

Spectrum MRI: Towards Diagnosis of Multi-Radio Interference in the Unlicensed Band

Akash Baid¹, Suhas Mathur^{*2}, Ivan Seskar¹, Sanjoy Paul³, Amitabha Das³, Dipankar Raychaudhuri¹

¹WINLAB, Rutgers University, {baid, seskar, ray}@winlab.rutgers.edu

²AT&T, suhas@att.com

³Infosys Technologies Ltd., {sanjoy_paul, amitabha_das}@infosys.com

Abstract—The increasing density and data rate of unlicensed band wireless devices in small office and home (SOHO) environments has led to significant inter- and intra-radio interference problems. Multiple competing standards such as the IEEE 802.11b/g, Bluetooth and ZigBee, all of which operate in the 2.4 GHz ISM band, can interfere with each other when used in typical indoor environments, potentially causing significant performance degradation. This paper presents detailed experimental results (using the ORBIT radio grid testbed) to quantify the effects of such interference in representative SOHO scenarios. In particular, different topologies, traffic loads and number of interfering devices are emulated to show the impact of multi-radio interference and to characterize each kind of interference. Further, a cross-layer, multi-radio interference diagnosis framework (called “spectrum MRI”) is described with the aim of isolating and classifying multi-radio interference problems using heuristic and model-based methods. A specific example of identifying interference problems which may affect an 802.11g video link is given to illustrate the proposed measurement and diagnosis framework.

I. INTRODUCTION

The evolution of wireless protocols and access technologies for the unlicensed bands has led to the rapid proliferation of consumer grade wireless devices that do not require spectrum configuration by end users. The typical digital home environment is increasingly moving towards dense deployment of multiple wireless devices using a variety of unlicensed band radio standards. Unfortunately, this has also meant that the unlicensed band is becoming interference limited, and in many cases, overcrowded with multiple radio access technologies competing for common spectrum. For example, the popular 802.11 standard, the Bluetooth standard and the ZigBee standard, all share the same chunk of radio spectrum, as shown in Figure 1, in addition to emitters such as cordless phones and leaking microwave ovens, also in the same band. As we show in the subsequent sections, uncoordinated sharing of unlicensed spectrum leads to significant interference related performance degradations. In particular, we study the multi-radio interference problem in detail in this paper, focusing on the performance loss under various scenarios typical in home environments, and we put forward the thesis that in many cases

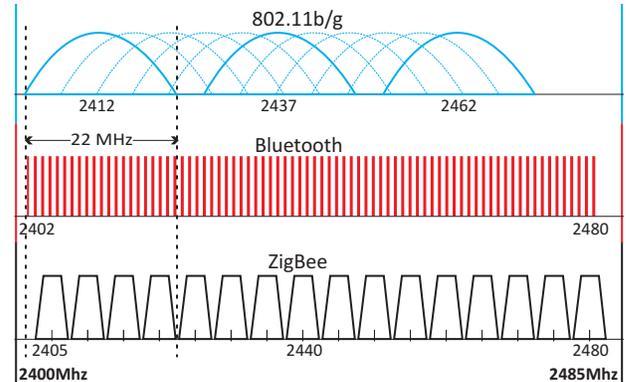


Fig. 1. 802.11, Bluetooth and ZigBee Channels in the 2.4 GHz ISM Band

it is possible to diagnose multi-radio interference problems by passive observation of symptoms that are produced as artifacts of the interference.

While there are both commercial products and recent studies around the interference problem in the 2.4 GHz spectrum (see [1] for a detailed survey), most of the work has been focussed on troubleshooting WiFi problems in large campus or enterprise environments. In comparison, our work differs on two counts - (a) We focus on small office and home environments which leads to different interference problems and different solution requirements compared to a large scale system, and (b) Rather than concentrating only on WiFi problems, our aim is to diagnose multi-radio interference since in a home environment, a user-owned Bluetooth or ZigBee device might be equally or even more important than a WiFi device. To this end, our system utilizes one or more *monitors* that capture the ongoing multi-radio transmissions passively and aggregate their observations into a database. From these combined traces, we can use a heuristic or learning algorithm which identifies interference problems and if possible, recommends configuration changes in one or more devices. Such a low cost monitoring and diagnosis system for multi-radio interference is intended to improve the performance of home networks, which are usually operated by non-expert users. In this paper, we first provide some qualitative and quantitative multi-radio interference examples typical in the SOHO environment and subsequently describe our framework and methodology for interference diagnosis and classification

The research reported in this paper was supported in part by a grant from Infosys Technologies Ltd. *This work was performed while Suhas Mathur was at WINLAB, Rutgers University.

in a multi-radio environment.

II. BACKGROUND AND RELATED WORK

As radio standards in the unlicensed bands have evolved over the years, a number of studies on inter-radio interference have been conducted typically following an approach that seeks to characterize one standard versus another. For example [2] and [3] analyze the impact of Bluetooth on 802.11b/g and suggest some techniques to improve co-existence between these two standards. Similarly, [4] provides a detailed analytical model for interactions between ZigBee and 802.11 and between ZigBee and Bluetooth. In a complex multi-radio environment having simultaneous interactions of several competing wireless standards, the approach of modeling and analysis becomes much harder.

A related area of research focuses on the more fundamental causes and effects of interference on PHY and MAC layer performances with the aim of designing techniques to overcome the problems involved (for example [5]). The authors in [6] go a step further to derive closed-form throughput expressions by creating an analytical framework for interactions between heterogeneous radios using such physical layer models. In contrast to previous studies on the nature and modeling of heterogeneous radio interference, our work focuses on the diagnosis aspect.

Due to the popularity of the 802.11 WLAN standard, most work in the network diagnosis and management domain has focussed on solving issues within this standard. The authors in [7] for example, provide details on an elaborate cross-layer trace collection and analysis system to address issues ranging from configuration problems to interference related problems. Similarly, a systematic approach in [8] focuses on the framework for collecting and analyzing traces for 802.11 networks. Some other approaches for 802.11 WLAN diagnosis include a structural and behavioral model based system [9], distributed physical layer anomaly detection [10] and fault diagnosis using signal error rate and RSSI parameters [11]. In this work, we take a more generic view of the network in terms of multiple standards and devices, and introduce an appropriate framework for multi-radio interference diagnosis.

III. MULTI-RADIO INTERFERENCE EXAMPLES

In this section, we present examples of some typical multi-radio interference problems in home networks. We classify the interference measurement experiments into the following categories:

- Intra 802.11 Interference
- Inter-radio Interference
 - 802.11-Bluetooth Interference
 - 802.11-ZigBee Interference
 - Bluetooth-ZigBee Interference
- Complex Multi-radio Interference

Emulation Methodology: All the experimental evaluations

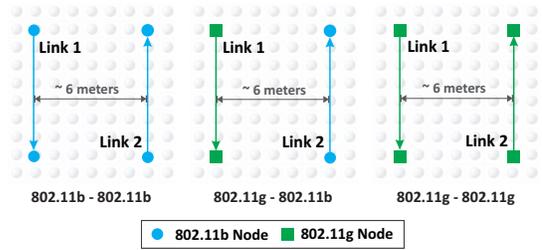


Fig. 2. Topology showing the effect of a co-channel slow link

Configuration	Link	1 Mbps	5.5 Mbps	11 Mbps	0 Mbps
11b-11b	L1	1.19	2.89	3.76	5.70
	L2	0.63	1.92	2.87	-
11g-11b	L1	17.60	22.32	25.10	31.80
	L2	0.36	0.78	1.74	-
11g-11g	L1	9.67	14.70	17.91	31.80
	L2	0.62	2.42	4.66	-

TABLE I
THROUGHPUT (MBPS) OF LINK 1 & LINK 2 UNDER DIFFERENT RATE OPTIONS FOR LINK 2. LINK 1 RATE IS SET TO MAXIMUM.

described in this section were conducted on the ORBIT testbed [12] which consists of 400 small form-factor PCs placed in a 20 x 20 regular grid with an inter-node separation of about 3ft, spanning a total area of 3600 sq. ft. Each of these nodes is equipped with two IEEE 802.11a/b/g wireless interfaces, with 40 nodes also equipped with Bluetooth dongles and 30 nodes also equipped with 802.15.4 TelosB motes. The Iperf tool is used for throughput measurement of TCP and UDP data in both 802.11 and Bluetooth radios and a customized ZigBee traffic generator built on the TinyOS platform is used for performance measurement on the ZigBee nodes. The throughput measurements in each of the experiments described in this section were averaged over ten or more readings spread in time and location inside the ORBIT grid to remove random effects of environmental changes and device specific variance. Unless otherwise mentioned, all the 802.11 nodes in our experiments were operated on Channel 1 which did not have any external interference as confirmed by a spectrum analyzer. In the following subsections, we identify and quantitatively analyze some commonly occurring multi-radio interference problems:

A. Slow Co-channel link in 802.11

When two 802.11 b/g links co-exist on the same channel, the slower link has a higher channel occupancy time causing the high rate link to undergo more backoffs and thus suffer a large drop in throughput. To quantify this effect, three cases of single link interference are emulated as shown in Figure 2. In all the three cases, the data-rate of Link 1 is set to the highest (11 Mbps for 802.11b and 54 Mbps for 802.11g) while the rate of Link 2 is changed in steps. All links in this experiment carry saturation TCP traffic with a buffer size of 8 KBytes and each reading is averaged over ten trials of 100 second duration. From Table I, we observe a substantial drop in Link 1 throughput when the interfering link (Link 2) data-rate drops down from 11 Mbps to 1 Mbps in all the three cases. This drop is about 32% in case of 802.11b-802.11b interference

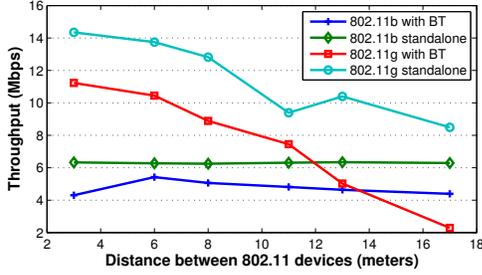


Fig. 3. Throughput(Mbps) of 802.11 link at varying distances with co-located Bluetooth transmitter

and 53% in case of 802.11g-802.11g interference.

From a practical point of view, this scenario is very common and can present itself in a number of ways: for example, if an old laptop with a slow 1 Mbps 802.11b radio link is connected to the AP, or if some sort of local interference triggers an automatic rate reduction scheme on one of the links.

B. 802.11-Bluetooth Interference

One of the most common types of inter-radio interference occurs between 802.11b/g and Bluetooth radios. The most severe 802.11-Bluetooth interference is observed in the co-located case where the Bluetooth and 802.11b/g radios are located on the same physical device such as smart phones and laptops. At such distances, a transmission on any of the roughly 22x1 MHz Bluetooth channels that overlap with a 802.11 channel will cause a packet error for the 802.11 transmission. The following two experiments exemplify some of the problems in the 802.11-Bluetooth interaction scenario:

1) *Co-located 802.11b/g and Bluetooth:* Dual radio nodes were used to emulate a co-located case in which the distance between the 802.11b/g and Bluetooth radios is about 25cm and the transmit power is 18dBm and 4dBm(Class 2 device) respectively. To create a worst-case interference scenario in this topology, the Bluetooth transmitter and the co-located 802.11b/g receiver operate concurrently. The 802.11 transmitter is located at varying distance which varies the received power levels at the receiver. TCP traffic is pushed through the 802.11b/g link with the rate set at 11 Mbps for 802.11b and 24 Mbps for 802.11g. The Bluetooth interferer carries a 512 kbps UDP load with a datagram size of 1 KB. From the observed throughput numbers in Figure 3 we can see that the impact of Bluetooth is greater for an 802.11g link with a steep drop with distance. We observe that when the 802.11g transmitter and receiver are separated by a distance of 15 meters or more, the 24 Mbps link throughput can drop to less than 3 Mbps.

2) *Effect of Autorate on Bluetooth-802.11 Interference:* Another interesting issue here is the behavior of the 802.11 autorate selection algorithm in presence of Bluetooth interference. To study this effect, we measure the throughput (using TCP traffic at 11 Mbps) of a 802.11b link with and without autorate enabled in the three configurations shown

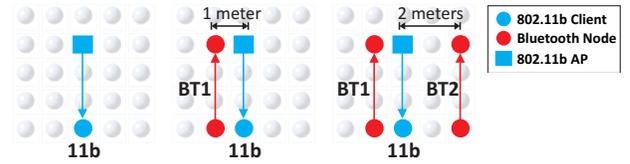


Fig. 4. Topology for autorate effect

Topology	Set Rate = 11Mbps			Autorate Enabled		
	11b	BT1	BT2	11b	BT1	BT2
Only 11b	5.40	-	-	4.92	-	-
11b, BT1	3.95	0.25	-	1.92	0.25	-
11b, BT1, BT2	2.81	0.24	0.63	0.42	0.30	0.69

TABLE II
LINK THROUGHPUT (MBPS) WITH AND WITHOUT AUTORATE ENABLED FOR DIFFERENT TOPOLOGIES

in Figure 4. While a number of WLAN automatic rate fallback algorithms have been proposed and tested, for our demonstration purpose we use the autorate feature in the MadWifi driver, which implements the Onoe bit-rate selection algorithm. From the throughput numbers shown in Table II we can clearly see that rate reduction in such a scenario causes longer collision windows and thus lower throughput in 802.11b. With autorate option enabled in most laptop 802.11b/g cards, this presents a very common example of a configuration problem that causes loss in throughput.

C. Complex Multi-radio Interference Environment

As a final example of multi-radio interference, we emulate a complex small office environment consisting of multiple 802.11b, 802.11g, Bluetooth and ZigBee radios distributed throughout the premises. As shown in Figure 5 an 802.11b node on channel 1 forms the main access point to which four clients are connected. Two additional 802.11g links on channels 1 and 11 respectively support point to point devices such as set-top box to TV or projector. A fifth 802.11b link on channel 11 emulates point to point file transfer in this scenario. As is common in commercial premises, there are neighboring APs within the interference region of the environment, here depicted by an 802.11g AP on the top-left with two clients. Some three Bluetooth constant bit rate transmissions add to the radio clutter with one of the Bluetooth nodes being co-located with the 802.11g node. A number of low cost ZigBee sensors form a part of the security/temperature control infrastructure and relay periodic readings to a central ZigBee concentrator in this scenario.

This topology was studied under varying traffic conditions and an example configuration is mentioned in Table III. The last column from the table shows that when all the links are active, a drop of more than 50% of the nominal throughput is observed for almost all links in the network. In this particular case, for example, throughput of link B1 drops by about 93%, while that of BT1 and all the ZigBee links drops by 89% and 97% respectively making them extremely problematic from a user's perspective.

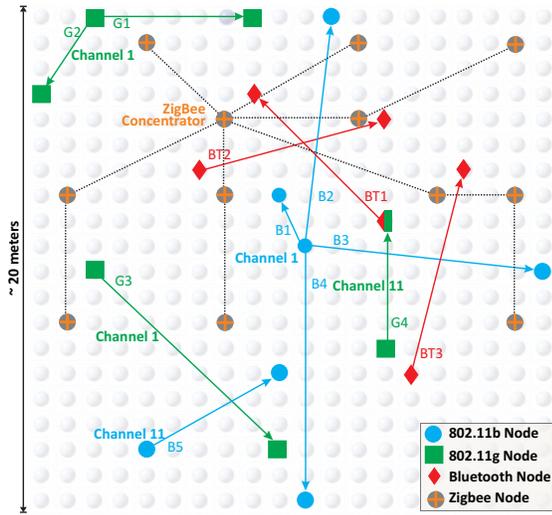


Fig. 5. Multi-radio Interference Topology

Links	Traffic Configuration	Throughput as % of nominal
B1, B2, B3, B4	Saturation TCP 11 Mbps data rate	6.9, 16.4 14.2, 21.5
B5	CBR UDP 10 Mbps offered load	66.9
G1, G2, G3, G4	Saturation TCP 54 Mbps data rate	50.9, 43.9 24.3, 25.9
BT1, BT2, BT3	CBR UDP 512 kbps offered load	11.9, 24.9 39.8
ZigBee Nodes	Periodic 50 Bytes at 250 kbps	3.1

TABLE III

TRAFFIC CONFIGURATIONS OF THE LINKS USED IN THE TOPOLOGY OF FIGURE 5

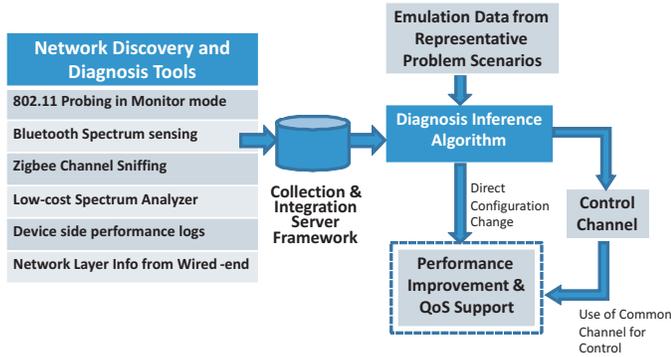


Fig. 6. System Level Model

IV. INTERFERENCE DIAGNOSIS SYSTEM

Figure 6 shows the overall structure of our diagnosis system which consists of a set of network monitoring tools which log traces to a central database based on which the diagnosis heuristics identifies and classifies different interference problems. Although each different kind of monitor shows a different view of the wireless environment producing output traces in different formats, the aim of our system is to converge the traces into a common sqlite database format to get full advantage of correlation algorithms that run across multiple

kinds of traces. Traces are collected and synchronized using the Orbit Measurement Library (OML) [13] framework which is built on a client-server architecture ideal for such modular additions. Some details on each of the monitoring tools are as follows:

- **802.11 Probing:** We use a modified form of the tcpdump tool using the libpcap packet capture library to log the headers of all packets received on the WiFi interface card. Multiple logs of the same packet received by different monitor nodes are purged and synchronized keeping only the timestamp and RSSI from these repeat traces for localization information.
- **Bluetooth Spectrum Sensing:** Since the frequency hopping scheme employed in Bluetooth transmissions make it hard to passively monitor, we employ a spectrum sensing technique as described in [14] to estimate of the number of active Bluetooth transmissions and their traffic load.
- **ZigBee Channel Sniffing:** We employ a passive frequency-hopping listener application built upon the TinyOS platform to log and aggregate the packets received over the ZigBee interface. These are then ported to the OML server keeping the tables in sync with other monitors.
- **Spectrum Analyzer:** A low-cost coarse resolution spectrum analyzer like [15] can provide a means to detect other sources of RF emissions, for example that from cordless phones, leaking microwave ovens, etc. Since these devices are commonplace in a SOHO environment, we employ this additional monitor to diagnose such interference problems.
- **Device-side Logs:** Interference being a receiver side phenomenon, can be best detected from the user device involved in the transmission. As such, we have an optional device side logging mechanism which records the changes in throughput, delay and received power.
- **Wired-end Information:** With some prior knowledge about the devices, the wired-side logs from the AP, for example, can provide information about the logical topology of the system which helps the diagnosis algorithm to narrow down on the interfering links.

V. DIAGNOSIS EXAMPLE: INTERFERENCE TO HD VIDEO

In this section, we provide a detailed example of how the proposed data-centric approach can be used to diagnose possible interference related problems in a specific real world application - that of high-rate video transmissions. Due to its high bandwidth and low delay requirements, Wireless HD streaming presents itself as an important problem in the SOHO environment that our work is focussed on. As an initial case study, we define four possible interferences to a video stream in a home environment and subsequently elaborate on the key parameters that can be used to characterize and diagnose each of these problems.

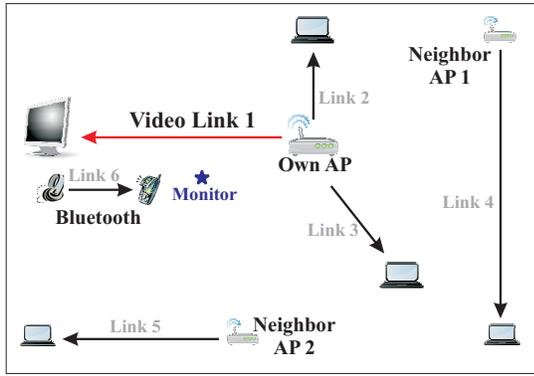


Fig. 7. Representative topology for Video Streaming Example

A. Topology Setup

To emulate a small home network, we ran our experiments based on the representative topology shown in Figure 7. In Figure 7, link 1 is the main video stream, with the distance between the transmitter and receiver being close to 10 meters. The VLC application is used to emulate the video link and a CBR 15 Mbps video at a transmit rate of 54 Mbps is used as the stream.

Links 2 and 3 represent other user devices connected to the same AP and both of them generate approximately 1 Mbps TCP data traffic at 54 Mbps link speed with constant inter-departure time and Pareto-distributed packet length with mean 450 bytes. Links 4 and 5 are independent AP-client pairs that emulate neighboring links outside the home, quite commonly found in such a setting. These links carry 2 Mbps data traffic at 54 Mbps each with the same characteristics as that of links 2 and 3. There is also a Bluetooth pair close to the video receiver which carries a CBR UDP 512 kbps on-off data stream. In the rest of this section, we work with the following four types of interferences to the video link under a reasonable assumption that only one of the four problems dominate the video quality at a given point of time:

- Bluetooth Interference
- Slow link on the same AP
- Slow link on a neighboring AP
- Channel Congestion

While this presents a small selection from the wide variety of interferers present in a home environment, all of the listed problems are very common for home networks and in this case study, we show how a passive observation of the various link transmissions can be used to filter out signatures that can be ascribed to each of these problems.

B. Classification Parameters

As the case with any other kind of diagnostics, in order to say ‘Problem X occurred due to reason Y’, we need to first identify the parameters that show the ‘symptoms’ of the problem and then classify the multi-dimensional parameter space with appropriate hyperplanes. We found that the two most important parameters for the problems mentioned above are the per-link percentage-retransmission and the occupancy.

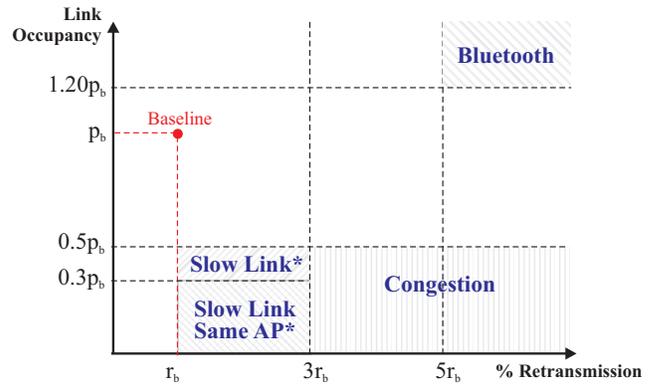


Fig. 8. Interference Diagnosis Regions on the Occupancy vs Retransmission plot. (*Requires Secondary Test)

Here we define link occupancy as the fraction of time for which this link was transmitting data and includes both primary and the retransmissions. These two parameters are aggregated using a 200ms sliding window to avoid unrelated artifacts and reduce storage requirement. The classification of the four problem scenarios, based on these two parameters are shown in Figure 8 which is explained later in this section. To have an idea of the parameter values for a healthy case of nominal traffic on all interfering 802.11 links and no Bluetooth interference, the baseline operating point is marked in the figure which shows 2.8% retransmissions and link occupancy ratio of about 0.24. In addition to this baseline traffic, the effect of each of the four problems show up in slightly different ways as follows:

a) *Bluetooth Interference*: : Since Bluetooth only corrupts some packets at random, the transmitter still sees the channel as unoccupied as before and tries to counter the errors by retransmissions which in turn increases the link occupancy. Thus in this case, there is an increase in both the occupancy as well as the percentage retransmission. In our experiments, we create this case by turning on the Bluetooth link along with the video link. Other 802.11 links (2,3,4 and 5) were also turned on with their nominal traffic flow.

b) *Channel Congestion*: : When there are a number of high-speed 802.11 links in the same contention region, each one gets only a fraction of the channel occupancy and also the number of collision induced errors increases. Hence this gives rise to a large drop in occupancy accompanied by an increase in the percentage retransmission. We emulate this by turning on all the 802.11 links with 15 Mbps UDP data traffic at 54Mbps link speed.

c) *Slow link on Same AP*: : As seen in section III, a slow link in the network can cause problems for other high-speed links by occupying the shared channel for a large amount of time. This is mostly followed by a slight increase in retransmissions as the number of collisions increase compared to the baseline case. For the emulation, link 2 of Figure 7 is converted to a slow 1 Mbps link while all other links carry their normal traffic.

Interference Type	% ReTx	Occu.	Diagnosis
Bluetooth - Distance 1m	18.0	0.39	✓
Bluetooth - Distance 6m	16.5	0.33	✓
Bluetooth - Distance 11m	13.6	0.24	x
Slow Link - Rate 1Mbps, Dist 1m	4.9	0.07	✓
Slow Link - Rate 1Mbps, Dist 15m	3.2	0.08	✓
Slow Link - Rate 11Mbps, Dist 1m	6.5	0.12	✓
Slow Link - Rate 11Mbps, Dist 15m	7.2	0.17	x
Congestion - 4 links, Traffic 20Mbps	16.4	0.10	✓
Congestion - 4 links, Traffic 5Mbps	10.1	0.17	x
Congestion - 3 links, Traffic 20Mbps	9.8	0.12	✓

TABLE IV
DIAGNOSIS OUTPUT FOR THE PROPOSED HEURISTIC ALGORITHM OVER
DIFFERENT INTERFERENCE CASES

d) *Slow link on Other AP*: : In this case, link 4 is made a slow link by changing its transmission rate to 1 Mbps, while the other links still carry the baseline traffic. Here, we observed that the fall in occupancy is less than that in case of the slow link, same AP case, but the effect on both occupancy and percentage retransmissions is qualitatively the same. This ambiguity between the same AP and other AP cases can be solved by using the information about AP MAC addresses in the diagnosis algorithm.

An important point in the above discussion and in Figure 8 is the value of the thresholds that divide the different regions. The values specified in Figure 8 are heuristic bounds that we observed from multiple trials under different traffic and rate conditions. The thresholds are also chosen relative to the baseline values, Occupancy = p_b and Retransmission percentage = r_b , which makes the thresholds robust for streaming video transmissions of any rate.

C. Experimental Verification

Table IV lists the different configurations that we tested our algorithm on and also shows the average values of the two relevant parameters. The interference configurations are as mentioned in the table, for example in the Bluetooth case, we vary the distance between the video receiver and the Bluetooth pair and Table IV shows the corresponding results for distances 1, 6 and 11 meters. Based on the thresholds, the algorithm correctly diagnoses the presence of Bluetooth interference in the first two cases while it fails in the 11m case due to the relatively less impact of the interference on the main link. Similarly, in case of a slow link interference, when the slow link data rate is set to 11 Mbps and its distance from the main link is 15m, the occupancy level goes above the threshold causing the algorithm to declare a ‘no-problem’ case. The congestion case is emulated using varying number of extra high-speed links and also varying the amount of TCP traffic that each link carries. The traces record an occupancy level of 0.17 in the case of 4 competing links with 5Mbps TCP traffic each and thus do not indicate the existence of a problem. These results are as expected given that the algorithm detects Bluetooth interference, slow links or congestion only when it is strong and thus significant enough to affect the HD video quality.

VI. CONCLUSION AND FUTURE WORK

The experimental results in this paper outline the impact of inter-radio interference and show that in typical topologies, the effect of one radio class on others causes significant reductions in the throughput. In particular, we show some common scenarios like slow link - fast link interference in 802.11, co-located Bluetooth and 802.11 radios, and a dense deployment of both Bluetooth and ZigBee radios, and benchmark the loss in usable throughput in each case. We have also introduced a framework for multi-radio monitoring which provides the basis for diagnosis of interference related problems. The basic building blocks of the system consists of a set of monitors and a database server that collects packet traces to be used for diagnosis. A simple multi-parameter threshold based classification method was shown to diagnose common types of problems quite effectively. Future work is planned on more robust diagnosis methods based on statistical analysis and machine learning techniques.

REFERENCES

- [1] D. Dujovne, T. Turetti, and F. Filali, “A taxonomy of IEEE 802.11 wireless parameters and open source measurement tools,” *Communications Surveys Tutorials, IEEE*, vol. 12, no. 2, pp. 249–262, 2010.
- [2] N. Golmie, N. Chevrollier, and O. Rebala, “Bluetooth and WLAN coexistence: Challenges and solutions,” *IEEE Wireless Communications Magazine*, vol. 10, pp. 22–29, 2003.
- [3] A. Arumugam, A. Doufexi, A. Nix, and P. Fletcher, “An investigation of the coexistence of 802.11g WLAN and high data rate Bluetooth enabled consumer electronic devices in indoor home and office environments,” *IEEE Transactions on Consumer Electronics*, vol. 49, no. 3, pp. 587–596, Aug. 2003.
- [4] S. Y. Shin, H. S. Park, S. Choi, and W. H. Kwon, “Packet error rate analysis of ZigBee under WLAN and Bluetooth interferences,” *IEEE Transactions on Wireless Communications*, vol. 6, no. 8, pp. 2825–2830, August 2007.
- [5] R. Gummadi, D. Wetherall, B. Greenstein, and S. Seshan, “Understanding and mitigating the impact of RF interference on 802.11 networks,” in *Proc. of SIGCOMM*, 2007, pp. 385–396.
- [6] A. Stranne, O. Edfors, and B.-A. Molin, “Energy-based interference analysis of heterogeneous packet radio networks,” *IEEE Transactions on Communications*, vol. 54, no. 4, pp. 761–761, 2006.
- [7] Y.-C. Cheng *et al.*, “Automating cross-layer diagnosis of enterprise wireless networks,” in *Proc. of SIGCOMM*, 2007, pp. 25–36.
- [8] R. Mahajan, M. Rodrig, D. Wetherall, and J. Zahorjan, “Analyzing the MAC-level behavior of wireless networks in the wild,” in *Proc. of SIGCOMM*, 2006, pp. 75–86.
- [9] B. Yan and G. Chen, “Model-based fault diagnosis for IEEE 802.11 wireless LANs,” in *Proc. of MobiQuitous*, Toronto, Canada, Jul. 2009, to appear.
- [10] A. Sheth, C. Doerr, D. Grunwald, R. Han, and D. Sicker, “MOJO: a distributed physical layer anomaly detection system for 802.11 w lans,” in *Proc. of MobiSys*, 2006, pp. 191–204.
- [11] S. Rayanchu, A. Mishra, D. Agrawal, S. Saha, and S. Banerjee, “Diagnosing wireless packet losses in 802.11: Separating collision from weak signal,” in *Proc. of INFOCOM*, April 2008, pp. 735–743.
- [12] D. Raychaudhuri *et al.*, “Overview of the orbit radio grid testbed for evaluation of next-generation wireless network protocols,” in *Proc. of IEEE WCNC*, vol. 3, March 2005, pp. 1664–1669.
- [13] M. Singh, M. Ott, I. Seskar, and P. Kamat, “Orbit measurements framework and library (oml): Motivations, design, implementation, and features,” in *Proc. of TRIDENTCOM*, 2005, pp. 146–152.
- [14] R. Miller, W. Xu, P. Kamat, and W. Trappe, “Service discovery and device identification in cognitive radio networks,” in *Proc. of IEEE SECON*, 2007, pp. 670–677.
- [15] Wi-Spy USB Spectrum Analyzer, <http://www.metageek.net/products/wi-spy>.