Troubleshooting Multihop Wireless Networks

Lili Qiu, Paramvir Bahl, Ananth Rao, and Lidong Zhou

Updated November 2004

Technical Report
MSR-TR-2004-11

Effective network troubleshooting is critical for maintaining efficient and reliable network operation. Troubleshooting is especially challenging in multi-hop wireless networks because the behavior of such networks depends on complicated interactions between many unpredictable factors such as RF noise, signal propagation, node interference, and traffic flows. In this paper we propose a new direction for research on fault diagnosis in wireless networks. Specifically, we present a diagnostic system that employs on-line trace-driven simulations to detect faults and perform root cause analysis. We apply this approach to diagnose performance problems caused by packet dropping, link congestion, external noise, and MAC misbehavior. In a 25 node multihop wireless network, we are able to diagnose over 10 simultaneous faults of multiple types with more than 80% coverage. Our framework is general enough for a wide variety of wireless and wired networks.

Microsoft Research
Microsoft Corporation
One Microsoft Way
Redmond, WA 98052
http://www.research.microsoft.com

1. INTRODUCTION

Network management in multihop wireless networks is a necessary ingredient for providing high-quality, reliable communications among the networked nodes. Unfortunately, it has received little attention until now. In this paper, we focus on network troubleshooting, the component of network management responsible for maintaining the "health" of the network and ensuring its smooth and continued operation [16].

Troubleshooting a network, may it be wired or wireless, is a difficult problem. This is because of complex interactions between many different network entities and between faults that occur in the different parts of the network. Troubleshooting a multihop wireless network is even more difficult because:

- In wireless networks, signal propagation is affected by fluctuating environmental conditions. Signal variations make network links unpredictable and unreliable causing the network topology to change rapidly and frequently. These changes impact protocol and application behavior.
- Multihop wireless networks have limited capacity. Scarcity of resources such as bandwidth and energy puts tight constraints on the amount of management traffic the network can tolerate. The tradeoff between performance improvement because of management and performance degradation because of control overhead requires careful attention.

To address these challenges, we propose a novel troubleshooting framework that integrates a network simulator into the management system for detecting and diagnosing faults occurring in an operational network. We collect traces; we clean them; and then we use them to recreate in the simulator the events that took place inside the real network.

For our system to work, we must solve two problems: (i) accurately reproduce inside the simulator what just happened in the operational network; and (ii) use the simulator to perform fault detection and diagnoses.

We address the first problem by taking an existing network simulator (e.g., Qualnet [37], a commercially available packet-level network simulator) and identify the traces to drive it with. (Note: although we use Qualnet in our study, our technique is equally applicable to other network simulators, such as ns-2 [30], OPNET [32] etc.). We concentrate on physical and link layer traces, including received signal strength, and packet transmission and retransmission counts. We replace the lower two networking layers in the simulator with these traces to remove the dependency on generic theoretical models that do not capture the nuances of the hardware, software, and radio frequency (RF) environment.

We address the second problem with a new fault diagnosis scheme that works as follows: the performance data emitted by the trace-driven simulator is considered to be the expected baseline ("normal") behavior of the network and any significant deviation indicates a potential fault. When a network problem is reported/suspected, we selectively inject a set of possible faults into the simulator and observe their effect. The fault diagnosis problem is therefore reduced to efficiently searching for the set of faults which, when injected into the simulator, produce network performance that matches the observed performance. This approach is significantly different from the traditional signature based fault detection schemes.

Our system has the following three benefits. First, it is flexible. Since the simulator is customizable, we can apply our fault detection and diagnosis methodology to a large class of networks operating under different environments. Second, it is robust. We are able to capture complicated interactions within the network and between the network and the environment, as well as among the different faults. This allows us to systematically diagnose a wide range and combination of faults. Third, it is extensible. New faults are handled independently of the other faults as the interaction between the faults is captured implicitly by the simulator.

We have successfully applied our system to detect and diagnose performance problems that arise from the following four faults:

- Packet dropping. This may be intentional or may occur because of hardware and/or software failure in the networked nodes. We care about persistent packet dropping.
- Link congestion. If the performance degradation is because of too much traffic on the link, we want to be able to identify this.
- External noise sources. RF devices may disrupt on-going network communications. We concern ourselves with noise sources that cause sustained and/or frequent performance degradation.
- MAC misbehavior. This may occur because of hardware or firmware bugs in the network adapter. Alternatively, it may be due to malicious behavior where a node deliberately tries to use more than its share of the wireless medium.

These faults are more difficult to detect than fail-stop errors (e.g., a node turns itself off due to power or battery outage), and they have relatively long lasting impact on performance, In this paper, we focus only on identifying the faults, and not on the corrective actions one might take.

We demonstrate our systems ability to detect random packet dropping and link congestion in a small multihop IEEE 802.11a network. We demonstrate detection of external noise and MAC misbehavior via simulations because injecting these faults into the testbed in a controllable manner is difficult. In a 25 node multihop network, we find that our troubleshooting system can diagnose over 10 simultaneous faults of multiple types with more than 80% coverage and very few false positives.

To summarize, the primary contribution of our paper is to show that a trace-driven simulator can be used as a real-time analytical tool in a network management system for detecting, isolating, and diagnosing faults in an operational multihop wireless network. To the best of our knowledge, we are the first to propose and evaluate such a system. In the context of this system, we make the following three contributions:

- We identify traces that allow a simulator to mimic the multihop wireless network being diagnosed.
- We present a generic technique to eliminate erroneous trace data.
- We describe an efficient search algorithm and demonstrate its effectiveness in diagnosing multiple network faults.

The rest of this paper is organized as follows. We describe the motivation for this research and give a high-level description of our system in Section 2. We discuss system design rationale in Section 3. We show the feasibility of using a simulator as a real-time diagnostic tool in Section 4. In Section 5, we present fault diagnosis. In Section 6, we describe the prototype of our network monitoring and management system. We evaluate the overhead and effectiveness of our approach in Section 7, and discuss its limitations and future research challenges in Section 8. We survey related work in Section 9, and conclude in Section 10.

2. SYSTEM OVERVIEW

There is widespread grassroots interest in community and rural-area wireless mesh networks [3, 17]. Mesh networks enable applications like Internet gateway sharing [11, 2], local content sharing, gaming etc. They grow organically as users buy and install equipment [38], but they often lack centralized network management. Therefore, self-management and self-healing capabilities as envisioned in [16], are key to the long-term survival of these networks. It is this vision that inspires us to research network troubleshooting in multihop wireless networks.

Our management system consists of two distinct software modules. An *agent*, that runs on every node, gathers information from various protocol layers and the wireless network card. It reports this information to a management server, called *manager*. The manager analyzes the data and takes appropriate actions. The manager may run on a single node (centralized architecture), or may run on a set of nodes (decentralized architecture) [41].

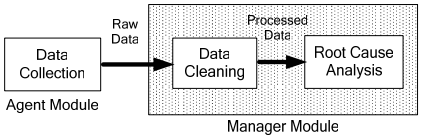

Figure 1: Troubleshooting process

Our three-step troubleshooting process is illustrated in Figure 1. The process starts by agents continuously collecting and transmitting their (local) view of the network's behavior to the manager(s). Examples of the information sent include traffic statistics, received packet signal strength on various links, and re-transmission counts on each link.

It is possible that the data the manager receives from the various agents results in an inconsistent view of the network. Such inconsistencies could be the result of topological and environmental changes, measurement errors, or misbehaving nodes. The *Data Cleaning* module of the manager resolves inconsistencies before engaging the analysis model.

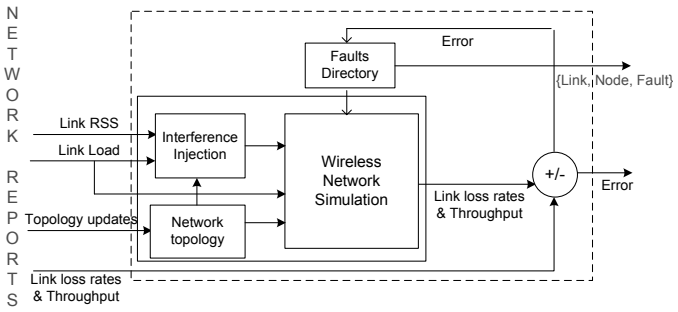

Figure 2: Root cause analysis module

After the inconsistencies have been resolved, the cleaned trace data is fed into the root-cause analysis module which contains a modified network simulator (see Figure 2). The analysis module drives the simulator with the cleaned trace data and establish the expected normal performance for the given network configuration and traffic patterns. Faults are detected when the expected performance does not match the observed performance. Root cause for the discrepancy is determined by efficiently searching for the set of faults that results in the best match between the simulated and observed network performance.

3. DESIGN RATIONALE

A wireless network is a complex system with many inter-dependent factors that affect its behavior. The factors include traffic flows, networking protocols, signal processing algorithms, hardware, RF propagation and, most importantly, the interactions between these impacts behavior. Additionally, network performance is also influenced by the interaction between nodes and external noise sources. We know of no heuristic or theoretical technique that captures these interactions and explains the behavior of such networks. In contrast, a high quality simulator provides valuable insights on what is happening inside the network.

As an example, consider a 7 * 3 grid topology network shown in Figure 3. Assume there are 5 long-lived flows F_1, F_2, F_3, F_4 and F_5 in the network, each with the same amount of traffic to communicate. All adjacent nodes can hear one another and the interference range is twice the communication range. The traffic between nodes A & O interferes with the traffic between nodes C & Q, and similarly traffic between nodes G & U interferes with the traffic between nodes E & S. However, neither traffic between G & U nor traffic between A & O interferes with traffic between D & R. Table 1 shows the throughput of the flows when each flow sends CBR traffic at a rate of 11 Mbps. As we can see, the flow F_3 receives much higher throughput than the flows F_2 and F_4.

F_1	F_2	F_3	F_4	F_5
2.50 Mbps	0.23 Mbps	2.09 Mbps	0.17 Mbps	2.55 Mbps

Table 1: Throughput of 5 competing flows in Figure 3

A simple heuristic may lead the manager to conclude that flow F_3 is unduly getting a larger share of the bandwidth, whereas an on-line trace-driven simulation will conclude that this is normal behavior. This is because the simulation takes into account the traffic flows and link quality, and based on the reported noise level it determines that flows F_1 and F_5 are interfering with flows F_2 and F_4, therefore allowing F_3 a open channel more often. Thus, the fact that F_3 is getting a greater share of the bandwidth will not be flagged as a fault by the simulator.

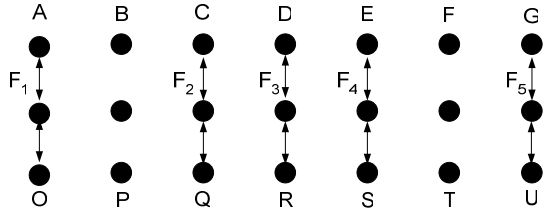

Figure 3: The flow F_3 gets a much higher share of the bandwidth than the flows F_2 and F_4, even though all the flows have the same application-level sending rate. A simple heuristic may conclude that nodes D and R are misbehaving, whereas simulation can correctly determine the observed throughput is expected.

Consequently, a good simulator is able to advise the manager on what constitutes normal behavior. When the observed behavior is different from what is determined to be normal, the manager can invoke the fault search algorithms to determine the reasons for the deviation.

In addition, while it might be possible to apply traditional signature-based or rule-based fault diagnosis approach to a particular type of network under a specific environment and configuration, simple signatures or rules do not capture the intrinsic complexity of fault diagnosis in general settings. In contrast, a simulator is customizable and with appropriate parameter settings, it can be applied to a large class of networks under different environments. Fault diagnosis built on top of such a simulator inherits its generality.

Finally, recent advances in simulators for multihop wireless networks, as evidenced in products such as Qualnet, have made the use of a simulator for real-time on-line analysis a reality. This is especially true for the relatively small-scale multihop wireless networks, up to a few hundred nodes, that we intend to manage.

4. SIMULATOR ACCURACY

We now turn our attention to the following question: "Can we build a fault diagnosis system using on-line simulations as the core tool?" The answer to the question cuts to the heart of our work. The viability of our system hinges on the accuracy with which the simulator can reproduce observed network behavior.

To answer this question, we quantify the challenge in matching the behavior of the network link layer and RF propagation. We then evaluate

the accuracy of trace-driven simulation. Finally we study how frequently the system needs to the adapt.

4.1 Physical Layer Discrepancies

Factors such as variability of hardware performance, RF environmental conditions, and presence of obstacles make it difficult for simulators to model wireless networks accurately [25]. To illustrate this problem, we conduct a simple experiment as follows.

We study the variation of received signal strength (RSS) with respect to distance for a variety of IEEE 802.11a cards (Cisco AIR-CB20A, Proxim Orinoco 8480-WD, and Netgear WAG511), and plot the results in Figure 4. The experiments are conducted inside a building with walls separating offices every 10 feet. As the distance increases, the number of walls (obstacles) between the two laptops also increases. The signal strength measurements are obtained using the wireless research API (WRAPI) [45]. For comparison, we also plot the RSS computed using the two-ray propagation model available in Qualnet [37]. This model is based solely on distance.

Note that the theoretical model does not estimate the RSS accurately. This is because it fails to take into account signal reflections from surrounding walls. Accurate modeling and prediction of wireless conditions is a hard problem to solve in its full generality but by replacing theoretical models with data obtained from the network we are able to significantly improve network performance estimation.

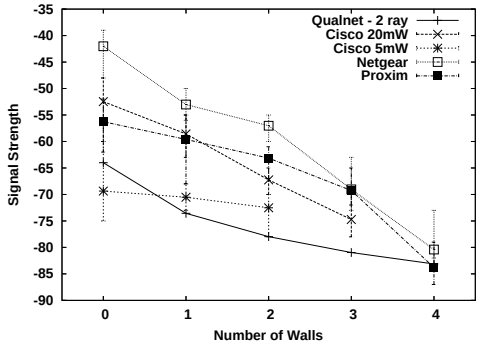

Figure 4: Comparing the simulator's two-ray RF wave propagation model for received signal strength with measurements taken from IEEE 802.11a WLAN cards from different hardware vendors.

In addition to the challenge of accurately modeling the physical layer and RF propagation, traffic demands from networked nodes are hard to predict. Fortunately, for the purpose of fault diagnosis, it is not necessary to have predictive models, and is sufficient to simulate what happened in the network *after the fact*. To do this, we require agents to periodically report information about the link conditions and traffic patterns to the manager. This information is processed and fed into the simulator. This approach overcomes the known limitations of RF propagation and traffic modeling in simulators in the context of fault diagnosis.

4.2 Baseline Comparison

Next we compare the performance of a real network to that of a simulator for a few simple baseline cases. We design a set of experiments to quantify the accuracy of simulating the overhead of the protocol stack as well as the effect of RF interference. The experiments are for the following scenarios:

1. A single one-hop UDP flow (1-hop flow)

2. Two UDP flows within communication range (2 flows - CR)

3. Two UDP flows within interference range (2 flows - IR)

4. One UDP flow with 2 hops where the source and destination are within communication range. We enforce the 2-hop route using static routing. (2-hop flow -CR)

5. One UDP flow with 2 hops where the source and destination are within interference range but not within communication range. (2-hop flow -IR)

All the throughput measurements are done using Netgear WAG511 cards and Figure 5 summarizes the results. Interestingly, in all cases the throughput from simulations are close to the real measurements. Case (1) shows that Qualnet simulator models the overheads of the protocol stack, such as parity bits, MAC-layer back-off, IEEE 802.11 inter-frame spacing and ACK, and headers accurately. The other scenarios show that the simulator accurately takes into account contention from flows within the interference and communication ranges.

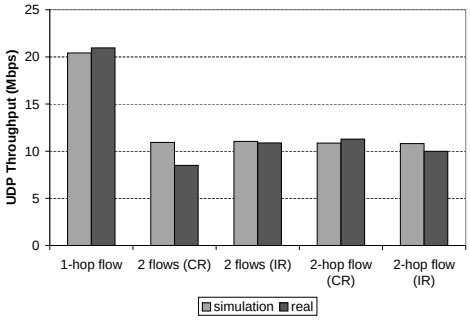

Figure 5: Estimated throughput from the simulator matches measured throughput in a real network when the RF condition of the links is good.

In the scenarios above, data are sent on high-quality wireless links, and almost never gets lost due to low signal strength. In our next experiment, we study how RSS affects throughput. We vary the number of walls between the sender and receiver, and plot the UDP throughput for varying packet sizes in Figure 6.

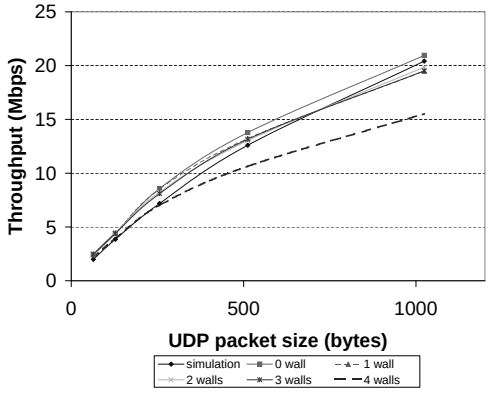

Figure 6: Estimated throughput matches with measured throughput when the RF condition of the links is good, and deviates when the RF condition of the links is poor (1-hop connection).

When the signal quality is good (e.g., when there are fewer than 4 walls in between), the throughput measured matches closely with the estimate from the simulator.

When the signal strength is poor, e.g., when 4 or more walls separate the two laptops, the throughput estimated by the simulator deviates from real measurements. The deviation occurs because the simulator does not take into account the following two factors:

# walls	loss rate	measured throughput	simulated throughput
4	11.0%	15.52 Mbps	15.94 Mbps
5	7.01%	12.56 Mbps	14.01 Mbps
6	3.42%	12.97 Mbps	11.55 Mbps

Table 2: Estimated and measured throughput match, when we compensate the loss rates due to poor RF in the real measurements by seeding the corresponding link in a simulator with an equivalent loss rate.

- Accurate packet loss as a function of packet-size, RSS, and ambient noise. This function depends on the signal processing hardware and the RF antenna within the wireless card.
- Accurate auto-rate control. On observing a large number of packet retransmissions at the MAC layer, many WLAN cards adjust their sending rate to something that is more appropriate for the conditions. The exact details of the algorithm used to determine a good sending rate differ from cards to cards.

Of the two factors mentioned above, the latter can be taken care of as follows. If auto-rate is in use, we again employ trace-driven simulation as follows: we monitor the rate at which the wireless card is operating, and provide it to the simulator (instead of having the simulator adapt in its own way). When auto-rate is not used (e.g., other researchers [10] have shown that auto-rate is undesirable when considering aggregate performance and therefore it should be turned off), the data rate is known.

The first issue is much harder to address because it may not be possible to accurately simulate the physical layer. One possible way to address this issue is through offline analysis. We calibrate the wireless cards under different scenarios and create a database to associate environmental factors with expected performance. For example, we carry out real measurements under different signal strengths and noise levels to create a mapping from signal strength and noise to loss rate. Using such a table in simulations allows us to distinguish between losses caused by collisions from losses caused by poor RF conditions. We evaluate the feasibility of this approach by computing the correlation coefficient between RSS and loss rates when the sending rate remains the same. We find the correlation coefficient ranges from -0.95 to -0.8. The high correlation suggests that it is feasible to estimate loss caused by poor RF conditions. [1]

Based on this idea, in our experiment we collect another set of traces in which we slowly send out packets so that most losses are caused by poor signal (instead of congestion). We also place packet sniffers near both the sender and receiver, and derive the loss rate from the packet-level trace. Then we take into the account the loss rate due to poor signal by seeding the wireless link in a simulator with a Bernoulli loss rate that matches the loss rate in real traces.

We find that after taking into account the impact of poor signal, the throughput from simulation matches closely with real measurements as shown in Table 2. Note that the loss rate and measured throughput do not monotonically decrease with the signal strength due to the effect of auto-rate, i.e., when the data rate decreases as a result of poor signal strength, the loss rate improves. (The driver of the wireless cards we use does not allow us to disable auto-rate.) Note that even though the match is not perfect, we do not expect this to be a problem in practice because several routing protocols try to avoid the use of poor quality links by employing some appropriate routing metrics (e.g., ETX [18], ETT [33]).

[1] Some researchers [5] report weak correlation between RSS and loss rates. They attribute multipath as the major reasons for packet loss. The difference between our finding and theirs is that they use RSS values reported by the receiver, whereas we use RSS reported by airopeek [7] running on a separate machine next to the receiver. Using RSS from the receiver introduces bias since RSS of corrupted packets is not available; in comparison, airopeek reports RSS for both received and corrupted packets, and the effects of multipath are reflected by the reported RSS values. A simple driver hack allows us to get the RSS value for corrupted packets.

4.3 Stability of Channel Conditions

So far, we have shown that with the help of trace collection a simulator is able to mimic reality. However, one question remains: how rapidly do conditions change and how often do we collect a trace? When the channel conditions are fluctuating very rapidly, collecting an accurate trace and shipping the trace to the manager may be difficult and costly. Figure 7 shows the temporal fluctuation in RSS over 10 minutes under the same measurement setup as described above. As expected, the RSS fluctuates over time. Fortunately, from our diagnosis perspective, the magnitude of the fluctuation is not significant, and the relative quality of the signals across different numbers of walls remains stable. This suggests that the environment is generally static, and the nodes may report only the average and standard deviation of the RSS to the manager every few minutes (e.g., 1 - 5 minutes).

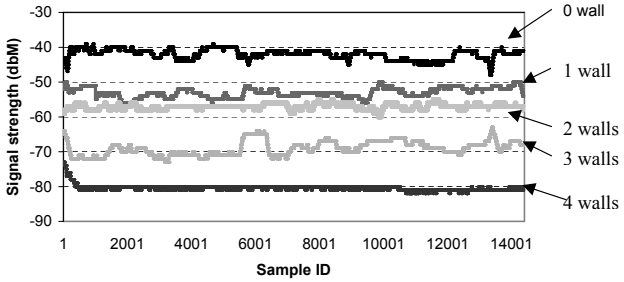

Figure 7: In good environmental conditions, received signal strength remains stable over time.

4.4 Remarks

In this section, we have shown that even though simulating a wireless network accurately is a hard problem, for the purpose of fault diagnosis, we can use trace-based simulations to reproduce what happened in the real network, after the fact. To substantiate this claim, we look at a number of simple scenarios and show that the throughput obtained from the simulator matches reality after taking into account information from real traces. We require only a small amount of data collected at the nodes, at a fairly low time-granularity.

5. FAULT ISOLATION AND DIAGNOSIS

We now present our simulation-based diagnosis approach. Our high-level idea is to re-create the environment that resembles the real network inside a simulator. To find the root cause, we *search over a fault space* to determine which fault or set of faults can re-produce performance similar to what has been observed in the real network.

In Section 5.1 we extend the trace-driven simulation ideas presented in Section 4 to reproduce network topology and traffic pattern observed in the real network.

Using trace-driven simulation as a building block, we then develop a diagnosis algorithm to find root-causes for faults. The algorithm first establishes the expected performance under a given set of faults. Then based on the difference between the expected and observed performance, it efficiently searches over the fault space to re-produce the observed symptoms. This algorithm can not only diagnose multiple faults of the same type, but also perform well in the presence of multiple types of faults.

Finally we address the issue of how to diagnose faults when the trace data used to drive simulation contains errors. This is a practical problem since data in the real world is never perfect for a variety of reasons, such as measurement errors, nodes supplying false information,

and software/hardware errors. To this end, we develop a technique to effectively eliminate erroneous data from the trace so that we can use good quality trace data to drive simulation-based fault diagnosis.

5.1 Trace-Driven Simulation

Taking advantage of trace data enables us to accurately capture the current environment and examine the effects of a given set of faults in the current network.

5.1.1 Trace Data Collection

We collect the following sets of data as input to a simulator for fault diagnosis:

- Network topology: Each node reports its neighbors. To be efficient, only changes in the set of neighbors are reported.
- Traffic statistics: Each node maintains counters for the volume of traffic sent to and received from its immediate neighbors. This data drives traffic simulation described in Section *5.1.2*.
- Physical medium: Each node reports its noise level and the signal strength of the wireless links from its neighbors. According to the traces collected from our testbed, we observe that the signal strength is relatively stable over tens of seconds. Slight variations in signal strength with time can be captured accurately through the time average, standard deviation, and other statistical aggregates.
- Network performance: To detect anomalies, we compare the observed network performance with the expected performance from simulation. Network performance includes both link performance and end-to-end performance, both of which can be measured through a variety of metrics, such as packet loss rate, delay, and throughput. In our work, we focus on link level performance.

Data collection consists of two steps: collecting raw performance data at a local node and distributing the data to collection points for analysis. For local data collection, we can use a variety of tools, such as WRAPI [45], Native 802.11 [28], SNMP [14], and packet sniffers (e.g., Airopeek [7], tcpdump [43]).

Distributing the data to a manager introduces overhead. In Section 7.1, we quantify this overhead, and show it is low and has little impact on the data traffic in the network. Moreover, it is possible to further reduce the overhead using compression, delta encoding, multicast, and adaptive changes of the time scale and spatial scope of distribution. For example, in the normal situation, a minimum set of information is collected and exchanged. Once the need arises for more thorough monitoring (e.g., when the information being collected indicates anomaly), then the manager requests more information and increases the frequency of data collection for the subset of the nodes that require intensive monitoring.

5.1.2 Simulation Methodology

We classify the characteristics of the network that need to be matched in the simulator into the following three categories: (i) traffic load, (ii) wireless signal, and (iii) faults. Below we describe how to simulate each of these components.

Traffic Load Simulation: A key step in replicating the real network inside a simulator is to re-create the same traffic pattern. One approach is to simulate end-to-end application demands. However, there can be potentially $N*N$ demands for an N-node network. Moreover, given the heterogeneity of application demands and the use of different transport protocols, such as TCP, UDP, and RTP, it is challenging to obtain end-to-end demands.

For scalability and to avoid the need for obtaining end-to-end demands and routing information, we use link-based traffic simulation. Our high-level idea is to adjust application-level sending rate at each link to match the observed link-level traffic counts. Doing this abstracts away higher layers such as the transport and the application layer, and allows us to concentrate only on packet size and traffic rate. However, matching the sending rate on a per-link basis in the simulator is non-trivial because we can only adjust the application-level sending rate, and have to obey the medium access control (MAC) protocol. This implies that we cannot directly control sending rate on a link. For example, when we set the application sending rate of a link to be 1 Mbps, the actual sending rate (on the air) can be lower due to back-off at the MAC layer, or higher due to MAC level retransmission. The issue is further complicated by interference, which introduces inter-dependency between sending rates on different links.

To address this issue, we use the following *iterative search* to determine the sending rate at each link. There are at least two search strategies: (i) multiplicative increase and multiplicative decrease, and (ii) additive increase and additive decrease. As shown in Figure 8, each link individually tries to reduce the difference between the current sending rate in the simulator and the actual sending rate in the real network. The process iterates until either the rate becomes close enough to the target rate (denoted as $targetMacSent$) or the maximum number of iterations is reached. We introduce a parameter α, where $\alpha \leq 1$, to dampen oscillation. In our evaluation, we use $\alpha = 0.5$ for $i \leq 20$, and $\frac{1}{i}$ for $i > 20$. This satisfies $\sum_i \alpha_i \rightarrow \infty$, and $\alpha_i \rightarrow 0$ as $i \rightarrow \infty$, and ensures convergence. Our evaluation uses multiplicative increase and multiplicative decrease, and we plan to compare it with additive increase and additive decrease in the future.

```
while (not converged and i < maxIterations)
    i = i + 1;
    if (option == multiplicative)
      foreach link(j)
        prevRatio = targetMacSent(j)/simMacSent(j);
        currRatio = (1 − α) + α ∗ prevRatio;
        simAppSent(j) = prevAppSent(j) ∗ currRatio;
    else // additive
      foreach link(j)
        diff = targetMacSent(j) − prevMacSent(j);
        simAppSent(j) = prevAppSent(j) + α ∗ diff;
    run simulation using simAppSent as input
    determine simMacSent for all links from simulation results
    converged = isConverge(simMacSent, targetMacSent)
```

Figure 8: Searching for the application-level sending rate using either multiplicative increase, multiplicative decrease or additive increase additive decrease.

Wireless Signal: Signal strength has a very important impact on wireless network performance. As discussed in Section 4, due to variations across different wireless cards and environments, it is hard to come up with a general propagation model to capture all the factors. To address this issue, we drive simulation using the real measurement of signal strength and noise, which can be easily obtained using newer generation wireless cards (e.g., Native 802.11 [28]).

Fault Injection: To examine the impact of faults on the network, we implement the ability to inject different types of faults into the simulator, namely (i) packet dropping at hosts, (ii) external noise sources, and (iii) MAC misbehavior [31].

- Packet dropping at hosts: a misbehaving node drops some traffic from one or more neighbors. This can occur due to hardware/software errors, buffer overflow, and/or malicious drops. The ability to detect such end-host packet dropping is useful, since it allows us to differentiate losses caused by end hosts from losses caused by the network.
- External noise sources: we support the ability to inject external noise sources in the network.
- MAC misbehavior: a faulty node does not follow the MAC etiquette and obtains an unfair share of the channel bandwidth. For

example, in IEEE 802.11 [31], a faulty node can choose a smaller contention window (CW) to send traffic more aggressively [26].

In addition, we also generate link congestion by putting a high load on the network. Unlike the other types of faults, link congestion is implicitly captured by the traffic statistics gathered from each node. Therefore trace-driven simulation can directly assess the impact of link congestion. For the other three types of faults, we apply the algorithm described in Section 5.2 to diagnose them.

5.2 Fault Diagnosis Algorithm

We now describe an algorithm to systematically diagnose root causes for failures and performance problems.

General approach: Applying simulations to fault diagnosis enables us to reduce the original diagnosis problem to the problem of searching for a set of faults such that their injection results in an expected performance that matches well with observed performance. More formally, given a network settings, NS, our goal is to find $FaultSet$ such that $SimPerf(NS, FaultSet) \approx ObservedPerf$, where the performance is a function value, which can be quantified using different metrics. It is clear that the search space is high-dimensional due to many combinations of faults. To make the search efficient, we take advantage of the fact that different types of faults often change one or few metrics. For example, packet dropping at hosts only affects link loss rate, but not the other metrics. Therefore we can use the metrics in which the observed and expected performance have significant difference to guide our search. Below we introduce our algorithm.

Initial diagnosis: We start by considering a simple case where all faults are of the same type, and the faults do not have strong interactions. We will later extend the algorithm to handle more general cases, where we have multiple types of faults, or faults that interact with each other.

For ease of description, we use the following three types of faults as examples: packet dropping at hosts, external noise, and MAC misbehavior, but the same methodology can be extended to handle other types of faults once the symptoms of the fault are identified.

As shown in Figure 9, we use trace-driven simulation, fed with current network settings, to establish the expected performance. Based on the difference between the expected performance and observed performance, we first determine the type of faults using a decision tree as shown in Figure 10. Due to many factors, simulated performance is unlikely to be identical with the observed performance even in the absence of faults. Therefore we conclude that there are anomalies only when the difference exceeds a threshold. The fault classification scheme takes advantage of the fact that different faults exhibit different behaviors. While their behaviors are not completely non-overlapping (e.g., both noise sources and packet dropping at hosts increase loss rates; lowering CW increases the traffic and hence increases noise caused by interference), we can categorize the faults by checking the differentiating component first. For example, external noise sources increase noise experienced by its neighboring nodes, but do not increase the sending rates of any node, and therefore can be differentiated from MAC misbehavior and packet dropping at hosts.

After the fault type is determined, we then locate the faults by finding the set of nodes and links that have large differences between the observed and expected performance. The fault type determines what metric is used to quantify the performance difference. For instance, we identify packet dropping by finding links with large difference between the expected and observed loss rates. We determine the magnitude of the fault using a function $g()$, which maps the impact of a fault into its magnitude. For example, under the end-host packet dropping, $g()$ function is the identity function, since the difference in a link's loss rate can be directly mapped to a change in dropping rate on a link (fault's magnitude); under the external noise fault, $g()$ is a propagation function of a noise signal.

1) Let NS denote the network settings
(i.e., signal strength, traffic statistics, network topology)
Let $RealPerf$ denote the real network performance
2) $FaultSet = \{\}$
3) Predict $SimPerf$ by running simulation with input $(NS, FaultSet)$
4) if $|Diff(SimPerf, RealPerf)| > threshold$
determine the fault type ft using the decision tree shown in Fig. 10
for each link or node i
if $(|Diff_{ft}(SimPerf(i), RealPerf(i))| > threshold)$
add $fault(ft, i)$ with
$magnitude(i) = g(Diff_{ft}(SimPerf(i), RealPerf(i))$

Figure 9: Initial diagnosis: one pass diagnosis algorithm

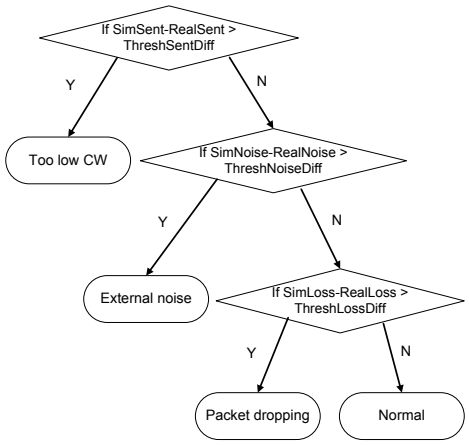

Figure 10: An algorithm to determine the type of faults

The algorithm: In general, we may have multiple types of faults interacting with each other. Even when all the faults are of the same type, they may still interact, and their interactions may make the above one pass diagnosis insufficient. To address these challenges, we develop an interactive diagnosis algorithm, as shown in Figure 11, to find root causes.

The algorithm consists of two stages: (i) initial diagnosis stage, and (ii) iterative refinements. During the initial diagnosis stage, we apply the one-pass diagnosis algorithm described above to come up with the initial set of faults; then during the second stage, we iteratively refine the fault set by (i) adjusting the magnitude of the faults that have been already inserted into the fault set, and (ii) adding a new fault to the set if necessary. We iterate the process until the change in fault set is negligible (i.e., the fault types and locations do not change, and the magnitudes of the faults change very little).

We use an iterative approach to search for the magnitudes of the faults. At a high level, the approach is similar to the link-based simulation, described in Section *5.1.2*, where we use the difference between the target and current values as a feedback to progressively move towards the target. In more details, during each iteration, we first estimate the expected network performance under the existing fault set. Then we compute the difference between simulated network performance (under the existing fault set) and real performance. Next we translate the difference in performance into change in faults' magnitudes using the function $g()$. After updating the faults with new magnitudes, we remove the faults whose magnitudes are too small.

In addition to searching for the correct magnitudes of the faults, we also iteratively refine the membership of the fault set by finding new faults that can best explain the difference between expected and observed performance. To control false positives, during each iteration we only add the fault that can explain the largest mismatch.

5.3 Handling Imperfect Data

1) Let NS denote the network settings
(i.e., signal strength, traffic statistics, and network topology)
Let $RealPerf$ denote the real network performance
2) $FaultSet = \{\}$
3) Predict $SimPerf$ by running simulation with input $(NS, FaultSet)$
4) if $|Diff(SimPerf, RealPerf)| > threshold$
go to 5)
else
go to 7)
5) Initial diagnosis:
initialize $FaultSet$ by applying the algorithm in Fig. 9
6) while (not converged)
a) adjusting fault magnitude
for each fault type ft in $FaultSet$ (according to decision tree in Fig. 10)
for each fault i in $(FaultSet, ft)$
$magnitude(i) -= g(Diff_{ft}(SimPerf(i), RealPerf(i)))$
if $(|magnitude(i)| < threshold)$
delete the fault (ft, i)
b) adding new candidate faults if necessary
foreach fault type ft (in the order of decision tree in Fig. 10)
i) find a fault i s.t. it is not in $FaultSet$
and has the largest $|Diff_{ft}(SimPerf(i), RealPerf(i))|$
ii) if $(|Diff_{ft}(SimPerf(i), RealPerf(i))| > threshold)$
add (ft, i) to $FaultSet$ with
$magnitude(i) = g(Diff_{ft}(SimPerf(i), RealPerf(i))$
c) simulate
7) Report $FaultSet$

Figure 11: A complete diagnosis algorithm: diagnose faults of possibly multiple types

In the previous sections, we describe how to diagnose faults by using trace data to drive online simulation. In practice, the raw trace data collected may contain errors for various reasons as mentioned earlier. Therefore we need to clean the raw data before feeding it to a simulator for fault diagnosis.

To facilitate the data cleaning process, we introduce *neighbor monitoring*, in which each node reports performance and traffic statistics not only for its incoming/outgoing links, but also for other links within its communication range. Such information is available when a node is in the promiscuous mode, which is achievable using Native 802.11 [28].

Due to neighborhood monitoring, multiple reports from different nodes are likely to be submitted for each link. The redundant reports can be used to detect inconsistency. Assuming that the number of misbehaving nodes is small, our scheme identifies the misbehaving nodes as the minimum set of nodes that can explain the discrepancy in the reports. Based on the insight, we develop the following scheme.

In our scheme, a sender i reports the number of packets sent and the number of MAC-level acknowledgements received for a directed link l as $(sent_i(l), ack_i(l))$; a receiver j reports the number of packets received on the link as $recv_j(l)$; in addition, a sender or receiver's immediate neighbor k also reports the number of packets and MAC-level acknowledgement it observes sent or received on the link as ($sent_k(l)$, $recv_k(l)$, $ack_k(l)$). An inconsistency in the reports is defined as one of the following cases.

1. The number of packets received on a link, as reported by the destination, is noticeably larger than the number of packets sent on the same link, as reported by the source. That is, for the link l from node i to node j, and given a threshold t:

$$recv_j(l) - sent_i(l) > t$$

2. The number of MAC-level acknowledgments on a link, as reported by the source, does not match the number of packets received on that link, as reported by the destination. That is, for the link l from node i to node j, and given a threshold t:

$$\mid ack_i(l) - recv_j(l) \mid > t$$

3. The number of packets received on a link, as reported by the destination's neighbor, is noticeably larger than the number of packets sent on the same link, as reported by the source. That is, for the link l from node i to node j, j's neighbor k, and given a threshold t:

$$recv_k(l) - sent_i(l) > t$$

4. The number of packets sent on a link, as reported by the source's neighbor, is noticeably larger than the number of packets sent on the same link, as reported by the source. That is, for the link l from node i to node j, i's neighbor k, and given a threshold t:

$$sent_k(l) - sent_i(l) > t$$

Since nodes do not send their reports strictly synchronously, we need to use a threshold $t > 0$ to mask the resulting discrepancies. Note that in the absence of inconsistent reports, the above constraints cannot be violated as a result of lossy links.

We then construct an *inconsistency graph* as follows. For each pair of nodes whose reports are identified as inconsistent, we add them to the inconsistency graph, if they are not already in the graph; we add an edge between the two nodes to reflect the inconsistency. Based on the assumption that most nodes send reliable reports, our goal is to find the smallest set of nodes that can explain all the inconsistency observed. This can be achieved by finding the smallest set of vertices that covers the graph, where the identified vertices represent the misbehaving nodes.

This is essentially the minimum vertex cover problem [19], which is known to be NP-hard. We apply a greedy algorithm, which iteratively picks and removes the node with the highest degree and its incident edges from the current inconsistency graph until no edges are left.

History of traffic reports can be used to further improve the accuracy of inconsistency detection. For example, we can continuously update the inconsistency graph with new reports without deleting previous information, and then apply the same greedy algorithm to identify misbehaving nodes.

6. SYSTEM IMPLEMENTATION

We have implemented a prototype of network monitoring and management module on the Windows XP platform. In this section, we present the components of the prototype implementation, the design principles, and its features.

Our prototype consists of two separate components: *agents* and *managers*. An agent runs on every wireless node, and reports local information periodically or on-demand. A manager collects relevant information from agents and analyzes the information.

The two design principles we follow are: simplicity and extensibility. The information gathered and propagated for monitoring and management is cast into performance counters supported on Windows. Performance counters are essentially (name, value) pairs grouped by categories. This framework is easily extensible.

Values in these performance counters are not always read-only. Writable counters offer a way for an authorized manager to change the values and influence the behavior of a node in order to fix problems or initiate experiments remotely.

Each manager is also equipped with a graphical user interface (GUI) to interact with network administrators. The GUI allows an administrator to visualize the network as well as to issue management requests.

The manager is also connected to the back-end simulator. The information collected is processed and then converted into a script that drives the simulation producing fault diagnosis results.

The capability of the network monitoring and management depends heavily on the information available for collection. We have seen welcoming trends in both wireless NICs and the standardization efforts to expose performance data and control at the physical and MAC layers, e.g., Native 802.11 NICs [28].

7. EVALUATION

In this section, we present evaluation results. We begin by quantifying the network overhead introduced by data collection and show its impact on the overall performance. Next, we evaluate the effectiveness of our diagnosis techniques and inconsistency detection scheme. We use simulations in some of our evaluation because this enables us to inject different types of faults in a controlled and repeatable manner. When evaluating in simulation, we diagnose traces collected from simulation runs that have injected faults. Finally we report our experience of applying the approach to a small-scale testbed. Even though the results from the testbed are limited by our inability to inject some types of faults (external noise and MAC misbehavior) in a controlled fashion, they demonstrate the feasibility of on-line simulations in a real system. Unless stated differently, all results from simulations are based on IEEE 802.11b. The testbed results are based on IEEE 802.11a.

7.1 Data Collection Overhead

For data collection, every node not only collects information locally, but also delivers the data to the manager. We evaluate the overhead involved in having all nodes in a network report the information to a manager. Since the primary goal of this section is to demonstrate the feasibility of distributing all the data to a manager at a modest cost, we only make *conservative assumptions* about the sizes of the reports. Further optimization is possible as described in Section *5.1.1*.

In our evaluation, we place nodes randomly in a square, and the manager is chosen at random amongst the nodes. We keep the average number of neighbors around 6 as we increase the network size. On average, a node takes 1 to 5 hops to reach the manager. As described in Section *5.1.1*, for each link we collect traffic counters (i.e., the number of packets sent and received), signal strength, and noise. The size of link report depends on whether redundant information is sent for consistency checking. Therefore we consider two scenarios: when each link is reported by one node (i.e., without data cleaning) and when each link is reported by all the observers (including the sender and receiver) to allow us to check for consistency (i.e., with data cleaning). Since every node has around 6 immediate neighbors, a conservative estimate of a link report is 72 bytes when no redundant link data is sent, and is 312 bytes when redundant link data is sent.

In Figure 12, we plot the average overhead of data gathering over 10 random runs for various network sizes. As it shows, even with data cleaning using 60 second report interval, the overhead remains low, around 800 bits/s/node. Moreover, the overhead does not increase much as the network size increases.

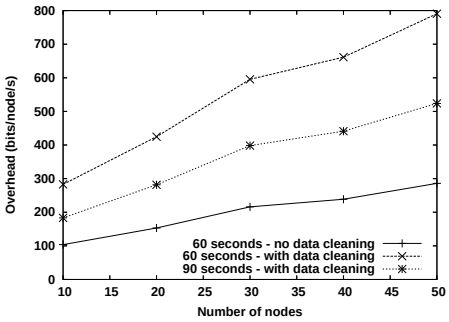

Figure 12: Management traffic overhead

Figure 13 shows the performance of FTP flows in the network with and without the data collection traffic. Ten simultaneous FTP flows are started from random sources to the manager. The graph shows the average throughput of these flows on the y-axis. As we can see, the data collection traffic (with and without sending redundant link data) has little impact on the application traffic in the network.

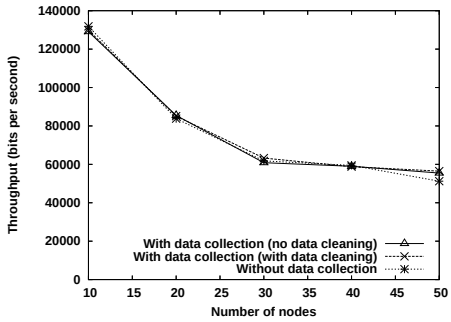

Figure 13: Effect of overhead on throughput

Summary: In this section, we evaluate the overhead of collecting traces, which will be used as inputs to diagnose faults in a network. We show that the data collection overhead is low and has little effect on application traffic in the network. Therefore it is feasible to use trace-driven simulation for fault diagnosis.

7.2 Evaluation of Fault Diagnosis through Simulations

In this section, we evaluate our fault diagnosis approach through simulations in Qualnet.

7.2.1 Diagnosing one or more faults of possibly different types

Our general methodology of using simulation to evaluate fault diagnosis is as follows. We artificially inject a set of faults into a network, and obtain the traces of network topology and link load under faults. We then feed these traces into the fault diagnosis module to infer root causes, and quantify the diagnosis accuracy by comparing the inferred fault set with the fault set originally injected.

We use both grid topologies and random topologies for our evaluation. In a grid topology, only nodes horizontally or vertically adjacent can directly communicate with each other, whereas in random topologies, nodes are randomly placed in a region. To challenge our diagnosis scheme, we put a high load on the network by randomly picking 25 pairs of nodes to send one-way constant bit rate (CBR) traffic at a rate of 1 Mbps. Under this load, the links in the network have significant network congestion loss, which makes diagnosis even harder. For example, identifying losses caused by packet dropping at hosts is more difficult when there is significant network congestion loss. Correct identification of dropping links also implies reasonable assessment of congestion loss. In addition, we randomly select a varying number of nodes to exhibit one or more faults of the following types: packet dropping at hosts, external noise, and MAC misbehavior. For a given number of faults and its composition, we conduct three random runs, which have different traffic patterns and fault locations. We evaluate how accurate our fault diagnosis algorithm, described in Section 5.2, can locate the faults. The time that the diagnosis process takes depends on the size of topologies, the number of faults, and duration of the faulty traces. For example, diagnosing faults in 25-node topologies takes several minutes. Such diagnosis time scale is acceptable for diagnosing long-term performance problems. Moreover, the efficiency can be significantly improved through code optimization.

We use coverage and false positive to quantify the accuracy of fault detection, where coverage represents the percentage of faulty locations that are correctly identified, and false positive is the number of (non-faulty) locations incorrectly identified as faulty divided by the total number of true faults. We consider a fault is correctly identified when both its type and its location are correct. For packet dropping and external noise sources, we also compare the inferred faults' magnitudes with their true

magnitudes.

Detecting packet dropping at hosts: We start by evaluating how accurately we can detect packet dropping. In our evaluation, we select a varying number of nodes to intentionally drop packets with the dropping rate varied between 0 - 100%. We vary the number of such misbehaving nodes from 1 to 6.

We apply the diagnosis algorithm, which first uses trace-driven simulation to estimate the expected performance (i.e., noise level, throughput, and loss rates) in the current network. Since we observe a significant difference in loss rates, but not in the other two metrics, we suspect that there is packet dropping on these links. We locate the dropping links by identifying links whose loss rates are significantly higher than their expected loss rates. We use 15% as a threshold so that links whose difference between expected and observed loss rates exceed 15% are considered as packet dropping links. We then inject the faults into the simulator, and find that this significantly reduces the difference between the simulated and observed performance.

Figure 14 shows the accuracy of detecting dropping links in a 5×5 grid topology. Note that some of the faulty links do not carry enough traffic to meaningfully compute loss rates. In our evaluation, we use 250 packets as a threshold so that only for the links that send over 250 packets, loss rates are computed. We consider a faulty link sending less than a threshold number of packets as a *no-effect fault* since it drops only a small number of packets.[2] As Figure 14 shows, under our diagnosis scheme, in most cases over 80% effective faulty links are identified correctly. The false positive (not shown) is 0 except for two cases in which one link is misidentified as faulty. Moreover the accuracy does not degrade with the increasing number of faults. When we compare the difference between the inferred and true dropping rates, we find the inference error, computed as $\sum_i |infer_i - true_i| / \sum_i true_i$, is within 25%. This error is related to the threshold used to determine if the iteration has converged. In our simulations, we consider an iteration converges when changes in loss rates are all within 15%. We can further reduce the inference error by using a smaller threshold at a cost of longer running time. Also, in many cases it suffices to know where packet dropping occurs without knowing precise dropping rates.

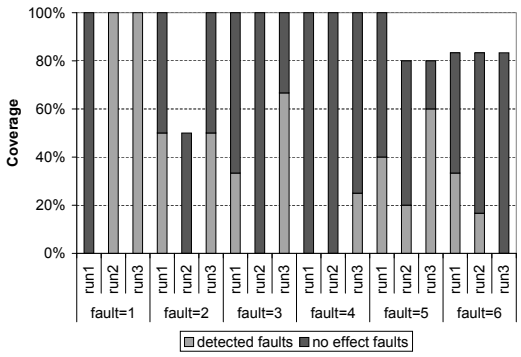

Figure 14: Accuracy of detecting packet dropping in a 5×5 grid topology

Detecting external noise sources: Next we evaluate the accuracy of detecting external noise sources. We randomly select a varying number of nodes to generate ambient noise at 1.1e-8 mW. We again use the trace-driven simulation to estimate the expected performance under the current traffic and network topology when there is no noise source. Note that simulation is necessary to determine the expected noise level, because the noise experienced by a node consists of both ambient noise and noise due to interfering traffic; accurate simulation of network traffic is needed to determine the amount of noise contributed by interfering traffic. The diagnosis algorithm detects a significant difference (e.g., over 5e-9mW) in noise level at some nodes, and conjectures that these nodes generate extra noise. It then injects noise at these nodes with magnitude derived from the difference between expected and observed noise level to the simulator. After noise injection, it sees a close match between the observed and expected performance, and hence concludes that the network has the above faults.

Figure 15 shows the accuracy of detecting noise generating sources in a 5×5 grid topology. As we can see, in all cases noise sources are correctly identified with at most one false positive link. We also compare the inferred magnitudes of noises with their true magnitudes, and find the inference error, computed as $\sum_i |infer_i - true_i| / \sum_i true_i$, is within 2%.

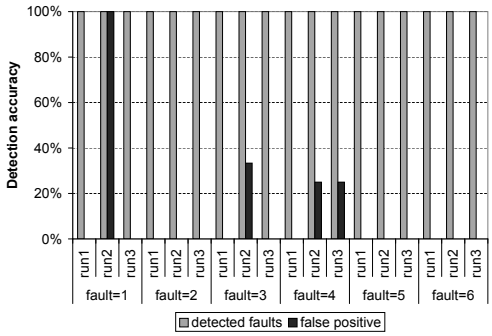

Figure 15: Accuracy of detecting external noise sources in a 5×5 grid topology

Detecting MAC misbehavior: Now we evaluate the accuracy of detecting MAC misbehavior. In our evaluation, we consider one implementation of MAC misbehavior. But since our diagnosis scheme is to detect unusually aggressive senders, it is general enough to detect other implementations of MAC misbehavior that exhibit similar symptoms. In our implementation, a faulty node alters its minimum and maximum MAC contention window in 802.11 (CWMin and CWMax) to be only half of the normal values. The faulty node continues to obey the CW updating rules (i.e., when transmission is successful, CW = CWMin, and when a node has to retransmit, CW = min((CW+1)*2-1, CWMax)). However since its CWMin and CWMax are both half of the normal, its CW is usually around half of the other nodes'. As a result, it transmits more aggressively than the other nodes. As one would expect, the advantage of using a lower CW is significant when network load is high. Hence we evaluate our detection scheme under a high load.

In our diagnosis, we use the trace-driven simulation to estimate the expected performance under the current traffic and network topology, and detect a significant discrepancy in throughput (e.g., the ratio between observed and expected throughput exceeds 1.25) on certain links. Therefore we suspect the corresponding senders have altered their CW. After injecting the suspected faults, we see a close match between the simulated and observed performance. Figure 16 shows the diagnosis accuracy in a 5×5 topology. We observe the coverage is mostly around 70% or higher. The false positive (not shown) is zero in most cases; the only case in which it is non-zero is when there is only one link misidentified as faulty.

Detecting mixtures of packet dropping and MAC misbehavior: Next we examine how accurately the diagnosis algorithm can handle multiple types of faults. First we consider mixtures of packet dropping and MAC misbehavior. To challenge the diagnosis scheme, we choose pairs of nodes adjacent to each other with one node randomly dropping one of its neighbors' traffic and the other node using an unusually small CW. We vary the number of node pairs selected to misbehave from 1 to 6

[2]These faulty links may have impact on route selection. That is, due to its high dropping rate, it is not selected to route much traffic. In this paper, we focus on diagnosing faults on data paths. As part of our future work, we plan to investigate how to diagnose faults on control paths.

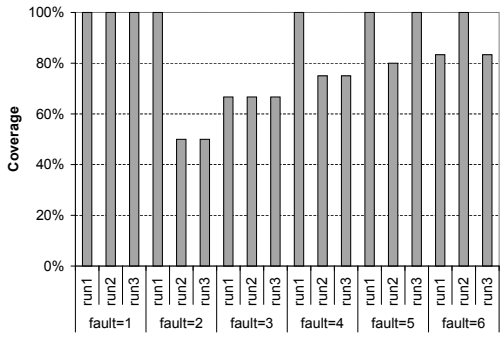

Figure 16: Accuracy of detecting MAC misbehavior in a 5×5 grid topology

Topology	# Faults	4	6	8	10	12	14
25-node random	Coverage	100%	100%	75%	90%	75%	93%
	False positive	25%	0	0	0	0	7%
7×7 grid	Coverage	100%	83%	100%	70%	67%	71%
	False positive	0	0	0	0	8%	0

Table 3: Accuracy of detecting combinations of packet dropping, MAC-misbehavior, and external noises in other topologies

(i.e., the total number of faults varies from 2 to 12 in the network). Figure 17 summarizes the accuracy of fault diagnosis in a 5×5 grid topology. As it shows, in most cases over 80% faults are correctly identified. Moreover, the false positive (not shown) is close to 0 in all cases. Comparing the inferred link dropping rates with their actual rates, we observe the inference error is within 30%.

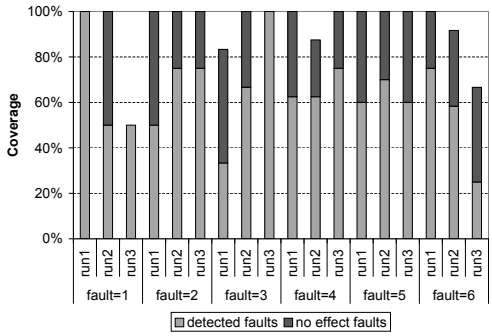

Figure 17: Accuracy of detecting combinations of packet dropping and MAC misbehavior in a 5×5 grid topology

Detecting mixtures of all three fault types: Finally we evaluate the diagnosis algorithm under mixtures of all three fault types as follows. As in the previous evaluation, we choose pairs of nodes adjacent to each other with one node randomly dropping one of its neighbors' traffic and the other node using an unusually small CW. In addition, we randomly select two nodes to generate external noise. Figure 18 summarizes the accuracy of fault diagnosis in a 5×5 topology. As it shows, the coverage is above 80%. The false positive (not shown) is close to 0. The accuracy remains high even when the number of faults in the network exceeds 10. The inference errors in links' dropping rate and noise level are within 15% and 3%, respectively.

To test sensitivity of our results to the network size and type of topology, we then evaluate the accuracy of the diagnosis algorithm using a 7×7 grid topology and 25-node random topologies. In both cases, we randomly choose 25 pairs of nodes to send CBR traffic at 1 Mbps rate. Table 3 summarizes results of one random run. As it shows, we can identify most faults with few false positives.

Summary: To summarize, we have evaluated the fault diagnosis approach using a variety of scenarios, and shown it yields fairly accurate results.

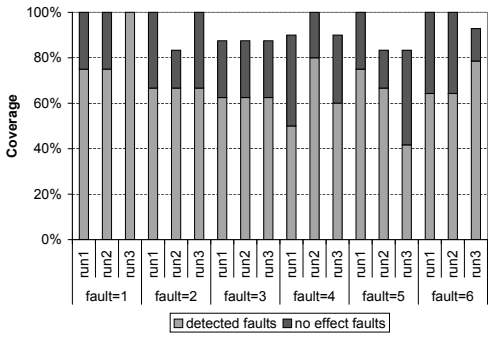

Figure 18: Accuracy of detecting combinations of packet dropping, MAC misbehavior, and external noises in a 5×5 grid topology

7.3 Data cleaning effectiveness and overhead

As mentioned earlier, to deal with data imperfectness, we need to process the raw data by applying the inconsistency detection scheme described in Section 5.3 before feeding them to the diagnosis module. In this section, we evaluate the effectiveness of this scheme using different network topologies, traffic patterns, and degrees of inconsistency.

- Network topologies: We use both random and grid topologies for evaluation. In the former, we randomly place nodes in a region while in the latter we place nodes in an $L \times L$ grid, where only the nodes horizontally or vertically adjacent can directly communicate with each other. We vary the size of the region to evaluate how node density affects the accuracy of inconsistency detection, while fixing the total number of nodes at 49 in all cases.

- Traffic patterns: We generate CBR traffic in the network. We consider two types of traffic patterns:

  1. Client-server traffic: in this case, we place one server at the center of the network to serve as an Internet gateway, and the other 48 nodes all establish connections from themselves to the gateway. We assume that the performance reports generated by the server are correct, and if a client's report deviates from the server's, it is the client that supplied incorrect information.
  2. Peer-to-peer traffic: we randomly select pairs of nodes from the network to transfer CBR traffic. We keep the number of connections the same as in client-server traffic.

- Inconsistent reports: We randomly select a varying fraction of nodes to report incorrect information. In addition, for every such node, we vary the fraction of its adjacent links that are reported incorrectly. We use d to denote the fraction, where $d = 1$ means that the selected node reports all the adjacent links incorrectly, while $d < 1$ means that the selected node reports a fraction of its adjacent links incorrectly.

We again use coverage and false positive to quantify the accuracy, where coverage denotes the fraction of misbehaving nodes that are correctly identified, whereas false positive is the ratio between the number of nodes that are incorrectly identified as misbehaving and the number of true misbehaving nodes.

Effects of node density: Figure 19 shows the effect of node density on the fraction of misbehaving nodes detected and false positives in random topologies. When the area is a 1400m × 1400m, a node has 7 to 8

neighbors within communication range on average, whereas in a 2450m $\times$ 2450m region, a node only has 2 to 3 neighbors on average.

We make the following observations. First, the detection accuracy is high: except for the lowest node density, in most cases the coverage is above 80% and false positive (not shown) is below 15%. Second, as one would expect, the detection accuracy tends to be higher in a denser topology than in a sparser topology. This is because in a denser topology, there are more observers for each link, and majority voting works better. Note that the accuracy does not strictly decrease with the network size due to random selection of misbehaving nodes.

Effects of traffic types: Also, in Figure 19, we see that with peer-to-peer traffic, the detection accuracy is lower than with client-server traffic. This is because for client-server traffic, we trust the server to report correct information; we can detect misbehaving clients whenever their reports deviate from that of the server. In comparison, in peer-to-peer traffic, all nodes are treated equally and we only rely on majority voting.

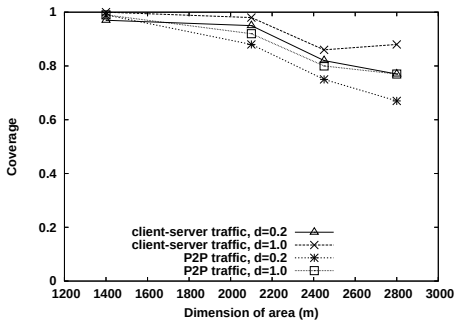

Figure 19: Detection accuracy in random topologies with varying d, node density, and traffic patterns.

Effects of number of misbehaving nodes: Figure 20 plots the detection accuracy versus the number of misbehaving nodes that report incorrect information in the network. The density is held constant by holding the region fixed at a 2450m $\times$ 2450m square. Here, we plot the accuracy for both the grid and the random topologies. As it show, the accuracy is high even when a large fraction (40%) of the nodes in the system are misbehaving. In all cases, the coverage is higher than 80%, and the false positives (not shown) are lower than 12%.

Effects of topology type: Next we examine the effects of network topologies on detection accuracy. We compare the detection accuracy in the grid topology against the random topology, both spanning 2450m $\times$ 2450m. In Figure 20, we can see that the grid topology almost always has a higher detection accuracy than the random topology. A closer look of the topology reveals that while the average node degree in the grid and random topologies are comparable, both around 2 to 3, the variation in node degree is significantly higher in the random topology. There are significantly more nodes with only one neighbor in the random topology. In this case, it is hard to detect which node supplies wrong information. In comparison, nodes in grid topologies have a similar number of neighbors (only corner nodes have fewer neighbors), and no nodes have fewer than 2 neighbors, which makes it easier for majority voting. This observation suggests that the minimum node degree is more important to detection accuracy than the average node degree.

Incorporating history: So far we have studied the case when every node submits one traffic report at the end of simulation. Now we evaluate the case in which nodes periodically send report. In this case, we can take advantage of history information as described in Section 5.3. As shown in Figure 21, we observe a higher coverage when using history information. Similarly, the false positive (not shown) is lower when history information is incorporated.

Summary: In summary, we show that the inconsistency detection scheme is able to detect most misbehaving nodes with very few false positives under a wide variety of scenarios we have considered. These results suggest that after data cleaning, we can obtain good quality trace data to drive simulation-based diagnosis.

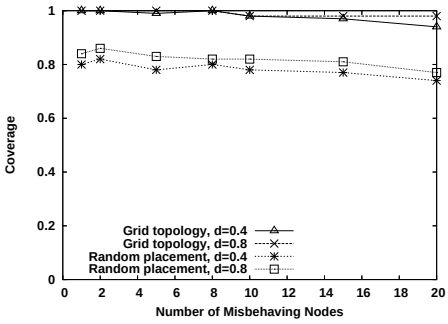

Figure 20: Detection accuracy when the nodes are placed in a 2450m $\times$ 2450m region with varying topology types and the number of misbehaving nodes.

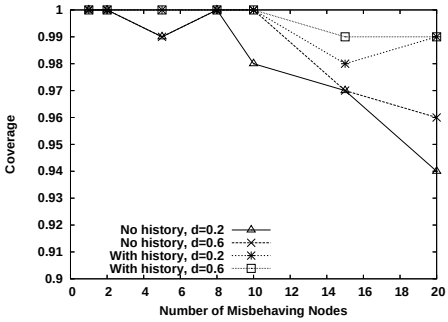

Figure 21: Comparing detection accuracy between with and without using history in a 2450m $\times$ 2450m grid topology using peer-to-peer traffic.

7.4 Evaluation of Fault Diagnosis in a Testbed

In this section, we evaluate our approach using experiments in a testbed. Our testbed consists of 4 laptops, each equipped with a Netgear WAG511 card operating in 802.11a mode. The laptops are located in the same office with good received signal strength. Each of them runs a routing protocol, similar to DSR [22], to determine the shortest hop-count paths to the other nodes. However, due to packet losses caused by high traffic load and artificial packet dropping, the nodes sometimes switch between 1-hop routes and 2-hop routes. The traffic statistics on all links are periodically collected using the monitor tool, described in Section 6. We randomly pick a node to drop packet from one of its neighbors, and see if we can detect it. To resolve inconsistencies in traffic reports if any, we also run Airopeek [7] on another laptop. (Ideally we would like to have nodes monitor traffic in the promiscuous mode, e.g., using Native 802.11 [28], but since we currently do not have such cards, we use Airopeek to resolve inconsistencies.)

First, we run experiments under low traffic load, where each node sends CBR traffic at a rate varying from 1 Mbps to 4 Mbps to another node. We find the collected reports are consistent with what has been observed from Airopeek. Then we feed the traces to the simulator (also running in the 802.11a mode), and apply the diagnosis algorithm in Figure 11. Since in the testbed one node is instructed to drop one of its neighbor's traffic at a rate varying from 20% to 50%, the diagnosis algorithm detects that there is a significant discrepancy between the expected and observed loss rates on one link, and correctly locates the dropping

link.

Then we repeat the experiments when we overload the network by having each node sending CBR traffic at a rate of 8 Mbps. In this case, we observe that the traffic reports often deviate from the numbers seen in Airopeek. The deviation is caused by the fact that the NDIS driver for the NIC sometimes indicates sending success without actually attempting to send the packet to the air [29]. This implies that it is not always possible to keep an accurate count of the packets sent locally. However, the new generation of wireless cards, such as Native 802.11 [28], will expose more detailed information about the packets, and enable more accurate accounting of traffic statistics. The inaccurate traffic reports observed in the current experiments also highlight the importance of cleaning the data before using them for diagnosis. In our experiment, we clean the data using Airopeek's reports, which capture almost all the packets in the air, and feed the cleaned data to the simulator to estimate the expected performance. Applying the same diagnosis scheme, we derive the expected congestion loss, based on which we correctly identify the dropping link.

8. DISCUSSION

To the best of our knowledge, ours is the first system that integrates a network simulator into a network management system to troubleshoot an operational multihop wireless network. The results are promising.

Our diagnosis system is not limited to the four types of faults discussed in this paper. Other faults such as routing misbehavior can also be diagnosed. Since routing misbehavior has been the subject of much previous work [27, 12, 21], we focus on diagnosing faults on the data path, which have not received much attention. In general, the fault to be diagnosed determines the traces to collect and the level of simulation.

Our system can be extended, and below, we discuss some remaining research challenges.

We focus on faults resulting from misbehaving but non-malicious nodes. What if the faults are because of malicious attacks? These are generally hard to detect as they can be disguised as benign faults. It would be interesting to study how security mechanisms (e.g., cryptographic schemes for authentication and integrity) and counter-measures such as secure traceroute [34] can be incorporated into our system.

Currently, our system works with a fairly complete knowledge of the RF condition, traffic statistics, and link performance. Obtaining such complete information is sometimes difficult. It would be useful to investigate techniques that can work with incomplete data, i.e. data obtained from a subset of the network. This would improve the scalability of the troubleshooting system.

Finally, there is room for improvement in our core system as well. Our system depends on the accuracy and efficiency of the simulator, the quality of the trace data, and the fault search algorithm. Improvement in any one of these will result in better diagnosis. For example, our system could benefit from fast network simulation techniques developed by [20, 23]. Further, Bayesian inference techniques could be useful for diagnosing faults that exhibit similar faulty behavior.

We are continuing our research and are in the process of enhancing the system to take corrective actions once the faults have been detected and diagnosed. We are also extending our implementation to manage a 50 node multihop wireless testbed. We intend to evaluate its performance when some of these nodes are mobile.

9. RELATED WORK

Many researchers have worked on problems that are related to network management in wireless networks. We broadly classify their work into three areas: (1) protocols for network management; (2) mechanisms for detecting and correcting routing and MAC misbehavior, and (3) general fault management.

In the area of network management protocols, Chen *et al.* [15] present Ad Hoc Network Management Protocol (ANMP), which uses hierarchical clustering to reduce the number of message exchanges between the manager and agents. Shen et al. [41] describe a distributed network management architecture with SNMP agents residing on every node. Our work differs from these two pieces of work in that we do not focus on the protocol for distributing management information, but instead on algorithms for identifying and diagnosing faults. Consequently, our work is complimentary to both [15] and [41].

A seminal piece of work in the area of detecting routing misbehavior is by Marti, Giullu, Lai, and Baker [27]. The authors address network unreliability problems stemming from selfish intent of individual nodes, and propose a *watchdog* and *pathrater* agent framework for mitigating routing misbehavior and improving reliability. The basic idea is to have a watchdog node observe its neighbor and determine whether it is forwarding traffic as expected. The pathrater assigns ratings for paths based on observed node behavior. Based on the rating, nodes establish routes that avoid malicious nodes resulting in overall increase in throughput. Other work following this thread of research includes [9, 12, 13].

Our work differs from watchdog-like mechanism in the following ways. First, the focus of the above research has been on detecting and punishing malicious nodes, whereas our focus is on detecting and diagnosing general performance problems. Second, we use reports from multiple neighbors and take historical evidence into account to derive more accurate link loss rates. Finally, most importantly, we use online simulation-based diagnosis to determine the root cause for high link loss rates. Different from a simple watch-dog mechanism, which considers end points as misbehaving when its adjacent links incur high loss rate, our diagnostic methodology takes into account current network configuration and traffic patterns to determine if the observed high loss rates are expected, and determines the root causes for the loss (e.g., whether it is due to RF interference, or congestion, or misbehaving nodes).

In the area of wireless network fault management, there exist a number of commercial products in the market. Examples include AirWave [8], AirDefense [6], Computer Associate's UniCenter [1], Symbol's Wireless Network Management System (WNMS) [42], IBM's Wireless Security Auditor (WSA) [46], and Wibhu's SpectraMon [44]. Recently, Adya et al. [4] present architecture and techniques for diagnosing faults in IEEE 802.11 infrastructure networks. Our work differs from these in that the above work target infrastructure or base station based wireless networks. Multihop wireless networks are significantly different.

10. CONCLUSION

Troubleshooting a multihop wireless network is challenging due to the unpredictable physical medium, the distributed nature of the network, the complex interactions between various protocols, environmental factors, and potentially multiple faults. To address these challenges, we propose online trace-driven simulation as a troubleshooting tool.

We evaluate our system in different scenarios and show that it can detect and diagnose over 10 simultaneous faults of different types in a 25-node multihop wireless network. This result suggests that our approach is promising. An important property of our system is that it is flexible and can be extended to diagnose additional faults. We hope that this paper will inspire other researchers to further investigate trace-driven simulation as a tool to diagnose and manage complex wireless and wireline networks.

11. REFERENCES

[1] The future of wireless enterprise management. http://www3.ca.com/.

[2] Promise of intelligent networks. http://news.bbc.co.uk/2/hi/technology/2787953.stm.

[3] NSF workshop on residential boradband revisited: Research challenges in residential networks, boradband access and

applications. http://cairo.cs.uiuc.edu/nsfbroadband/, October 2003.

[4] A. Adya, P. Bahl, R. Chandra, and L. Qiu. Architecture and techniques for diagnosing faults in IEEE 802.11 infrastructure networks. In *In Proc. of ACM MOBICOM*, Sept. 2004.

[5] D. Aguayo, J. Bicket, S. Biswas, G. Judd, and R. Morris. Link-level measurements from an 802.11b mesh network. In *In Proc. of ACM SIGCOMM*, Aug. 2004.

[6] AirDefense: Wireless LAN security and operational support. http://www.airdefense.net/.

[7] Wildpackets Airopeek. http://www.wildpackets.com/products/airopeek.

[8] Airwave, a wireless network management solution. http://www.airwave.com/.

[9] B. Awerbuch, D. Holmer, and H. Rubens. Provably secure competitive routing against proactive Byzantine adversaries via reinforcement learning. In *JHU Tech Report Version 1*, May 2003.

[10] G. Berger-Sabbatel, F. Rousseau, M. Heusse, and A. Duda. Performance anomaly of 802.11b. In *Proc. of IEEE INFOCOM*, June 2003.

[11] P. Bhagwat, B. Raman, and D. Sanghi. Turning 802.11 inside-out. In *Workshop on Hot Topics in Networks (HotNets-II)*, November 2003.

[12] S. Buchegger and J. Y. Le Boudec. Nodes bearing grudges: Towards routing security, fairness, and robustness in mobile ad hoc networks. In *Proceedings of the Tenth Euromicro Workshop on Parallel, Distributed and network-based Processing*, pages 403 – 410. IEEE Computer Scoiety, January 2002.

[13] S. Buchegger and J.-Y. Le Boudec. The effect of rumor spreading in reputation systems for mobile ad-hoc networks. In *Proceedings of Modeling and Optimization in Mobile, Ad Hoc and Wireless Networks*, Sophia-Antipolis, France, March 2003.

[14] J. Case, M. Fedor, M. Schoffstall, and J. Darvin. A simple network management protocol (SNMP). In *Internet Engineering Task Force, RFC 1098*, May 1990.

[15] W. Chen, N. Jain, and S. Singh. ANMP: Ad hoc network management protocol. In *IEEE Journal on Selected Areas in Communications*, volume 17 (8), August 1999.

[16] D. D. Clark, C. Patridge, J. C. Ramming, and J. T. Wroclawski. A knoweldge plane for the internet. In *Proceedings of ACM SIGCOMM*, Karlsruhe, Germany, August 2003.

[17] Wireless networking reference—community wireless/rooftop systems. http://www.practicallynetworked.com/.

[18] D. D. Couto, D. Aguayo, J. Bicket, and R. Morris. A high-throughput path metric for multi-hop wireless routing. In *Proc. of ACM MOBICOM*, Sept. 2003.

[19] M. Garey and D. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W.H. Freeman and Company, 1979.

[20] Y. Gu, Y. Liu, and D. Towsley. On integrating fluid models with packet simulation. In *IEEE INFOCOM*, Mar. 2004.

[21] Y. Hu, A. Perrig, and D. B. Johnson. Wormhole Detection in Wireless Ad Hoc Networks. Technical Report TR01-384, Rice University, Computer Science, Houston, TX, December 2001.

[22] D. B. Johnson, D. A. Maltz, and J. Broch. DSR: The dynamic source routing protocol for multihop wireless ad hoc networks. In *Ad Hoc Networking*, 2001.

[23] H. Kim and J. C. Hou. A fast simulation framework for ieee 802.11-operated wireless lans. In *ACM SIGMETRICS*, Jun. 2004.

[24] A. V. Konstantinou, D. Florissi, and Y. Yemini. Towards self-configuring networks. In *DARPA Active Networks Conference and Exposition (DANCE '02)*, San Francisco, CA, May 2002.

[25] D. Kotz, C. Newport, and C. Elliott. The mistaken axioms of wireless-network research. Technical Report TR2003-467, Dartmouth College, Computer Science, Hanover, NH, July 2003.

[26] P. Kyasanur and N. Vaidya. Detection and handling of MAC layer misbehavior in wireless networks. In *Dependable Computing and Communications Symposium (DCC) at the International Conference on Dependable Systems and Networks (DSN)*, pages 173–182, San Francisco, California, June 2003.

[27] S. Marti, T. Giuli, K. Lai, and M. Baker. Mitigating routing misbehavior in mobile ad hoc networks. In *Proceedings of ACM MOBICOM*, Boston, MA, August 2000.

[28] Native 802.11 framework for IEEE 802.11 networks. Windows Platform Design Notes, March 2003. http://www.microsoft.com/whdc/hwdev/tech/network/802x/Native80211.r

[29] Network devices and protocols: Windows DDK. NDIS library functions.

[30] The network simulator – ns-2. http://www.isi.edu/nsnam/ns/.

[31] L. M. S. C. of the IEEE Computer Society. Wireless LAN medium access control (MAC) and physical layer (PHY) specifications. *IEEE Standard 802.11*, 1999.

[32] OPNET modeler. http://www.opnet.com.

[33] J. Padhye, R. Draves, and B. Zill. Routing in multi-radio, multi-hop wireless mesh networks. In *In Proc. of ACM MOBICOM*, Sept. 2004.

[34] V. N. Padmanabhan and D. R. Simon. Secure traceroute to detect faulty or malicious routing. In *ACM SIGCOMM Workshop on Hot Topic in Networks (HotNets-I)*, Oct. 2002.

[35] N. Pissinou, B. Bharghavati, and K. Makki. Mobile agents to automate fault management in wireless and mobile networks. In *Proceedings of the IEEE International Parallel and Distributed Processing Systems Workshop*, pages 1296–1300, 2000.

[36] R. Prakash and M. Singhal. Low cost checkpointing and failure recovery in mobile computing systems. In *IEEE Trans. on Parallel and Distributed Systems*, volume 7 (10), pages 1035–1048, 1996.

[37] The Qualnet simulator from Scalable Networks Inc. http://www.scalable-networks.com/.

[38] MIT Roofnet. http://www.pdos.lcs.mit.edu/roofnet/.

[39] M. Sabin, R. D. Russell, and E. C. Freuder. Generating diagnositc tools for network fault management. In *Integrated Network Management*, pages 700–711, 1997.

[40] G. K. Saha. Transient fault-tolerant mobile agent in mobile computing. In *IEEE Transactions on Computers*, USA, 2003. IEEE Computer Society Press.

[41] C.-C. Shen, C. Jaikaeo, C. Srisathapornphat, and Z. Huang. The Guerrilla management architecture for ad hoc networks. In *Proc. of IEEE MILCOM*, Anaheim, California, October 2002.

[42] SpectrumSoft: Wireless network management system, Symbol Technolgies Inc. http://www.symbol.com/.

[43] Tcpdump. http://www.tcpdump.org/.

[44] SpectraMon, Wibhu Technologies Inc. http://www.wibhu.com/.

[45] Wireless research API. http://ramp.ucsd.edu/pawn/wrapi/.

[46] Wireless security auditor (WSA). http://www.research.ibm.com/gsal/wsa/.