Hindawi Publishing Corporation
EURASIP Journal on Wireless Communications and Networking
Volume 2010, Article ID 527847, 12 pages
doi:10.1155/2010/527847

Review Article

Techniques for Improving the Accuracy of 802.11 WLAN-Based Networking Experimentation

Marc Portoles-Comeras, Josep Mangues-Bafalluy, and Manuel Requena-Esteso

Parc Mediterrani de la Tecnologia (PMT), Centre Tecnològic de Telecomunicacions de Catalunya (CTTC),
Avenue Carl Friedrich Gauss 7, Castelldefels 08860, Barcelona, Spain

Correspondence should be addressed to Marc Portoles-Comeras, marc.portoles@cttc.cat

Received 17 June 2009; Revised 19 October 2009; Accepted 15 December 2009

Academic Editor: Faouzi Bader

Wireless networking experimentation research has become highly popular due to both the frequent mismatch between theory and practice and the widespread availability of low-cost WLAN cards. However, current WLAN solutions present a series of performance issues, sometimes difficult to predict in advance, that may compromise the validity of the results gathered. This paper surveys recent literature dealing with such issues and draws attention on the negative results of starting experimental research without properly understanding the tools that are going to be used. Furthermore, the paper details how a conscious assessment strategy can prevent placing wrong assumptions on the hardware. Indeed, there are numerous techniques that have been described throughout the literature that can be used to obtain a deeper understanding of the solutions that have been adopted. The paper surveys these techniques and classifies them in order to provide a handful reference for building experimental setups from which accurate measurements may be obtained.

1. Introduction

The widespread availability of wireless products at low cost, mainly IEEE 802.11 WLAN cards, has allowed wireless networks, both commercial and end-user driven, to appear all over the world at an unprecedented pace. Additionally, mismatches between theoretical/simulation and experimental work made the wireless networking research community realize the need to turn-highly experimental, which was also fostered by the availability of such low-cost equipment. The use of data from research testbeds and experimental measurement campaigns has revealed numerous unforeseen aspects about the actual functionality of WLAN networks.

However, experimentation results have to be handled with care. Placing wrong assumptions on the actual behavior of experimentation tools can lead to dismissing useful research results because they do not match the expected theoretical behavior. An improper calibration of networking tools or, simply, a misunderstanding of their actual behavior can render invalid conclusions out of measurement results.

This paper surveys recent literature related to wireless network experimentation and presents the multiple performance issues that have been encountered when dealing with actual networking tools. Such issues may range from card misbehavior to unexpected interaction of the wireless hardware with the environment (e.g., interference). In some cases, these findings were not predicted in advance because they are specific to particular implementations. However, in other cases, it turns out that even the simplest wireless setups place additional requirements to networking nodes and tools with respect to those posed to wired ones. As a consequence, they need a conscious calibration before operation, which is even more important in more complex wireless setups that include multihop and/or multiradio scenarios.

In many cases, the detection of specific performance issues related to wireless networking starts from unexpected measurement observations. Unpredicted throughput or loss measurements guide researchers to start a deeper analysis of the behavior of wireless devices. This suggests the usage of workload and loss measurements to detect and/or diagnose potential problems that may appear in wireless deployments.

For this reason, this paper describes reference setups that can be used to analyze the behavior of the solutions to be used in order to detect in advance any performance anomaly.

Once performance issues are detected, researchers have been using a variety of techniques to diagnose the exact origin of the behavior observed. The paper continues with a survey and classification of these diagnosis techniques. Some of the techniques described involve costly equipment or complex setups but, in the majority of cases, they can be fulfilled with typical equipment present in WLAN laboratories.

The main contributions of this paper follow:

(i) raising awareness of the need for calibration when selecting the hardware and software to be used in wireless networking deployments,

(ii) discussing negative effects that the lack of understanding of wireless networking components might have on the interpretation of results,

(iii) surveying and classifying the techniques to be used to characterize the components used in a wireless networking testbed with these techniques being used to either detect or diagnose unexpected behavior of networking components,

(iv) proposing a series of steps to methodologically apply prior to any wireless networking deployment in order to evaluate the adequacy and limitations of the solutions adopted.

Up to our knowledge, there has been no previous work on categorizing the artifacts that may arise in wireless network experimentation due to hardware misbehavior or its unexpected interaction with the wireless environment. This paper compiles and tries to provide a systematic view of this field with the goal of helping wireless networking experimenters to obtain more reliable setups and measurements.

The rest of the paper is organized as follows. Section 2 reviews the sources of potential problems found in the wireless networking experimentation literature. Section 3 describes the basic techniques proposed for detecting the consequent performance issues. Section 4 presents the methods used to diagnose the origin of such performance issues. Finally, Section 5 provides some recommended practices to be applied when deploying a wireless networking testbed and Section 6 concludes the paper.

2. About Hardware Behavior: Sources of Potential Problems

Recent experimental research activity around wireless networking devices has revealed the presence of a number of issues that can affect the accuracy or, at least, the interpretation of the results gathered. In some cases, these findings were not predicted in advance, as they are specific to particular implementation decisions (e.g., cards noncompliant with the standards). In other cases, they correspond to new requirements that wireless devices and networks pose to computing systems (e.g., adequate control of interference in multiradio setups).

This section presents a review of the performance issues encountered in the wireless networking experimentation literature. Such issues may range from card misbehavior to unexpected interaction of the wireless hardware with the environment (e.g., interference). Furthermore, this section also identifies the specific impact that these issues may have on the expected performance of a given wireless networking deployment; see Table 1 for a summary.

Wireless networking devices present a series of performance issues that can be broadly classified into three main categories. First, wireless networking solutions present in some cases particular implementations that do not *comply with standard specifications.* This originates interoperability problems or unfairness in some scenarios.

Second, wireless networking technologies pose a series of *new requirements for computing* systems that previous communication technologies (such as the wired ones) did not. These new requirements have to be consciously tackled when moving systems into wireless networking, as they are prone to lead to misinterpretation of the results gathered.

Finally, the implementation of wireless networking solutions has not undergone the optimization effort that some wired technologies have received. As such, they are prone to present problems derived from *suboptimal or constrained developments.*

2.1. Compliance of Standard Specifications. Recent studies have shown how particular implementations of wireless technologies present particular interpretations of the standard specification. These specific implementations constitute an important source of deviations between experimental results and those predicted in analytical and simulation models.

The authors of [1] conduct the first reported measurement campaign over an operational WLAN network. One of the first issues they encounter is what they refer to as "type loss," where different WLAN implementations lost a large number of packets due to interoperability problems. Even though the authors did not find the exact origin of the problem, they managed to identify that particular implementations of the standard from different vendors were not compatible.

The authors of [2] provide a detailed study of the behavior of actual IEEE 802.11 devices and reveal how different card models present a rather different behavior with respect to some of the standardized timing values defined in the specification (e.g., EIFS). The authors extend the study in [5], where they provide a detailed analysis of the implementation of time-dependent processes (e.g., backoff) in multiple wireless device models.

The literature presents examples of how these issues affect the performance of wireless networks. As an example, the study presented in [3] reveals how the absence of "post-backoff" in specific IEEE 802.11 implementations leads to observing abnormally high throughput values when interconnecting wireless cards. Additionally, the authors of [17] show how the actual hardware behavior, when deviating from standard specifications, compromises some of the assumptions taken in theoretical and simulation-based analyses. They show, for example, how the incorrect

Table 1: Classification of the performance issues related to wireless networking: Origin and possible consequences.

Category	Source of potential problems	Observed behavior of WLAN node
Compliance of standards	"type loss" [1]	Stations from different brands present interoperability problems. Significant loss of packets.
	Post/Pre-backoff [2, 3]	Packet rates obtained might be higher than maximum ones allowed by standard.
	Standard timing implementations (slot time, backoff distribution, IFS, etc.) [2, 4, 5]	Some devices present unexpected time distributions or do not perform backoff at all. Unfairness and instability might appear.
New requirements for the system	Adjacent-channel and cochannel interference [6, 7]	Nonexpected traffic appears in measurements (close transmitters in other channels, stations in neighboring networks using same channel). Throughput may degrade.
	Power consumption [8]	Max. power allowed by low-end nodes might limit number of simultaneous wireless cards. Max. operational throughput may be lower.
	Processing power [8, 9]	Wireless cards or nodes may present lower processing power. Packet rates might be affected in highly demanding experimental settings.
	Diversity of antennas [10]	Failing to appropriately configure antenna output (with no connected antenna) leads to packet losses.
	Tolerance to delay spread [11]	When using card in scenario for which it was not designed (e.g., indoor versus outdoor), unexpected losses as well as unexpected good channels may occur.
Implementation issues (suboptimal or constrained development)	Beta drivers [12]	The use of drivers under development may affect the performance of system. Throughput and loss measurements can be distorted.
	Interrupt-handling and device memory limitations [9, 13]	Frequent interrupts or limitations in hardware memory rapidly become the bottleneck in a networking device affecting the overall performance.
	Isolation of cards [14]	Performance is affected (e.g., less available bandwidth, more collisions).
	Carrier sensing accuracy [13]	Depending on the implementation of the carrier sense mechanism, throughput and loss may degrade in multihop scenarios.
	Leakage and impedance matching of RF components [15]	Connectors leak and packets from neighboring networks might affect performance reducing throughput or incrementing loss.
	Transmission power control [16]	In case it does not work as expected, adaptation algorithms cannot be executed in order to optimize functionality.

implementation of standard time-related functions has a direct impact on the level of fairness between contending stations.

2.2. New Requirements for the Design of the System. Wireless technologies present a series of particular characteristics that have to be taken into account when developing and deploying networking infrastructures. Some of these characteristics compromise common networking practices (usually inherited from the wired networking world), which should be carefully tackled. Another source of problems may come from assuming that cards and nodes in complex wireless scenarios (e.g., multiradio) behave in the same way as in simple ones (e.g., single radio).

A clear example of these requirements is presented in [14], where the authors show how using more than one wireless card in a single computing machine is not straightforward due to leakage and coupling problems. Further, they show how carefully preparing multiradio nodes, taking into account radiation properties of devices, can effectively save misinterpretations of unexpected measurement results (e.g., reduced throughput). The study presented in [15] further reveals how specific design properties of devices, such as impedance values, are an important issue to be considered when preparing multiradio nodes. The study reveals how the performance of a multiradio node can closely approach that of multiple single-radio nodes when the interaction between all components used is clearly understood.

Wireless devices pose important power consumption requirements for computing devices. These requirements not only affect battery-powered machines but also might constitute a problem to other computing machines. The power required to operate, as an example, multiple WLAN cards at the same time has been shown to be prohibitive for certain computing machines [8]. This issue entails the risk of misinterpreting throughput or loss measurements, even more thinking that the same computing machines are able to successfully operate multiple higher throughput Ethernet devices.

Antenna diversity is another example of this category. Some wireless cards offer the option of configuring different transmission mechanisms in order to take advantage of antenna diversity. However, as shown in [10], wireless networking nodes can suffer severe performance degradation (in terms of loss) when this issue is not taken into account during the design stage.

Finally, when deploying a certain type of wireless network, one should also consider the scenario for which it was designed. As shown in [11], IEEE 802.11 WLAN cards are designed to work with values of the delay spread that are consistent with indoor transmissions. However, when using these same WLAN card models in outdoor scenarios, such delay spreads can be easily surpassed, causing unexpected loss measurements. Recent research approaches propose tuning the transmission bandwidth of wireless cards in order to circumvent such problems [18].

2.3. Implementation Constraints. The low-end profile of some wireless devices promotes the development of cost-effective solutions. However, this design strategy usually leads to generating networking solutions with specific performance limitations that have to be considered when preparing experimental setups.

An example of such a case is the one described in [9]. The authors show how some WLAN cards present different response times depending on the amount of information that they have to process. This has a direct impact on some popular bandwidth measurement tools widely used in the research community. In fact, not only the wireless cards, but also beta developments of drivers [12] or suboptimal processing of system interrupts can also lead to similar observations. Also related to the ability of a node to support a certain configuration is the access delay or speed of the hard drive in use, which may compromise data gathering capabilities, or the maximum bus transfer speed and processing power, which may be an issue for low-end devices commonly found in wireless experimental setups.

Even though WLAN specifications define certain minimum PHY layer requirements (e.g., transmission masks, maximum emission levels), some of the commercial solutions present different behaviors. A typical example of this is the protection against RF leakage of some WLAN cards [15]. RF leakage has a direct impact on the throughput and loss performance of wireless networking solutions. Another example would be the accuracy of the carrier sense mechanism. Even though the standard defines several mechanisms to determine the busy/idle state of a channel, it has been shown that adaptation strategies are optimal to avoid performance degradation in certain scenarios [19]. WLAN cards from different vendors implement particular solutions for the carrier sense, which show different performance levels [13].

Another potential source of problems is the very rough support of certain functionalities, such as modifying the transmission power control [16]. This may hinder the application of accurate mechanisms to dynamically adapt the topology or other kinds of cross-layer algorithms that need accurate power control. Furthermore, in some cases, functionalities offered by the driver are not really supported by the hardware (or vice versa), as pointed out in [16].

Finally, it should be considered that WLAN cards are still relatively recent and have not undergone the same optimization process as other technologies, such as Ethernet. As an example, state-of-the-art implementations of Ethernet devices support polling or batch interrupt options (e.g., the NAPI in Linux), which eliminate the bottleneck that represents attending frequent interrupts in a computing system. Such a solution would be highly beneficial in multiradio settings where multiple wireless cards operating at high bit rates generate frequent interrupts requesting service.

3. Basic Techniques for Detecting Hardware Misbehavior

The previous section has shown how the performance issues associated to WLAN networking devices have a direct impact on throughput or loss measurements. Indeed, in many cases the observation of an unexpected behavior in terms of losses or throughput constitutes the original indication that wireless devices are not working as expected. This section elaborates on this observation. In particular, it describes how to use workload and loss measurements in order to assess, prior to deployment, the correct performance of a wireless networking solution. In this sense, the techniques described here should be applied before the ones explained in Section 4, as the first step is to detect the existence of a problem. Then, those explained in Section 4 are used to diagnose its cause.

3.1. Detecting Hardware Misbehavior by Assessing Supported Workloads. It is common to find results in the literature that are based on characterizing the workload that a networking solution supports. This type of characterization has also been extended to wireless networking environments in order to assess the performance of WLAN solutions. Specifically, these studies consist in finding the number of packets or bits per second that the platform under analysis is able to sustain in response to an offered workload.

As described in the previous section, there are several issues that may limit the workload that a node can sustain. Among them, one can list the processing power of the computing box (or the wireless hardware itself), the maximum storage capacity of the wireless hardware, the interrupt handling process in the intercommunication between the wireless device and the board, the degree of standard and protocol compliance of the wireless device, and the (in)correct use or calibration of the wireless node. Additionally, when multiple wireless interfaces are considered, interference among them or power consumption requirements might also be limiting factors.

It is worth mentioning here that the characterization of supported workloads must be done in the absence of channel propagation errors, as the focus is on assessing the performance of the hardware itself regardless of environment conditions. This is typically achieved by using coaxial cables to interconnect wireless devices or by establishing close Line-of-Sight (LoS) communications in isolated environments (i.e., without external interference). Section 5.3 provides further discussion about this.

The characterization of supported workloads might take many forms, depending on the target scenario and planned functionality of the wireless node under test (NUT). Relevant to wireless networking nodes, researchers have analyzed (1) *traffic generation and reception rates* of a node, (2) its *forwarding rates,* and (3) the *concurrent rates* that it is able to sustain.

Figure 1 summarizes the experimental setup used to obtain the supported load characterization of a node (NUT in the figure). The figure represents the three types of characterization methods described.

The *traffic generation and reception rates* of a node measure the maximum rate at which a node can generate or receive traffic. This characterization technique has been used, for example, in [3] and typically consists in setting constant bit rate flows of an increasing intensity between a controlled testing node and the node under test (NUT). The test proceeds until measuring the maximum load that the NUT is able to sustain without losses. This has been used in [3] to characterize traffic generation between pairs of cards from different vendors by plotting packet_rate_generated versus packet_rate_requested curves. In this particular case, the setup used was the one tagged as (1.a) in Figure 1. The shape of the obtained curves is a slope 1 line that becomes flat (i.e., slope = 0) as the maximum rate is reached. As explained in [3], when one compares the maximum rates obtained for each pair of cards with the theoretical one, the conclusion is that some pairs surpass this limit, and thus, there may be standard compliance issues. The different values obtained depend on usage of short/long preamble and on whether pre-backoff is used or not. To further study this behavior, some techniques are presented in Section 4 for really diagnosing the cause of this noncompliance.

Reference [3] also used traffic reception tests for characterizing the performance of multiple sniffers installed in the same machine. In this case, the setup used was (1.b) in Figure 1. When compared to the curve obtained when a single sniffer is installed, one may observe a degradation of the capture rate, starting at a packet rate of 1000 packets/s, when multiple Prism cards are simultaneously sniffing the same flow. This may allow detecting problems in the isolation of cards, which may be appropriately diagnosed by means of techniques explained in the next section.

Measuring the *forwarding rates* of a node is a classic benchmarking technique, commonly applied to switching and routing nodes [20]. It accounts for the number of frames per second that a NUT can send to an intended destination in response to an offered load. This characterization technique is especially useful in those nodes that are to be used in the backhaul of a wireless mesh network, for example, [21]. Setup 2 of Figure 1 would be used in this case, but the same kind of curves as those obtained in traffic generation tests may be used to detect problems in the forwarding node.

Finally, when measuring the *concurrent rates* that a node is able to sustain, a researcher is determining the traffic load that a station is able to reliably handle when it is both dealing with incoming and outgoing traffic. Reliable here means that no packet drops are observed in neither the incoming nor the outgoing traffic of the measured wireless node. This technique can be thought of as joining together the previous two. Therefore, traffic is generated in the same way as in the previous cases. This characterization technique is described in [13] and the setup used is 3 in Figure 1. The reader should notice that flow 2 is not the result of forwarding flow 1, and what is represented is receiving_rate_of_NUT versus sending_rate_of_NUT. In some cases, the NUT may generate more traffic than it should if busyness of medium was perfectly assessed. This allows detecting problems in carrier sensing accuracy [13].

Several experimental research studies give practical evidence of the convenience of using supported load characterization (e.g., [6, 14, 15]) to detect unexpected issues. In such studies the authors generally compare the measured values in a supported load characterization with an expected (theoretical) performance.

There exist numerous studies in the literature that provide reference numbers for the throughput that can be achieved. These analyses are applied both in regular reference conditions (e.g., [22]) and in optimal (zero propagation loss) conditions (e.g., [23, 24]).

3.2. Detecting Hardware Misbehavior by Assessing Loss Measurement Performance.

The measurement of loss is popular in experimental wireless networking research. It is a valuable tool for the design of efficient solutions to combat the high variability of wireless channel conditions. Further, loss measurements allow the design of efficient network deployment and management strategies.

However, as shown in Section 2, care should be taken when gathering loss measurements, as wireless devices may present some performance issues affecting this type of measurements. Examples of this are the processing power of wireless devices, the proper use of RF components, memory limitations of the elements used (WLAN card or computing machine), or the compliance of standard specs and its effects on fairness.

Typically, assessing the performance of a solution to support loss measurements does not involve bringing the

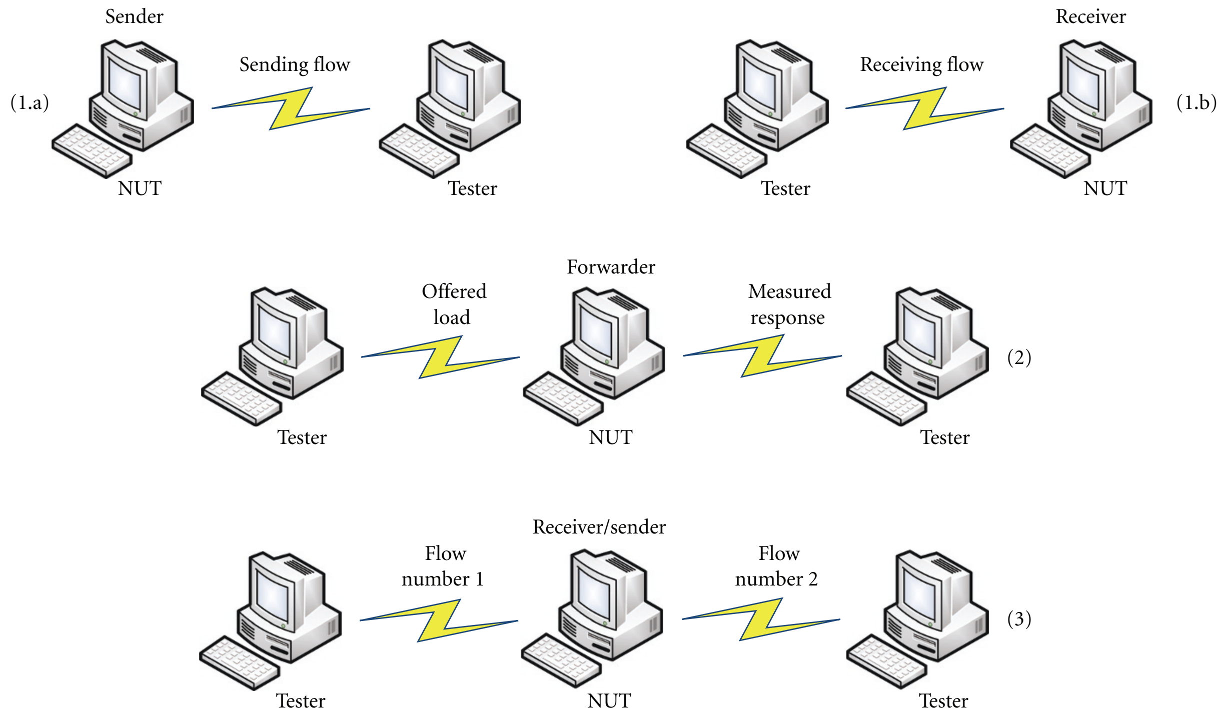

Figure 1: Experimental setup used to assess the *traffic generation* (1.a) *and reception* (1.b) *rates,* the *forwarding rates* (2), and the *concurrence rates* (3) of a wireless node under test (NUT).

wireless node to its performance limits. In such a case, unpredicted losses would lead to unexpected measurements of supported workloads.

Rather, assessing the ability of a node to produce accurate loss measurements typically requires running the node in favorable conditions. From a physical layer perspective, this leads to interconnecting wireless components with coaxial cables or establishing close LoS communications. The objective of this is preventing propagation losses in order to identify other possible sources of loss in the wireless solutions used. Alternatively, some approaches [11] use channel emulators to evaluate the behavior of wireless solutions in typical reference channel conditions. This helps them evaluate whether the behavior of wireless devices is the one expected from design specifications. This second approach, however, is more expensive.

There are a number of features that are common in loss measurements gathered in WLAN network deployments (e.g., [1, 11]). In general, they are long experiments and they rely on broadcast packets of different sizes sent at different rates. Such measurements can be repeated in a controlled manner in order to assure that measurement nodes have sufficient memory capabilities and, under optimal channel propagation conditions, they present no losses. The presence of unexpected losses is an indicator of the presence of specific performance issues that may affect the correct functionality of a given network deployment.

4. Approaches Adopted to Diagnose the Origin of Hardware Misbehavior

Previous sections have presented typical sources of unexpected behavior in wireless networking experimentation (Section 2) and common tests to detect potential anomalies in the equipment used (Section 3). This section aims at closing the loop and presenting the methodologies that researchers use to determine the origin of unexpected behavior. In other words, this section presents the methods that researchers have been using to diagnose what the specific problem associated to unexpected performance observations is.

Furthermore, these techniques are in general found in the context of measurement studies that have a certain goal, hence determining the exact measurement of interest, for example, the way in which workload is increased, or the parameter to be represented in the X-axis of the curve. In this sense, it is difficult to highlight one single curve for each technique that is able to solve all issues in all contexts. Therefore, the focus is on describing the technique and its interest, and also on giving some hints on its potential application as well as some representative proposed plots.

Table 2 lists these techniques, grouping them into (1) *hardware design issues*, (2) *wireless communications and protocols*, and (3) *environment awareness*. Additionally, the table shows the equipment required, if any, in order to successfully apply each one of the techniques and the source/s of potential problems (as introduced in Section 2) each technique tackles.

4.1. Characterizing Issues Related to Hardware Design. Hardware design options deeply shape the performance of networking devices. The complexity and efficiency of the components and the way they interact influence the overall functionality of the networking node. As previously mentioned, wireless networking devices did not yet undergo all performance optimizations that wired nodes did. Current networking architectures offer a wide range of features that are not currently exploited by wireless networking commercial products.

Besides, while wireless devices follow this optimization process, one can find some characterization proposals that have been designed and used to extract information about specific hardware-related issues. These techniques are not exclusive to wireless networking devices, but they are particularly relevant and have been successfully used for WLAN solutions.

The authors of [9] propose a methodology that aims at characterizing the time it takes for a wireless node (which is carrying a particular wireless device) to generate a response to a packet probing request. This is done in the context of bandwidth measurements, but it has general applicability when the *response time* of wireless networking devices must be characterized. For instance, such parameter may have a noticeable influence in some bandwidth measurement techniques. The methodology proposed there, but of general application to other contexts, consists in using a traffic generator (D-ITG) that can signal the instant of a packet generation through the serial port. This signal is used both at the sender and receiver to trigger energy level measurements in an oscilloscope that serve to obtain exact timestamps of packet generation instants. The *response time* is affected by several factors concerning hardware (and software) design options and it may deeply impact the performance of several packet probing-based algorithms. In particular case, the curve obtained was response time versus packet size, but there may be another X-axis of interest depending on the scenario under evaluation.

An important issue differentiating wireless networking devices and their wired counterparts is *power consumption*. This is especially relevant when standalone low-power machines are planned to be used in wireless deployments. The study in [8] presents a methodology to characterize the power that multiple wireless interfaces demand to the power supply. In this case, a controllable power supply is used as power source for the board and the methodology consists in progressively adding wireless devices and progressively stressing the workload handled by the wireless networking node and measuring the current consumed at any time. In this case, a relevant figure is power consumption versus number of wireless cards. This kind of measurements is important because circuitry designed to support wired communications may not be correctly dimensioned to support the power consumption requirements of wireless networking communications and may have insufficient resources to support multiple wireless networking devices.

4.2. Characterizing Issues Related to Wireless Communications and Protocols. The characterization techniques presented above are not specifically designed for wireless networking nodes, but might be applied, in general, to any networking node. However, wireless networking presents a series of particularities that require specific characterization techniques specifically designed for it. This document categorizes these techniques as *wireless communications and protocols* and the more specific category *environment awareness*.

The category *wireless communications and protocols* encloses those characterization techniques that determine the ability of a node to communicate using the procedures of the protocol it claims to be using. On the other side, the category *environment awareness* refers to those techniques that characterize the ability of the node to be aware of the environment it is communicating in and to adapt to it. In a sense, the former refers to characterizing the ability of the node to actively use the environment to communicate and the latter characterizes the ability of the node to detect and be protected against issues related to wireless transmissions. This section discusses techniques within the *wireless communications and protocols* category. The next section deals with the *environment awareness* category.

The first draft proposals of the recommended practice 802.11T [25] include the performance metric *throughput versus coverage/range and multipath*. This characterization technique is similar to that of *traffic generation and reception rates* described above but including, in the measurement, effects of channel propagation. Among the several factors shaping the result of this measurement, one can identify the signal distortion induced by hardware and wireless node case shape, the transmission power and antenna gain, correct impedance matching between the components used, and the mutual interference between the several interfaces that a wireless node might be carrying. A reduced version of this technique is used in [11]. A potential curve of interest in this case is bit-rate received versus distance (or versus total path loss) when allowing automatic rate adaptation.

Recently, a study by Ben Abdesslem et al. [16] proposed the characterization of the *power control support* that commercial wireless devices offer. The characterization technique consists in placing a measurement node close to the wireless node that is going to be characterized (NUT). Then, the NUT starts a ping with the measurement node. The ping process forms, in this case, a constant flow that is used, at the measurement node, to evaluate the received signal strength level. From here on, the measurement methodology consists in progressively changing (with a certain granularity) the output power level of the NUT (from 1 mW up to the maximum allowed one) and recording the received signal strength at the measurement node. The result reveals whether the NUT supports power control options and with which granularity. Thus, the curve of interest is in this case received signal strength versus output power level at NUT.

The authors in [5] present a technique to characterize the correctness of a wireless device in following interpacket generation times specified by the standard. Constant bit rate flows with various packet sizes were tested. Using a precise timestamping (down to the microsecond) to measure

Table 2: Survey of characterization techniques used in the literature to diagnose the origin in the presence of unexpected performance of WLAN devices.

Category	Characterization techniques	Equipment required	Source of potential problems
Hardware design issues	Response time [9]	Oscilloscope	Beta drivers, processing power
	Power consumption [8]	Controllable Power supply	Power demand
Wireless communications and protocols	Throughput versus coverage/range and multipath [11]	Channel emulator	Standard timing compliance, diversity of antennas, transmission power control
	Transmission power [16]	Power meter	Transmission power control
	Packet interarrival [5]	None	Post/pre-backoff, standard timing compliance
	Interoperability tests [1]	None	Type loss
Environment awareness	Concurrent rates [13]	None	Carrier sensing accuracy
	Adjacent and cochannel interference [6] (alt. [7])	Spectrum analyzer (alt. none)	Adjacent and cochannel interference
	Shielding and antenna separation [13, 14]	None	Isolation of cards

the packet intergeneration time, the authors are able to characterize the random backoff process of wireless cards and the CWmin they use. The methodology determines the time elapsed since a packet is ready to be transmitted in a computing machine until it is actually sent under various conditions. The authors in [4] use a similar technique to obtain the packet interarrival time. They use this technique to characterize several commercial wireless nodes, how they conform to the Standard, and how they interact. The cumulative probability distribution for each measured interpacket spacing is the curve of reference in this case.

A recent study presented in [10] introduces a methodology to understand the specific procedure that a WLAN card follows when it has more than one antenna available for transmission. It differentiates between unicast and broadcast frame transmission, as different transmit diversity algorithms may be implemented depending on the presence or not of ACK frames. In unicast transmissions, they may help in determining how good a certain channel is. Again, constant bit rate flows are generated with a large packet size. For broadcast transmissions the curve of interest is SNR versus time, with an appropriate averaging window, which is set to 40 ms in [10]. Periodic variations of SNR at regular intervals indicate the presence of antenna diversity algorithms. To further confirm this throughput versus time plots with and without antenna diversity enabled may be of use. As for unicast frames, the plots of reference are the SNR and retry distributions. Non-Gaussian SNR distributions and nonmonotonically decreasing retry distributions may indicate antenna diversity issues. Again, throughput versus time plots with and without diversity techniques enabled would definitely confirm this. Overall, the technique is able to identify whether the card periodically switches the transmission antenna (without assessment) or whether it follows a specific pattern (e.g., switch when there are two consecutive unsuccessful retransmissions).

Finally, one can find other characterization efforts in the literature related to characterizing specific interoperability issues between nodes. Firstly, a pioneering study by Yeo et al. [1] on wireless network measurements captures traces of the communications between WLAN devices in order to assess interoperability issues between them. Interoperability problems may be present when there are substantial packet losses observed in the communication between two nodes in a controlled environment in terms of channel propagation. Secondly, an initiative from the University of New Hampshire [26] offers the possibility to vendors to conduct exhaustive conformance tests to validate and certify 802.11 products.

4.3. Characterizing the Environment Awareness of a Node. This category includes all those characterization techniques that aim at determining the interaction of the node being characterized and the environment in which it is placed in terms of wireless transmissions. This category gains particular importance in wireless mesh networking environments, where one should characterize the effects of the interactions between multiple radios inside a single box, the possibility to use several channels at the same time, or the effectiveness and conditions of using several antennas.

The authors of [6] propose a methodology to characterize the interference that a node generates in adjacent channels. Using a spectrum analyzer, they measure the power spectral density of a wireless node as seen in close proximity. The current transmission mask specified in 802.11 standard does not allow in practice coexistence of transmitters and receivers simultaneously operating in adjacent channels in a single computer, as adjacent channel interference may be substantial. However, this technique may be used to determine which wireless nodes are more suitable to build up a multiradio node, the distance required between devices and/or antennas, and even the most suitable technology to build up multiradio mesh testbeds.

Assuring isolation from external interferences or harmful leakages has been the object of several calibration processes in wireless experimental research studies. Traditionally, sniffing out zero channel activity with a monitoring node during a certain time is considered acceptable to consider that an experiment is free from external interference. However, nowadays, finding spots with zero channel activity, especially when using Wi-Fi nodes, is becoming harder and harder. Several methods have been proposed in the literature to shield and isolate wireless networking nodes from the influence of external interferences or from issues such as adjacent and cochannel interferences. The authors of [14] propose the use of throughput measurements to characterize the effectiveness of a shielding solution. The most representative figure to characterize this is throughput versus numerous channels of separation when constant bit rate UDP flows with different packet sizes are used. Additionally, the authors of [6] use a spectrum analyzer to characterize the shape of the output signal of a wireless node. This same technique can be used to estimate the effectiveness of a shielding solution.

The methodology proposed in [13] aims at characterizing the accuracy of a wireless node to detect transmissions from other nodes in multihop environments. The methodology obtains a characterization of the accuracy of the node when gathering low-level energy measurements and its potential impact on sustaining wireless communications. A new metric is proposed, called α, that accounts for the probability that a node is sensing the medium as idle when it is actually being used. The metric is computed by measuring the difference between the expected ideal protocol behavior and the actual behavior of a wireless node. As explained in Section 3.1, setup 3 in Figure 1 is used and constant bit rate traffic is sent for both flows. The reference curve in this case is receiving rate of NUT versus sending rate of NUT. If the curve obtained is within the operational region defined, it indicates that the NUT is constrained by the (in)accuracy of medium sensing. On the other hand, if it falls outside this region, it indicates that the node is constrained by the inability of handling incoming and outgoing traffic at the same time at the configured rate.

5. Some Recommended Practices

There follows a list of recommended practices, regarding the characterization of nodes, to be applied when deploying a wireless networking testbed. The list, as a whole, constitutes a methodology to characterize the wireless hardware to be used to set up a wireless networking testbed. These recommended practices build on the discussion followed throughout the paper and enclose most of the ideas explained and some of the concluding remarks found in the papers referenced.

5.1. Always Consider a Characterization Step in Your Testbed Design and Implementation Process. Characterization of wireless networking nodes has not been much documented in the recent literature on experimentation. However, as shown throughout this paper, planning a conscious characterization of the wireless networking nodes that are going to be used can be very helpful. This can avoid taking nonappropriate assumptions that can affect the validity of the results gathered. An experimenter must be aware of which are the features of the underlying hardware that are of essential importance for the performance of the algorithms and protocols to be tested and carry out characterization tests that help him understand whether the hardware solution adopted is the appropriate one. This is referred to as *selection of relevant parameters* in Figure 2.

Once parameters of interest are selected, the characterization of wireless nodes can be divided into two main steps. In the first step, the researcher uses the characterization techniques presented in Section 3 to detect possible misbehaviors of the wireless nodes to be used. In the second step, and in case unexpected results arise during the first one, the researcher uses specific techniques (among those presented in Section 4) to diagnose the origin of these previously unforeseen observations. Once the origin of the problem is diagnosed, appropriate corrective actions in terms of *new design choices* may be taken. The new *node prototype* is built and a new *characterization* process is carried out, and so on, until no problem is detected, which gives green light to the eventual *deployment of the testbed*. See Figure 2 for an overview of the process.

5.2. Use a Supported Load Characterization to Assess Whether the Node Prototype Follows the Expected Behavior. As shown in Section 3.1, supported load characterizations are a simple and useful means to detect unexpected behaviors of the wireless networking nodes that are being characterized. They have been repeatedly used in the literature [6, 14, 15] to detect unexpected interactivity issues between components, nonstandard compliance, and so forth.

However, though supported loads are simple to set up, the researcher might also be cautious in choosing and tuning the tools used (e.g., the software tool used to generate traffic) in order to prevent gathering invalid observations.

5.3. In Case Hardware Is not Behaving as Expected, Use Specific Characterizations to Diagnose the Origin of Misbehaviors. There are two main ways to diagnose the origin of unexpected hardware behavior. On one side, the researcher can use a specific characterization technique designed to tackle a particular aspect of the wireless node. Examples of this are the methodology used in [6] to analyze the level of adjacent channel interference using a spectral analyzer or the methodology described in [9] to obtain a measurement of the response time of a wireless node.

On the other side, the researcher can also make use of supported load characterizations to diagnose the origin of misbehaviors. This is possible when the researcher has the possibility of isolating the impact of each one of the components during the supported load characterization. As an example, the authors in [15] characterize the interference between wireless devices in a multiradio node using supported load characterizations. The study shows and

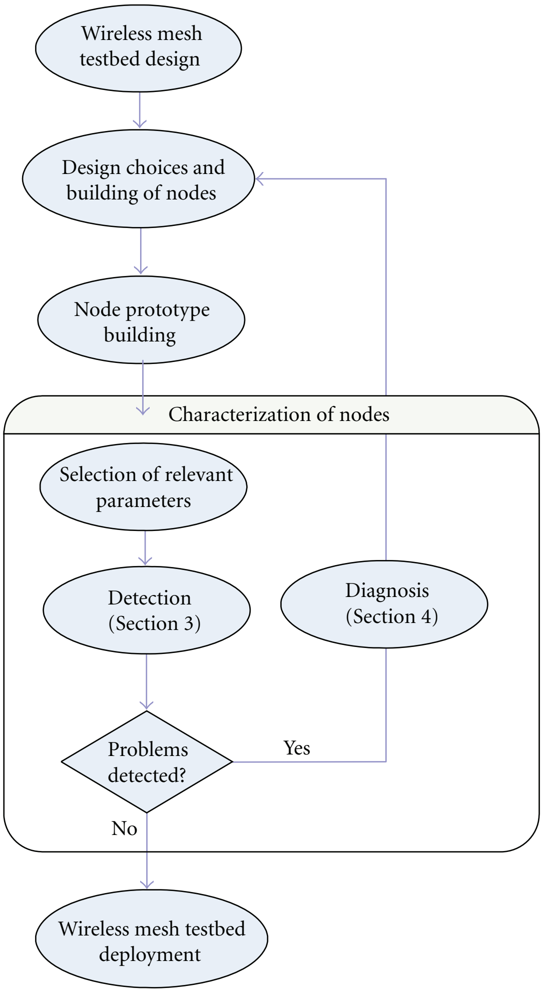
Wireless mesh testbed design
Design choices and building of nodes
Node prototype building
Characterization of nodes
Selection of relevant parameters
Detection (Section 3)
Diagnosis (Section 4)
Problems detected?
Yes
No
Wireless mesh testbed deployment

Figure 2: Consider a characterization step in your testbed design and implementation process.

characterizes how the attenuation between antennas and the correct match adaptation of components is behind misbehaviors of multiradio settings.

There follows a list of hints that can help accelerating the diagnosis process and avoid undergoing an extensive characterization of the node to find specific performance issues.

(1) Whenever Possible, Start First by Characterizing and Deciding on the Wireless Hardware to Use. While building a wireless node, one can use either compact commercial solutions or assemble different components (typically a computing form factor and one or more wireless devices). When this second solution is adopted, it is interesting to obtain a separate characterization of the wireless hardware prior to characterizing the wireless node as a whole.

Recent research has shown (e.g., [1]) that not all wireless hardware solutions behave comparably, regardless of the board functionality, and it is an interesting practice to characterize the devices to be used prior to building the wireless node itself. This serves also to identify or discard possible wireless hardware related bottlenecks when characterizing the performance of the whole node.

Separate characterization of the wireless hardware can be done by using a low entry server as the form factor support in order to assure that processing power of the board is not a limiting factor.

(2) Devote Special Attention on Tuning the Operational Characteristics of Devices and Components. Conscious tuning of the devices used should be always conducted, even more during the characterization process. Attention should be paid to the following issues. (1) Choose appropriately adapted RF components to your working frequency. In general, the higher the working frequency the more cautionary one should be. (2) Devise an appropriate data gathering methodology. Issues such as hard drive writing access delay or speed might compromise data collection. (3) The maximum bus transfer speed and processor power should always be taken into account when evaluating results. (4) Some drivers offer configuration options that are not really implemented in the wireless hardware that they control. Double check this before starting measurements.

(3) Whenever Possible, Use Cables, Otherwise, Avoid Interferences and Gather Measurements Using Close LoS Communications. Performance characterization of the hardware used should be done in the absence of complex propagation losses. The objective here is characterizing the hardware used regardless of the environment.

RF cables are a very convenient solution to conduct measurements. Devices are both isolated from external interferences and safe from unintended propagation losses. However, not all wireless hardware solutions offer the possibility to connect cables or external antennas. In these cases, wireless nodes that establish a communication during any test run should be placed close to each other and within Line of Sight in order to minimize unintended propagation effects. Experimental runs show that either using cables or close LoS communications (in the absence of external interferences) show close results even in indoor environments [15].

(4) Bear in Mind That Single-Radio Performance Upper-Bounds Multiradio Performance. As was shown in [15], the performance of a single-radio setting upper-bounds the performance of the same device in a multiradio setting. A prior characterization of the single-radio node performance can be used as a reference when calibrating multiradio nodes. The closer the performance of the multiradio node is to the multiple single-radio case, the better calibrated your setting is.

5.4. Once a Hardware Misbehavior Is Detected and Diagnosed, Either Reconsider the Network Design or Take It into Account in the Results. Whenever unexpected issues arise, they might lead to reconsidering the hardware to be used. As an example, when two wireless cards turn out not to correctly interoperate, a hardware device change must be considered.

However, the process of characterizing unexpected behaviors of the wireless node might help understanding the experimental results and correctly assess the tolerance to be accepted in the results.

6. Conclusions

From a general point of view, the novelty of this paper resides on the fact that it brings attention to the potential dangers of drawing conclusions without appropriately considering the behavior of the specific hardware in use in a wireless networking testbed. There is a risk that a researcher makes a series of assumptions about the hardware/software tools used that may not be sufficiently accurate.

In order to support this observation, the paper surveys recent literature and identifies a number of performance issues that can constitute sources of problems during experimentation. Indeed, a substantial part of the referenced literature is motivated by unexpected observations when operating WLAN testbeds.

The paper also highlights the fact that unexpected observations are generally related to throughput and loss measurements. This suggests the idea of using a priori workload and loss tests in controlled environments to assess the performance of the solutions to be used. This will help detecting in advance any potential issue that may arise during the operation of an experimental deployment.

Once the presence of an unexpected performance issue is detected, the literature offers numerous techniques to diagnose the origin of the problem; they are surveyed and categorized in this paper. Despite requiring in some cases of complex setups, these techniques will allow determining the need for calibrating or changing a given WLAN solution or, at least, using it but handling the results with caution.

Finally, some illustrative examples are provided to exemplify how a researcher should make use of characterization on her/his own benefit.

Acknowledgment

This work was supported in part by the Spanish Ministry of Science and Innovation under Grant number TEC2008-06826/TEC (project ARTICO).

References

[1] J. Yeo, M. Youssef, and A. Agrawala, "A framework for wireless LAN monitoring and its applications," in *Proceedings of the Workshop of Wireless Security (WiSe '04)*, 2004.

[2] A. Di Stefano, G. Terrazzino, L. Scalia, I. Tinnirello, G. Bianchi, and C. Giaconia, "An experimental testbed and methodology for characterizing IEEE 802.11 network cards," in *Proceedings of the IEEE International Symposium on a World of Wireless Mobile and Multimedia Networks (WoWMoM '06)*, Niagara Falls, Canada, June 2006.

[3] M. Portoles-Comeras, M. Requena-Esteso, J. Mangues-Bafalluy, and M. Cardenete-Suriol, "Monitoring wireless networks: performance assessment of sniffer architectures," in *Proceedings of the IEEE International Conference on Communications (ICC '06)*, Istanbul, Turkey, June 2006.

[4] G. Berger-Sabbatel, Y. Grunenberger, M. Heusse, F. Rousseau, and A. Duda, "Interarrival Histograms: A Method for Measuring Transmission Delays in 802.11 WLANs," Research report, LIG lab, Grenoble, France, October 2007.

[5] G. Bianchi, A. D. Stefano, C. Giaconia, L. Scalia, G. Terrazzino, and I. Tinnirello, "Experimental assessment of the backoff behavior of commercial IEEE 802.11b network cards," in *Proceedings of the Annual Joint Conference of the IEEE Computer and Communications Societies (INFOCOM '07)*, Anchorage, Alaska, USA, May 2007.

[6] C.-M. Cheng, et al., "Adjacent channel interference in dual-radio 802.11a nodes and its impact on multi-hop networking," in *Proceedings of IEEE Global Telecommunications Conference (GLOBECOM '06)*, 2006.

[7] P. Fuxjager, D. Valerio, and F. Ricciato, "The myth of non-overlapping channels: interference measurements in IEEE 802.11," in *Proceedings of the 4th Annual Conference on Wireless on Demand Network Systems and Services (WONS '07)*, Obergurgl, Austria, 2007.

[8] M. Portoles-Comeras, M. Requena-Esteso, and J. Mangues-Bafalluy, "Framework for characterizing hardware deployed in wireless mesh networking testbeds," in *Proceddings of*

the International Conference on Testbeds and Research Infrastructures for the Development of Networks and Communities (TridentCom '07), Orlando, Fla, USA, May 2007.

[9] L. Angrisani, A. Botta, A. Pescape, and M. Vadursi, "Measuring wireless links capacity," in *Proceedings of the 1st International Symposium on Wireless Pervasive Computing*, January 2006.

[10] D. Giustiniano, G. Bianchi, L. Scalia, and I. Tinnirello, "An explanation for unexpected 802.11 outdoor link-level measurement results," in *Proceedings of the Annual Joint Conference of the IEEE Computer and Communications Societies (INFOCOM '08)*, Phoenix, Ariz, USA, 2008.

[11] D. Aguayo, J. Bicket, S. Biswas, G. Judd, and R. Morris, "Link-level measurements from an 802.11b mesh network," in *Proceedings of the ACM Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications (SIGCOMM '04)*, 2004.

[12] The MadWifi project, http://madwifi-project.org.

[13] M. Portoles-Comeras, A. Krendzel, and J. Mangues-Bafalluy, "Methodology to characterize the performance of IEEE 802.11 nodes to be deployed in multi-hop environments," in *Proceedings of the 3rd International Workshop on Wireless Network Measurement (WiNMee '07)*, 2007.

[14] J. Robinson, K. Papagiannaki, C. Diot, X. Guo, and L. Krishnamurthy, "Experimenting with a multi-radio mesh networking testbed," in *Proceedings of the 1st International Workshop on Wireless Network Measurement (WiNMee '05)*, April 2005.

[15] M. Portoles-Comeras, J. Mangues-Bafalluy, and M. Requena-Esteso, "Multi-radio based active and passive wireless network measurements," in *Proceedings of the 2nd International Workshop on Wireless Network Measurement (WiNMee '06)*, Boston, Mass, USA, April 2006.

[16] F. Ben Abdesslem, L. Iannone, M. D. de Amorim, K. Kabassanov, and S. Fdida, "On the feasibility of power control in current IEEE 802.11 devices," in *Proceedings of the IEEE Percom Workshop on Pervasive Wireless Networking (PWN '06)*, Pisa, Italy, March 2006.

[17] D. Malone, I. Dangerfield, and D. Leith, "Verification of common 802.11 MAC model assumptions," in *Proceedings of the Passive and Active Measurement Conference (PAM '07)*, April 2007.

[18] R. Chandra, R. Mahajan, T. Moscibroda, R. Raghavendra, and P. Bahl, "A case for adapting channel width in wireless networks," in *Proceedings of the ACM Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications (SIGCOMM '08)*, 2008.

[19] J. Zhu, X. Guo, S. Roy, and K. Papagiannaki, "CSMA self-adaptation based on interference differentiation," in *Proceedings of IEEE Global Telecommunications Conference (GLOBECOM '07)*, Washington, DC, USA, November 2007.

[20] R. Mandeville, "Benchmarking terminology for LAN switching devices," IETF RFC 2285, February 1998.

[21] R. Draves, J. Padhye, and B. Zill, "Routing in multi-radio, multi-hop wireless mesh networks," in *Proceedings of the Annual International Conference on Mobile Computing and Networking (MobiCom '04)*, Philadelphia, Pa, USA, 2004.

[22] J. del Prado Pavon and S. Choi, "Link adaptation strategy for IEEE 802.11 WLAN via received signal strength measurement," in *Proceedings of the IEEE International Conference on Communications (ICC '03)*, Anchorage, Alaska, USA, May 2003.

[23] Y. Xiao and J. Rosdahl, "Performance analysis and enhancement for the current and future IEEE 802.11 MAC protocols," *ACM SIGMOBILE Mobile Computing and Communications Review*, vol. 7, no. 2, pp. 6–19, 2003.

[24] G. Bianchi, "Performance analysis of the IEEE 802.11 distributed coordination function," *IEEE Journal on Selected Areas in Communications*, vol. 18, no. 3, pp. 535–547, 2000.

[25] The IEEE 802.11T Task Group, http://grouper.ieee.org/groups/802/11/Reports/tgt_update.htm.

[26] University of New Hampshire, "InterOperability Laboratory," http://www.iol.unh.edu/services/testing.