas at Dallas
Richardson, TX 75080, U.S.A.
Contact E-mail: ehsaan@utdallas.edu

Abstract—Most researchers conduct wireless networking experiments in their laboratory or similar indoor environments. Such environments are veritable *RF jungles*, especially when we consider the ISM bands. In this paper we examine and test several common explicit and implicit assumptions that researchers tend to make about the wireless environment. Although these assumptions are acknowledged by most researchers, the extent of their impact is often underestimated. We find that because the environment is always in flux, it is almost impossible to reproduce the results of an experiment. Hence, there is a high risk of misinterpreting the data obtained from such experiments. Through this paper we try to caution experimenters against such risky assumptions when they venture into the *RF jungle*. After a successful proof-of-concept experiment, we advocate the use of wireless networking testbeds that provide experimenters better control over the RF environment by using coaxial cables, programmable attenuators and power dividers/combiners.

I. INTRODUCTION

Wireless networking research has been significantly influenced by the tools used by researchers for their experiments. Any problem with these experiments can lead to wrong assumptions and conclusions. For example, the ns-2 [1], Glomosim [2] and Opnet [3] simulators have had a profound impact on Mobile Ad-Hoc Network research. However, Cavin et al. [4] reported significant differences in the results generated by these simulators for the same experiment, thus raising questions about the accuracy of all these simulators. Kotz et al. [5] address some of these simplifying assumptions in simulations.

Is running experiments on an actual wireless network testbed more reliable? After all, the assumptions that simulators make about various components of the network can be done away with. If this was as straightforward as it sounds, why do different experimenters obtain significantly different results, and draw conflicting conclusions, from similar experiments? For example, Mishra et al. [6] stated that contrary to prevailing views, satisfactory network throughput can be obtained with spatially adjacent IEEE 802.11 networks operating on overlapping (and, therefore, interfering) channels. On the other hand, Gummadi et al. [7] reported that a small amount of interference can severely degrade the performance of commodity IEEE 802.11 devices. Similarly Fuxjager et al. [8] demonstrate the presence of detrimental interference between non-overlapping IEEE 802.11 channels.

Interestingly enough, these seemingly conflicting inferences are all correct — under certain circumstances. Often factors

external to the experiment (like background traffic or noise, unanticipated interaction between various devices, measurement inaccuracies, etc.) affect the results significantly. While most researchers assert that they are aware of the impact of these factors, we feel that in general they are underestimated and frequently assumed to be insignificant when analyzing experiment data. The goal of this paper is to identify and question the validity of the following common environmental assumptions, explicit and implicit, that may adversely affect the repeatability of a wireless experiment.

- 1) My building represents a typical home/office environment.
- 2) Experiments conducted at night benefit from less wireless interference.
- 3) No wireless networks detected equals no interference.

In the following sections we experimentally test each assumption and caution against the use of the assumption due to the varying results that we obtained. For each assumption we provide a recommendation for strengthening the experiment’s repeatability and accuracy. Finally, we advocate an interference resistant framework for conducting repeatable wireless networking experiments. Hopefully, this will pave the way for careful design, planning and instrumentation of wireless networking experiments so that researchers can independently conduct the same experiment and reproduce each others’ results.

II. ASSUMPTION: MY BUILDING REPRESENTS A TYPICAL HOME/OFFICE ENVIRONMENT.

In order to observe the impact of the building environment, we conducted simple wireless experiments in different locations. We used the same hardware and same physical topological set-up to minimize the number of variables involved. Furthermore, we started the 8-9 hour experiments at approximately the same time of day. The experiments were comparable to the ones performed by [6] and focused on the performance of wireless networks as a function of channel overlap and physical distance.

A. Experiment Setup

The hardware included two pairs of nodes, each pair consisting of a Linksys WRT54GL wireless access point and a Hawking HWUG1 Wireless-G USB Wi-Fi adapter, both operating in IEEE 802.11b mode. This access point was selected for its

controllability. The Wi-Fi adapter was selected because it uses the same Ralink RT73 chipset as many common Linksys cards use. The distance between the access point and the wireless adapter was roughly 1 meter for both pairs at all times. One pair (denoted as Pair A) was stationary at a particular location and was always operational on Wi-Fi channel 6. The other pair (denoted as Pair B) communicated over channel x , where x was selected from Wi-Fi channels 1 through 11. As a result, the channel separation between the two pairs is equal to $x - 6$, and varies between -5 and $+5$. We increased the physical separation between the two pairs in steps of 10 meters by moving Pair B. At each distance, we performed the following experiment three times and then used the average of these results.

The laptop with the Hawking adapter in Pair A generated 11 Mbps of CBR (Constant Bit Rate) UDP traffic for a duration of 60 seconds. The traffic data was generated by Iperf and sent wirelessly to the access point where the data was relayed to a server via Ethernet. Similarly, the laptop in Pair B sent 11 Mbps CBR UDP traffic to a different server via the second AP. Thus, there were two concurrent wireless sessions in progress.

B. Experiments

To evaluate the assumption that our building represents a typical office environment we chose to conduct our experiment in a long hallway that lies partly within the Computer Science building (CS) and partly within the Electrical Engineering building (EE). The two buildings are connected with a 25 meter catwalk-bridge and both buildings are equipped with a university managed IEEE 802.11b wireless network. We tested the assumption by conducting the experiment twice; once with Pair A fixed in the CS building and again with Pair A fixed in the EE building.

1) *CS to EE*: We stationed Pair A at the CS end and moved Pair B towards the EE building. We present the results of this experiment as a 3D graph in Figure 1. The graph compares channel separation on the X-axis, increasing distance on the Y-axis and sum of the throughput for Pair A and Pair B on the Z-axis. Clearly, the impact of channel separation was not symmetric. When channel separation was $-i$, the combined throughput tended to be consistently higher than when the channel separation was $+i$ even though the amount of channel overlap was the same in both cases. This is not the intuitive result that we expected.

2) *EE to CS*: Hoping to reproduce the trends of Figure 1, we stationed Pair A at the EE side and moved Pair B towards the CS building in the same hallway. The experiments' results are shown in Figure 2.

C. Evaluation

As is evident from Figures 1 and 2, we were unable to reproduce the trends in the same hallway for the same set of experiments by simply reversing the direction. Since the only variable that changed between experiments was the location of the fixed pair, we can see that a small change in the location of the participating nodes had a substantial impact

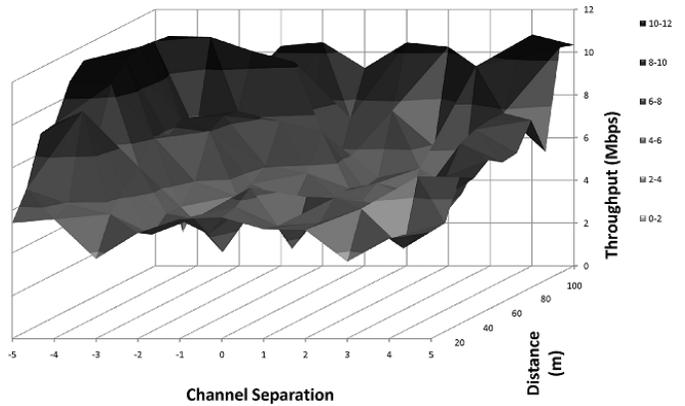


Fig. 1. Indoor CS to EE Experimental Results

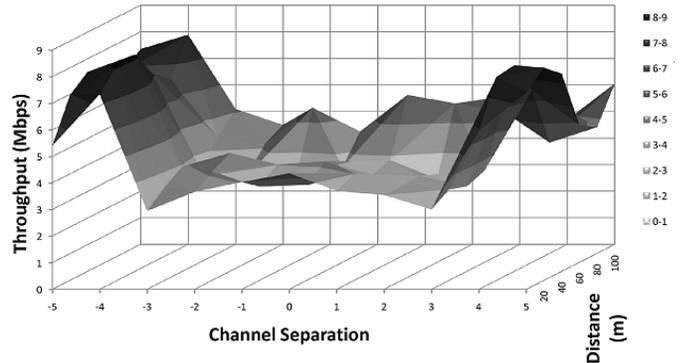


Fig. 2. Indoor EE to CS Experimental Results

on the outcome. Clearly, one cannot arbitrarily claim that one of these buildings, both with the same university managed wireless network, represents a typical office environment while the other does not. The problem can be exacerbated when one considers buildings at different institutions with heterogeneous wireless networks.

D. Recommendation

We caution against the use of the assumption that a researcher's building represents a typical environment. From our observations we realized the need to closely monitor the environment to avoid the pitfalls and wrong conclusions made in wireless networking research. We further recommend that researchers conduct experiments in multiple locations and buildings to verify that the results are comparable.

III. ASSUMPTION: EXPERIMENTS CONDUCTED AT NIGHT BENEFIT FROM LESS WIRELESS INTERFERENCE.

All the experiments described in Section II-B were conducted during the day. Would the results be different during the night? That seems to be the implicit assumption made by several experimenters. It is common to find statements similar to "...we ran our experiments when there was little external traffic ..." [7] [9]. For some studies there is no information about the level of background interference (for example [6]).

A. Experiment

To measure temporal changes in background wireless activity we used a Wi-Spy [10] spectrum analyzer to record activity continuously for one week in a laboratory in the CS department, adjacent to the CS-EE hallway. We found Wi-Spy to be a useful tool in this situation as it provides the amount of energy in terms of Received Signal Strength (RSS) in dBm present in the 2.4 GHz ISM band regardless of the PHY modulation used. The recordings are shown in Figure 3. The vertical axis is RSS, while the horizontal axis is the time of day. The figure shows the RSS value averaged over the entire 2.4 GHz ISM band, as well as the value for IEEE 802.11 channel 11 (the most active channel in our location). Each data point represents the average value for a one minute interval.

For the same week we also measured the wireless activity on the 2.4 GHz band in an apartment. This apartment complex is occupied primarily by university students; a community that is a heavy user of Wi-Fi networks. The activity level in the apartment and the laboratory are compared in Figure 4.

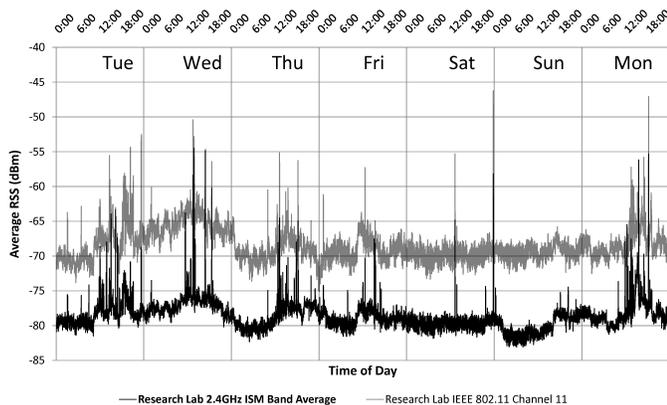


Fig. 3. Week of wireless activity in a laboratory.

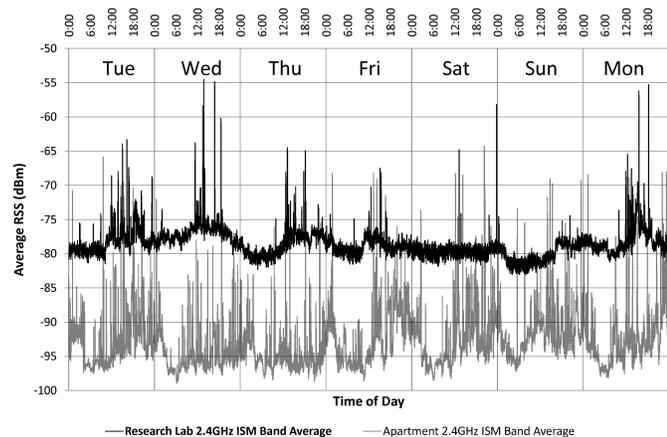


Fig. 4. Week of wireless activity in a laboratory and apartment.

B. Evaluation

Immediately, a few things become apparent. First, for the entire week the average RSS for the ISM band is above

-85 dBm, which is significantly above the sub -110 dBm background activity we later (Section V-A2) measure outdoors. Hence, there was no consistently good time of the day to conduct indoor experiments with minimal interference.

From Figure 4 we see that the average RSS in the apartment is about 15dB lower than that in the laboratory, yet about 20dB higher than what was recorded later outdoors (Section V-A2). Even the quietest time period at night in our laboratory corresponds to a very high level of background activity that can significantly bias the experiments. Second, regardless of the time of day wireless activity in the residential environment exhibited significantly larger fluctuations over very short periods of time, compared to the laboratory.

C. Recommendation

From our observations the indoor environment is inherently noisy, with unpredictable activity fluctuations, regardless of the day of week or time of day. Rather than assuming that experiments conducted at night avoid interference we recommend a three phased monitoring solution for experiments to be conducted at any time. 1) Characterize the environment to determine: average background activity, heavily occupied channels, etc. before an experiment is conducted. 2) Monitor the environment for changes during an experiment. 3) Verify that critical experiment data points do not correlate to unintentional changes in the environment.

IV. ASSUMPTION: NO WIRELESS NETWORKS DETECTED EQUALS NO INTERFERENCE.

Sometimes what we measure during an experiment can significantly influence what we conclude or do not conclude from the experiment. To illustrate this we characterize various measurement tools into three categories: 1) MAC SNR, 2) Network packet capture, 3) Spectrum analysis. We then conduct a sample experiment to compare the information gathered with each tool.

NetStumbler, used with a Prism wireless card, is an example of a MAC SNR measurement tool. NetStumbler scans Wi-Fi channels capturing beacons and reports the signal to noise ratio of that frame. Wireshark is an example of a network packet capture tool. Lastly, Wi-Spy is an example of a spectrum analysis tool because it scans the 2.4GHz ISM band in 328kHz increments and measures the raw received signal power. We compare the output of these tools in the following experiment.

A. Experiment

We placed a Wi-Fi access point and a Wi-Fi equipped laptop about six meters apart in our laboratory. We also had a pair of nodes, about six meters apart, using the IEEE 802.15.4/ZigBee protocol stack.

For the first sixty seconds of the experiment both pairs of nodes were idle. During seconds 60-120 only the Wi-Fi nodes were communicating with 11Mbps CBR traffic sent from the laptop to the wireless AP. The Wi-Fi pair was on channel 11, same as the nearest university AP (we refer to this AP as the remote dominant AP). At the 120 second mark the pair

of ZigBee nodes started exchanging maximum CBR traffic on the IEEE 802.15.4 channel 23, which has the maximum overlap with Wi-Fi channel 11. At the 180 second mark we started a microwave oven at a distance of about five meters from the laptop. We conducted three runs of this experiment and observed channel activity using two Wi-Spys (one near each Wi-Fi node), Wireshark on each of the Wi-Fi nodes, and NetStumbler running on a PC placed in the middle of the laboratory. Using NetStumbler we monitored the university network’s dominant AP as well as our AP. We report the observations of one run of our experiment in Figure 5. The other two runs had similar results.

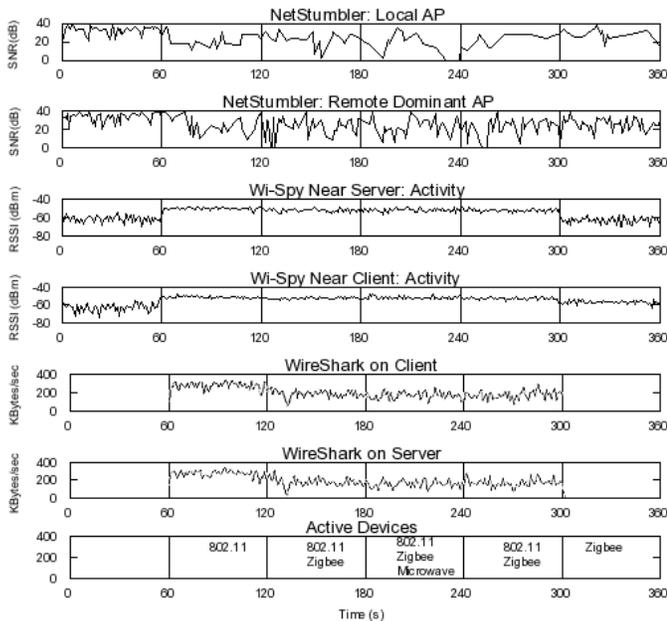


Fig. 5. Comparison of Different Wireless Measurement Tools

B. Evaluation

During seconds 120-180, when the ZigBee pair is contending for the channel with the Wi-Fi pair, Wireshark shows almost a 33% decline in activity (in kBytes/s). However, Wi-Spy shows there is no decline in overall wireless activity on the channel. Any perception of reduced activity is explained by the fact that Wireshark cannot detect ZigBee traffic. If the experimenter unknowingly conducts an experiment during this period, the decreased Wi-Fi throughput may be incorrectly attributed to the protocol under test.

During the 180-240 second period there is maximum contention/interference on the medium because in addition to the two communicating pairs the microwave oven is also active. However, Wireshark fails to convey this information. While the Wi-Spy trace for the entire 2.4 GHz band does show increased energy over the entire band, if one were to focus only on the RSS value for Wi-Fi channel 11 (as shown in Figure 5), the interference due to the microwave oven is likely to be missed.

Finally, the limitations of NetStumbler as a tool for measuring wireless activity also becomes apparent from this ex-

periment. When the node pairs start communicating, or the microwave oven is operational, the SNR values of both access points show significant fluctuations. But, there is no additional information to reflect the level of channel activity.

C. Case Study

During both the indoor hallway experiments (Section II-B) we encountered difficulty obtaining measurements for wireless channels 1-3 when Pair B was within a 20m region of the EE building. Within this region the wireless client was not able to maintain a steady connection to its AP placed 1 meter away. We immediately acknowledged there must be very strong interference in this area. To test this hypothesis we used NetStumbler to scan for nearby APs with a very strong signal around these channels. Our scan turned up no APs.

Next, using Wi-Spy we walked the halls looking for the source of a signal near Wi-Fi channel 2. We found a wireless video transmitter used for security in a distant EE laboratory. The camera consumed the available bandwidth, without performing CSMA/CA as Wi-Fi devices do, causing communication to suffer on Wi-Fi channels 1-3. *This further illustrates that external factors, usually unknown to the experimenter, can severely bias experiment results.*

D. Recommendation

No single metric is guaranteed to truly reflect the level and nature of wireless activity on a given channel. Furthermore, as we found in the case study, MAC SNR measurement and packet capture tools provide only standard related measurements, which is the cause of the implicit assumption “No wireless networks detected equals no interference”. So, experimenters need to monitor several network and channel parameters and look at all of them concurrently to draw inferences from their experiments.

V. COAX TESTBED

In some cases, recommendations from previous sections may fall short of providing a repeatable experiment because the interference may be unavoidable. Due to this we advocate the use of an interference-resistant coax based testbed for wireless experiments. In this discussion we walk the reader through the setup and execution of a proof-of-concept experiment. The experiment’s result trends are verified via comparison to results from the same experiment conducted outdoors.

A. Experiment

The coax testbed was verified through a proof-of-concept experiment designed to mimic the hallway experiments conducted indoors in Section II-B.

1) *RF Topology:* As shown in Figure 6, the coax test platform is symmetrical around the attenuator A1. Each set consists of a 4-way Wilkinson style RF splitter/combiner (JFW Industries part 50PD-560) which is connected to the Wi-Fi card and AP. A Wi-Spy is also connected to each side in order to measure the interference from the external environment that leaks in. The ‘parent’ port of each splitter is connected to a

programmable attenuator (JFW Industries part 50P-1501) with attenuation range of 0dB to 127dB.

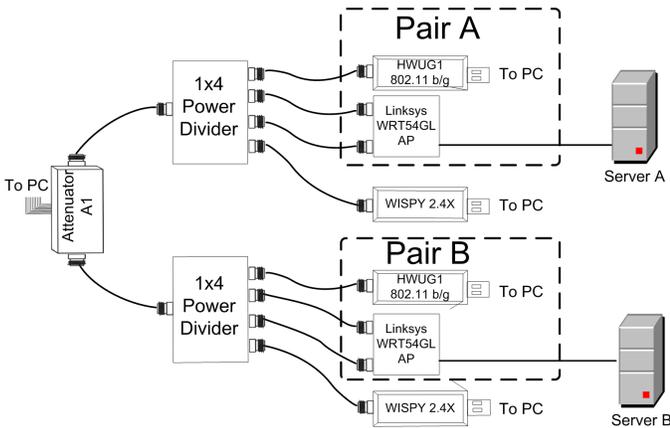


Fig. 6. Coax Test Platform

With all devices connected but powered off, we monitored the two Wi-Spys. One Wi-Spy showed a -90dBm interfering signal leaking into the coax network. By wrapping aluminum foil around connected wireless devices, we reduced the interference to less than -110dBm , which is the weakest signal accurately detected by Wi-Spy.

2) *Path loss model*: Recall that the experiment in Section II-B calls for an increasing distance between Pair A and Pair B. In order to emulate distance, we have derived a path loss model by utilizing three laptop computers and an access point on the university softball field. The softball field was verified to be sufficiently free of 2.4GHz ISM interference by using Wi-Spy, which recorded no signals stronger than -110dBm . Additionally, no wireless networks were detected. Two of the laptops were configured to communicate with 11Mbps CBR traffic. On the third laptop Wi-Spy recorded the signal level received when carried in 10m increments from 0m to 350m. To apply this model to our testbed we measured and recorded coax cable and splitting losses as system losses. Therefore, for a given distance the attenuator is set to be the outdoor loss for that distance minus the system loss of 21dB.

3) *Automated Testbed Programming Tool (ATPT)*: With a verified ‘clean’ coax network, experiments were initiated in the same manner as hallway experiments in Section II-B. However, the verification procedure had the time saving benefit of a custom Java application, ATPT. We developed ATPT to take a script file defining events (start Iperf, change attenuation, change AP’s channel, etc.), and run all the events in the specified order. Once initiated, ATPT performed 28 hours worth of data collection with no human involvement.

B. Evaluation

To this point we have constructed a coax-based wireless testbed and emulated the hallway experiments from Section II-B. However, because the coax environment is interference free we do not expect to receive identical results to our previous unshielded indoor experiments. We feel that to

evaluate the behavior of the coax testbed we should compare its results to a similar interference free environment, such as what we found on the softball field. Figures 7 and 8 compare the results of these two experiments. We see two very similar trends between the experiments. We later verified the outdoor experiments by conducting them in a public park, which we also found free of wireless interference. The results of both outdoor experiments were comparable.

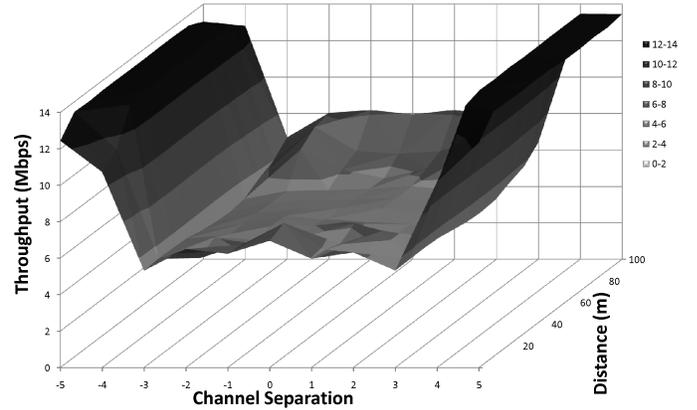


Fig. 7. Coaxial Testbed Experimental Results

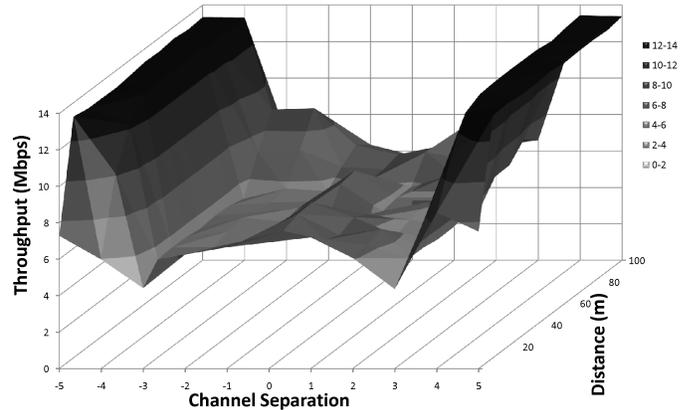


Fig. 8. Outdoor Experimental Results

C. Recommendation

The coax testbed proved to be a valuable tool for conducting wireless experiments due to its interference free environment and time saving nature. Furthermore, due to the lack of interference, we can assert that the coax testbed facilitates highly repeatable experiments. It is our recommendation that researchers seek similar interference free and repeatable environments for experiments. Experiments in repeatable environments eliminate the uncertainty in parsing results and attributing trends to protocol behavior. Lastly, experiments conducted on this coax testbed avoid all of the misleading assumptions that we explored in the preceding sections. Not all experiments can be conducted on the simple coax RF topology we presented here. To that end we are working on a larger,

topologically flexible version of the testbed. Furthermore, we are investigating controlled methods to inject noise and emulate multi-path fading on the coax testbed.

VI. RELATED WORK

Several wireless networking testbeds have been developed for research purposes. These include Netbed [11], Kansei [12], Trio [13] and other testbeds at U.C. Berkeley [14]. These testbeds, while providing valuable insights about wireless networking, are not isolated from the surrounding environment and are susceptible to background noise and interference. For a thorough study of wireless networking testbeds refer to [15].

We concentrate on testbeds that explicitly address repeatability of wireless experiments. ORBIT [16] project implements a 20×20 grid of 802.11 wireless nodes. Some studies like [17] have addressed the repeatability issues of experiments in ORBIT and proposed improvements. The Illinois Wireless Wind Tunnel seeks repeatability by scaling down RF experiments so that they may be conducted in a single anechoic chamber that shields interference and resists multipath [18].

A different approach has been taken by Judd et al. in the wireless emulator built at Carnegie Mellon University [19]. The emulator converts RF signals into digital streams, manipulates the digital streams according to propagation models using an FPGA+DSP engine, and then converts the bit stream back to RF signals. All the hardware is housed in a single server rack. However, there are some scalability concerns. Emulating a ten node network requires 180 Gbps of data to be streamed from/to the digital signal processor and 15 billion scaling operations have to be performed per second [19]. While the wireless emulator occupies significantly less space than ORBIT, it can be difficult to expand it to a very large size. Also, if some unknown RF characteristics have not been coded into the logic of the DSP engine, then their impact cannot be studied.

In [20] a similar approach for crafting a small testbed using coaxial components is discussed.

VII. CONCLUSION

We have shown that several commonly used explicit and implicit assumptions about wireless experiments are often false and may lead to unpredictable and unreliable results. This is due to unpredictable changes in background traffic and interference. We caution that researchers should overlook the role of environmental factors at their own peril. Also, unless we are careful, the tools we use to monitor our experiments may fail to record critical information and may even lead us to draw wrong conclusions. Hence, there is a critical need to develop rigorous experimental methodology so that different wireless networking research groups can independently reproduce results of experiments and validate various theories. Towards this end we advocate the use of wireless networking testbeds that provide experimenters control over the environment by using a combination of coaxial cables, programmable attenuators and power dividers/combiners.

REFERENCES

- [1] "The Network Simulator - ns-2." [Online]. Available: <http://www.isi.edu/nsnam/ns/>
- [2] "GloMoSim: Global Mobile Information Systems Simulation Library." [Online]. Available: <http://pcl.cs.ucla.edu/projects/gloimosim/>
- [3] "OPNET Technologies." [Online]. Available: <http://www.opnet.com>
- [4] D. Cavin, Y. Sasson, and A. Schiper, "On the Accuracy of MANET Simulators," in *Proceedings of the Second ACM International Workshop on Principles of Mobile Computing (POMC)*, 2002, pp. 38–43.
- [5] D. Kotz, C. Newport, R. S. Gray, J. Liu, Y. Yuan, and C. Elliott, "Experimental evaluation of wireless simulation assumptions," in *MSWiM '04: Proceedings of the 7th ACM international symposium on Modeling, analysis and simulation of wireless and mobile systems*, 2004, pp. 78–82.
- [6] A. Mishra, V. Shrivastava, S. Banerjee, and W. Arbaugh, "Partially Overlapped Channels Not Considered Harmful," in *Proceedings of SIG-Metrics/Performance: Joint International Conference on Measurement and Modeling of Computer Systems*. ACM, 2006.
- [7] R. Gummadi, D. Wetherall, B. Greenstein, and S. Seshan, "Understanding and mitigating the impact of rf interference on 802.11 networks," in *Proceedings of SIGCOMM'07*. ACM, 2007.
- [8] P. Fuxjager, D. Valerio, and F. Ricciato, "The myth of non-overlapping channels: interference measurements in IEEE 802.11," *Proceedings of the Conference on Wireless on Demand Network Systems and Services (WONS)*, pp. 1–8, 2007.
- [9] J. Padhye, S. Agarwal, V. Padmanabhan, L. Qiu, A. Rao, and B. Zill, "Estimation of link interference in static multi-hop wireless networks," in *IMC '05: Proceedings of the 5th ACM SIGCOMM conference on Internet Measurement*, 2005, pp. 28–28.
- [10] "Wi-Spy 2.4x: Spectrum Analyzer for 2.4GHz ISM Band." [Online]. Available: <http://www.metageek.net/products/wi-spy>
- [11] B. White, J. Lepreau, and S. Guruprasad, "Lowering the Barrier to Wireless and Mobile Experimentation," in *In Proceedings of HotnetsI*, 2002.
- [12] E. Ertin, A. Arora, R. Ramnath, V. Naik, S. Bapat, V. Kulathumani, M. Sridharan, H. Zhang, H. Cao, and M. Nesterenko, "Kansei: a testbed for sensing at scale," *Proceedings of the fifth international conference on Information processing in sensor networks*, pp. 399–406, 2006.
- [13] P. Dutta, J. Hui, J. Jeong, S. Kim, C. Sharp, J. Taneja, G. Tolle, K. Whitehouse, and D. Culler, "Trio: enabling sustainable and scalable outdoor wireless sensor network deployments," in *Proceedings of IPSN '06*, 2006, pp. 407–415.
- [14] "Soda Hall wireless and sensor network testbeds." [Online]. Available: <http://www.millennium.berkeley.edu/sensornets/>
- [15] W. Kiess and M. Mauve, "A survey on real-world implementations of mobile ad-hoc networks," *Ad Hoc Netw.*, vol. 5, no. 3, 2007.
- [16] 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," *Wireless Communications and Networking Conference, 2005 IEEE*, vol. 3, 2005.
- [17] S. Ganu, H. Kremo, R. Howard, and I. Seskar, "Addressing repeatability in wireless experiments using ORBIT testbed," in *Proceedings of TRIDENTCOM 2005*, 2005, pp. 153–160.
- [18] N. Vaidya, J. Bernhard, V. Veeravalli, P. Kumar, and R. Iyer, "Illinois wireless wind tunnel: a testbed for experimental evaluation of wireless networks," in *E-WIND '05: Proceedings of the 2005 ACM SIGCOMM workshop on Experimental approaches to wireless network design and analysis*, 2005, pp. 64–69.
- [19] G. Judd, "Using Physical Layer Emulation to Understand and Improve Wireless Networks," Ph.D. dissertation, Carnegie Mellon University, School of Computer Science (available as Tech Report CMU-CS-06-164), 2006.
- [20] J. Kaba and D. Raichle, "Testbed on a Desktop: Strategies and Techniques to Support Multi-hop MANET Routing Protocol Development," in *Proceedings of MobiHoc '01*. ACM, 2001.