IEEE ICNC 2013 - International Conference on Computing, Networking and Communications, San Diego, USA, 28-31 Jan 2013

Low RSSI in WLANs: Impact on Application-Level Performance

Markus Tauber, Saleem N. Bhatti
University of St Andrews, UK
{markus,saleem}@cs.st-andrews.ac.uk

***Abstract*—Widespread use of wireless LAN (WLAN) may soon cause an over-crowding problem in use of the ISM spectrum. One way in which this manifests itself is the low Received Signal Strength Indication (RSSI) at the WLAN stations, impacting performance. Meanwhile, the IEEE 802.11 standard is being evolved and extended, for example with new coding schemes and the 802.11n standard, which makes use of 5GHz and 2.4GHz. We report on measurements of the upper and lower bounds of performance with good and poor RSSI in 802.11g and 802.11n. We find that in operation under poor (low) RSSI, performance is indeed impacted. In some cases the impact is such that there may be little benefit in using the newer 802.11n over the mature 802.11g.**

I. Introduction

IEEE 802.11 Wireless LAN (WLAN) technology is increasingly used to provide connectivity for mobile computing devices and applications [1]. The maturity and widespread deployment of 802.11 infrastructure, and the availability of cheap, integrated chip-sets makes it a popular choice for many devices, including smart-phones, low-cost consumer devices such as net-books and hand-held game consoles or computer peripherals. This means that in many situations, the WLAN spectrum is 'crowded': there is the potential for increased interference of signals from the many end-systems using WLAN in close proximity. This could reduce Received Signal Strength Indication (RSSI) values at a WLAN end-system, and this will impact performance [2], [3]. Although there may be some debate about the efficacy of the impact of RSSI, the current 802.11 standards use RSSI values to select the Modulation Coding Scheme (MCS) for transmission, which in turn determines maximum transmission rates. So, RSSI has a direct impact on the performance of 802.11. IEEE 802.11n was designed to operate in both the 5GHz ISM band and the 2.4GHz ISM band. However, given the popularity of WLAN, it is likely that the 5GHz band may suffer the same 'Tragedy of of the Commons' [4] that has beset the 2.4GHz band. Also, even for 802.11n, 2.4GHz operation remains popular in cheaper devices, such as net-books and lower-end smart-phones, which is likely to compound the over-crowding problem for 2.4GHz. WLAN technology is also increasingly used in non-mobile scenarios, e.g. for providing a cheap back-haul for rural areas [2]. As well as interfering with existing WLAN deployments, such usage itself may be subject to certain environmental conditions that impact RSSI, e.g. tide levels and path lengths in [2], or signal interference in [3].

A. Motivation and Approach

As many variants of 802.11 are likely to operate simultaneously, perhaps in the same geographical area, it is important to consider how such 'crowded' operation may impact on use and deployment of such systems. Performance issues may have a direct influence on: (i) deployment and upgrade planning; (ii) day-to-day monitoring and management; (iii) capacity planning; (iv) systems performance; (v) network configuration.

There is increased uptake of 5GHz 802.11n, in parallel with 2.4GHz 802.11n and the mature 2.4GHz 802.11g usage. Enhancements to 802.11n include higher rate modulation coding schemes (MCSs). The maximum throughput of the MCSs employed in 802.11n can be further increased by using new features [5] such as MIMO, which allows multiple streams via individual antennae (or, if available, multiple NICs). Another feature is the use of packet aggregation techniques to reduce the frame overhead. These features, as well as the choice of the frequency band, are optional.

5GHz may, for now, be less crowded in use than 2.4GHz, but with the increasing number of application and devices using WLAN, 5GHz is likely also to be crowded in the near future. So, we have investigated scenarios in which 802.11g or 802.11n equipment has to operate in an environment with reduced RSSI, and so potentially suffer reduced performance. We are, in particular, interested in performance implications on *out-of-the-box* configurations, as most users may not have the expertise to fine-tune their equipment for optimal performance.

Specific applications may allow, for example, buffering and retransmission of packets, to compensate for impairments (delay/jitter and loss). However, applicability of such adaptation may depend on the specific uses case as well as on the performance requirements of application data flows [6]. The performance of the flows, including energy efficiency, in turn, is determined by packet size and the data (packet) transmission rate for a specific application data flow [7]. By experimenting with packet size and (packet) transmission rate we are able to evaluate the upper and lower bounds of performance – a *performance envelope* – under 'good' and 'poor' RSSI conditions, within which real applications operate. This will show the scope over which management actions or adaptation policies will be effective, narrowing down the solution space as well as providing bounds on the potential benefits [8].

To investigate the impact of poor RSSI, we have attenuated transmission signal strength below levels that might normally be experienced in an office environment. We have compared the resulting performance for a range of traffic patterns specified by (i) the (packet) transmission rate; and (ii) the packet size, in a single WLAN cell with good and poor RSSI for 802.11g (2.4GHz) and 802.11n, (5GHz, 20MHz channels).

We have found that, as expected, 802.11n allows for higher throughput under good conditions (high RSSI), but under poor conditions (low RSSI) the performance gain is negligible compared to 802.11g. In most cases, application-level loss is lower with 802.11g than with 802.11n. This may make 802.11g more suitable for poor RSSI conditions, such as a noisy RF environment, or where larger distances are involved. We have also found that for applications with lower data rates (e.g. individual VoIP or ViIP flows) the difference in performance between 802.11g and 802.11n is negligible. We have studied a number of IEEE 802.11 variants, but due to space constraints we restrict our paper to report only on 11n and 11g as outlined above.

B. Structure of this Paper

We present a summary of related work in Section II. In Section III, we explain our methodology, describe our testbed, and define observables and metrics. In Section IV, we present our results and discussion, concluding in Section V, and finish with a list of future work in Section VI.

II. Related Work

In contrast to our results, most other studies emulate poor RF conditions (e.g. interference) in specific use-cases [9]–[12]. We wish to provide a more generally useful result, and focus on upper and lower bounds represented by poor and good RSSI to define a general *performance envelope* in order to assess use of specific 802.11 variants.

The work most closely related is [12], which provides results of measurements of application-specific performance under IEEE 802.11a/b/g/n. The authors conclude that the benefits of a specific standard, feature (and subsequently the frequency band) depend on the use-case. However, in all considered scenarios, relatively good RSSI values are measured for their specific in-door test-bed setup. No measurements for general upper and lower performance bounds are provided with respect to RSSI, but only a few specific traffic patterns.

Decreasing RSSI and its impact on WLAN performance in a crowded 2.4GHz spectrum due to interference with Bluetooth was measured and reported previously, e.g. in [9].

In [10] performance degradation in 802.11n due to various interferers and levels of attenuation is experimentally evaluated. The authors focus on the former and conclude that theoretical performance gains due to MIMO cannot be reached in practise. Their findings with respect to performance degradation due to increased attenuation is aligned with our findings in the poor RSSI scenario. However, they do not evaluate the upper and lower bounds of operation in either 2.4GHz or 5GHz.

In [11], the authors experimentally evaluate the impact on performance of 802.11n features like MIMO, channel bonding and frame aggregation. They consider a scenario in which the presence of an 802.11g cell causes interference, in a specific office environment and configuration. They report that depending on the location of specific clients in their setup, features like MIMO, had varying benefit with respect to throughput. In [12]–[16] the authors report on empirical measurements of performance in IEEE 802.11 networks. They do not consider situations with poor RSSI: all the RSSI values are above the minimum RSSI required for selecting the fastest Modulation and Coding Scheme (MCS) used in 802.11n, i.e. effectively a 'best case' for RSSI. We explicitly consider RSSI values that are low enough to select slower MCS operation.

III. Experiment Design and Metrics

This paper is part of a body of work carried out in the context of performance- and energy-related topics. Hence, we use the same testbed, experimental harness and, partially, the approach as already described in [7], [8].

As most users do not have the expertise to fine-tune their equipment, we consider that most deployed systems are used in 'out-of-the-box' configurations, without performance tuning. Specifically, our assumptions were:

- *Standard WLAN configuration.* We used only standard, un-tuned WLAN radio-channel configurations. While many WLAN NIC drivers do permit various controls of the hardware, this is not easily accessible or comprehensible for modification by most users. Hence MCS/rate adaptation algorithms (e.g. automatic rate selection) or 802.11n performance enhancements (e.g. MIMO) were used in the standard OS configuration and in the default equipment configuration for our testbed.
- *Packet flow behaviour.* To measure application specific performance (throughput and loss), we used a range of UDP flows specified by packet rate and packet size to represent the upper and lower performance bounds. We also emulate traffic representative of a few popular applications to put our results into context. As our experiment is to examine the behaviour of the WiFi transmission using performance measures like throughput and loss, using TCP directly would modulate the behaviour we observe due to the congestion control and flow control behaviour for TCP. Indeed, as there are various different versions of TCP (e.g. [17]), all with different behaviour and cross interactions (e.g. see [18]), we find that UDP is more suitable for examining performance in a reproducible and unbiased manner. UDP allows us to define upper and lower performance bounds without traffic being 'choked back' by mechanisms like congestion control.
- *Number of Antennas.* As is the case in many low-end, small form-factor and cheaper devices, single antennas are used for 802.11g, and dual antennas are used for 802.11n. So, we have used the same in our testbed.

A. Overview

We have experimentally evaluated packet level performance in our WLAN testbed, with 802.11g and 802.11n, using offthe-shelf equipment. We generated packet flows of various bit-rates and packets sizes, using *iperf*[1] and collected throughput and loss measurements from sequential UDP flows. Our testbed (Fig. 1) consisted of a single client host, a host running a wireless access-point (AP) and an experimental control unit for providing storage for measurement data, *ntp*[2] services and system configuration. The WLAN hosts were setup in a teaching lab in the University of St Andrews with a distance of $\sim 24 \ \pm 0.5$ m between the 2dBi antennae for good RSSI, with poor RSSI emulated using 10dB attenuators. The actual RSSI values were measured at the receiver.

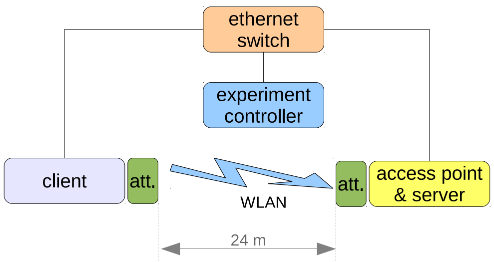

Fig. 1. Schematic of testbed showing physical connectivity. The testbed was configured separately for experiments with 802.11g (2.4GHz) and 802.11n (5GHz, 20MHz channels). The experiment controller uses Ethernet for control messages and shared file-system access. The separation between the antennas of the client and access point/server is 24m, attenuators (att.) and transmission power control were used to adapt the signal strength. Data packets generated by *iperf* were transferred across the WLAN.

B. Workloads: packet flow configuration

The range of workload traffic control varibales is summarised in Table I. These were applied with separate measurements taken for 802.11g and 802.11n.

TABLE I
UDP CONTROL VARIABLES.

Packet size in offered load	64; 1460 bytes
Offered load's bit rate	32; 256; 512 Kbps 1; 5; 10; 20; 30 Mbps

Each packet size was combined with each bit-rate (16 combinations); 20 flows measured with each combination executed for each of 802.11g and 802.11n with 20 MHz (320 flows for each); each flow had a duration of 4 minutes, giving a total of ~43 hours of measurements for each of poor and good RSSI conditions.

The ranges of data rates and packet sizes (Table I) were determined by upper and lower traffic rates possible in *both*, 802.11n and 802.11g in initial experiments. The 64 byte packet is the smallest size for which we have observed that *iperf* is able to generate server reports, and very few applications will have packets smaller than this. The 1460 byte packet is chosen as that is a common TCP maximum segment size to operate over the popular MTU size of 1500 bytes.

Traffic emulating a Skype (VoIP) flow was based on previous studies [19], [20], as was traffic emulating a YouTube (ViIP) flow [21], [22]. We have deduced HTTP-specific downstream traffic profiles from preliminary experiments using *wget*[3] to generate HTTP flows from http://mirror.ox.ac.uk/ for downloading of an Ubuntu ISO CD image file. For each of the above application-specific traffic profiles we have emulated 20 sequential UDP flows with *iperf*. For comparative assessment with VoIP/ViIP flows, we have used a flow duration of 4 minutes, i.e. all workloads have the same duration flows, but will have different packet sizes and (packet) transmission rates.

TABLE II
APPLICATION UDP WORKLOAD EMULATION.

Skype	300 byte packets, 65 Kbps
YouTube	1431 byte packets, 639 Kbps
HTTP	1420 byte packets, 11 Mbps

20 measurements with each flow (60 flows); flow duration of 4 minutes; ×4 gives ∼16 hours of measurements. The emulated workloads are based on [19]–[22], as well as on preliminary measurements.

C. Observed variables

In each experiment we have measured the observables as described below:

- *Performance:* throughput and loss, as recorded by *iperf*'s server reports, on the client for each UDP flow.
- *WLAN rate:* the RSSI value and selected MCS, as recorded periodically at the client, using *iwconfig*[4].

This latter observation – selected MCS and measured RSSI – is important. Currently, 802.11 standards define MCS selection as a function of RSSI, and, therefore, the maximum data rate that is achievable. So, changing RSSI values will have a direct impact on the maximum achievable data rate, regardless of other factors such as packet loss. This is another reason that consideration of RSSI is still important today, even though there is debate amongst the research and technical community about the efficacy of RSSI as a measure of signal quality: RSSI values are used to determine selection of MCS, so RSSI has a direct impact on overall system performance.

D. RSSI control

Table III shows the vendor-defined RSSI-to-MCS mapping of the chip-set that was used in our testbed[5], in which MCS_0 represents the slowest data rate. We have controlled the RSSI by mounting a 10dBm attenuator[6] between the WLAN NIC outputs at the client and the AP, and varied the transmission (TX) power. In preliminary experiments, we found that 10mW (10dBm) is the smallest TX power setting which still allowed packets to be transmitted with 802.11n. This resulted in RSSI values of about -88dBm, which is in contrast to the specification in Table III which suggests packet transfer being possible with lower RSSI values. For 802.11g we have found, in our setup, that a TX level of 3mW (5dBm) at both AP and client resulted in $\sim$ -85dBm and still allows packet transfer. For good conditions we used configurations which resulted in similar RSSI for 11n and 11g and allowed usage of the fastest

[1] https://sourceforge.net/projects/iperf/
[2] http://www.ntp.org/
[3] http://www.gnu.org/software/wget/
[4] http://man.he.net/man8/iwconfig
[5] http://www.compex.com.sg/Datasheets/WLM200NX_DSv3.2.9.pdf
[6] VAT10 attenuators from SSB http://www.ssb.de/

MCS, with a TX power of 50mW (17dBm) for 11n and 3mW (5dBm) for 11g (no attenuators used).

TABLE III
CHIP-SET SPECIFIC RSSI (DBM) TO MCS MAPPING

standard	MCS_0	MCS_1	MCS_2	MCS_3	MCS_4	MCS_5	MCS_6	MCS_7
802.11g	-94	-94	-93	-90	-86	-83	-80	-78
802.11n(20MHz)	-93	-91	-87	-85	-82	-78	-77	-74

RSSI to MCS mapping (all ±2dBm); values for 802.11g with 2.4GHz and 802.11n at 5GHz; MCS_0 is the slowest, MCS_7 is the fastest data rate

E. Equipment

All machines were of identical hardware: a Shuttle X (*XPC Barebone SS56G*[7]) with Intel®Pentium®4 CPU 3.00GHz, 1GB RAM, 112GB HD. Each was equipped with a wireless LAN NIC[8] based on the popular Atheros chipset. We used attenuators from SSB (Germany) which where mounted using jumper cables for connecting to proprietary, reverse-SMA connectors used for the WLAN NICs, with standard SMA connectors at the attenuator.

All machines used Ubuntu 10.04, a minimal server distribution, with the default kernel 2.6.32-24-generic-pae, and the latest WLAN modules (compat-wireless-2011-05-02). For running the AP we have used the *hostapd*[9] package with default parameters. Ubuntu 10.04 contains *hostapd* in version 0.6.9. We have configured channel 40 for 802.11n and channel 6 for 802.11g and used both with the nl80211 driver (an abstraction over the WLAN module mentioned above). The default *hostapd* parameters included a beacon interval of 100ms. To avoid overhead and bias due to link encryption and security mechanisms we disabled encryption and security. To prevent experiments being disturbed by other users, our WLAN cell did not broadcast the SSID in the beacon.

All nodes in the testbed ran in an isolated network. The system clocks of all the nodes where synchronised (using NTP [23]) before each individual experimental run.

IV. RESULTS AND DISCUSSION

A. Overview

In this section the effects on performance of the various experimental conditions are presented in comparison (Fig. 2–5) and in isolation (Fig. 6 and Fig. 7). An overview is presented in Table IV: the values represent the operational limits in the specific scenario. It can be summarised that, as expected, 802.11n allows higher throughput than 802.11g under good conditions. The higher throughput comes at the cost of an increased loss. The throughput improvement of 802.11n over 802.11g under poor conditions is, however, negligible. So, particularly for applications operating at low packet rates, 802.11n offers only very little improvement in comparison to 802.11g.

[7]http://www.shuttle.eu/_archive/old/es/www.shuttle.eu/html/index-416.html
[8]http://www.compex.com.sg/Datasheets/WLM200NX_DSv3.2.9.pdf
[9]http://hostap.epitest.fi/hostapd/
[11]Please note: This high loss rate (17%) was due to outliers (see Fig. 6). The offered load of 30Mbps is less than the reported throughput of 31.5Mbps, but this is within the error of measurement in the use of *iperf*.

TABLE IV
RSSI DEPENDENT UPPER AND LOWER PERFORMANCE BOUNDS

	$throughput_{max}$			$throughput_{max}$		
standard	64 B	1460 B	RSSI	64 B	1460 B	RSSI
802.11g	0.7 Mbps	2.4 Mbps	∼-85dBm	1.7 Mbps	19 Mbps	∼-65 dBm
802.11n	1 Mbps	2.9 Mbps	∼-88dBm	5.8 Mbps	31.5 Mbps[11]	∼-63 dBm
	$loss_{max}$			$loss_{max}$		
	64 B	1460 B	RSSI	64 B	1460 B	RSSI
802.11g	0.4 %	17.2 %[11]	∼-85dBm	0.1 %	0 %	∼-65 dBm
802.11n	0.3 %	0 %	∼-88dBm	29.8 %	0 %	∼-63 dBm

Maximum throughput and loss in all experiments under poor (left columns) and good (right columns) RSSI conditions for 802.11g and 802.11n.

B. Details

We present the application-specific variables – throughput and loss – to analyse the effects of the application-level changes in bit rate and packet size. Firstly, we show the differences in performance between 802.11g and 802.11n in good (Fig. 2) and poor conditions (Fig. 3) as well as the individual standards in both conditions (Fig. 4 and 5).

For analysing the difference between 802.11g and 802.11n, we have computed (Δ throughput), the normalised value of $throughput_{11g}/throughput_{11n}$. As the loss is already a normalised value we have simply computed the difference (Δ loss) of $loss_{11g} - loss_{11n}$. For analysing the difference between poor and good conditions (with both standards, 802.11g and 802.11n) we have computed (Δ throughput) the normalised value of $throughput_{poor}/throughput_{good}$. For analysing differences in loss we have again simply computed the difference (Δ loss) of $loss_{poor} - loss_{good}$. We see that the difference in either good or poor RSSI conditions is determined by the nature of the traffic (i.e. the data rate and packet size).

For raw measurements of each variable, we plot the mean and standard error (with 95% confidence) over 20 runs, per packet size for the offered load at each data rate (Fig. 6 and Fig. 7). In the majority of the experiments, only very small error bars were calculated for throughput and loss, so error bars may not always be visible even though they have been plotted. Also, we have made measurements at discrete values of the control variables, so lines on plots should be considered only as a visual aid, and do not represent an interpolation of results. We also provide information on which MCS was dynamically selected by the driver. As well as indicating the changing configuration of the WLAN operation with changing RSSI, this information also allows comparison with related work by other researchers, where work on performance focuses on the MAC level.

V. CONCLUSIONS

We have experimentally evaluated the upper and lower bounds of performance at the packet level under good and poor RSSI conditions for 802.11g (2.4GHz) and 802.11n (5GHz, 20MHz channels). Depending on their flow characteristics the performance of ‘real applications’ will lie somewhere within these bounds. We have used *out-of-the-box* configurations to represent most common use cases. We found that in an office environment with clear line-of-sight between nodes and

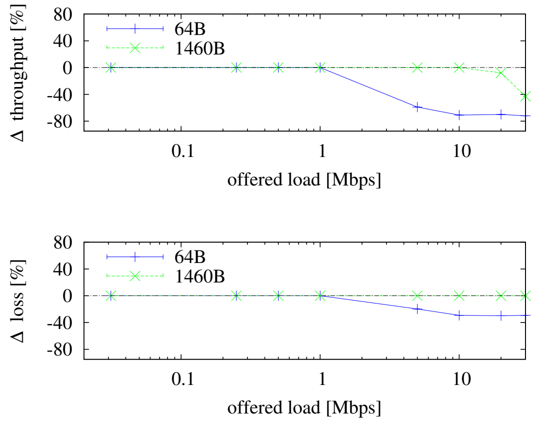

Fig. 2. Differences in throughput and loss of IEEE 802.11g between and IEEE 802.11n with high (∼-65dBm) RSSI. Horizontal zero line is a visual aid. Positive values indicate where 11g has higher values.

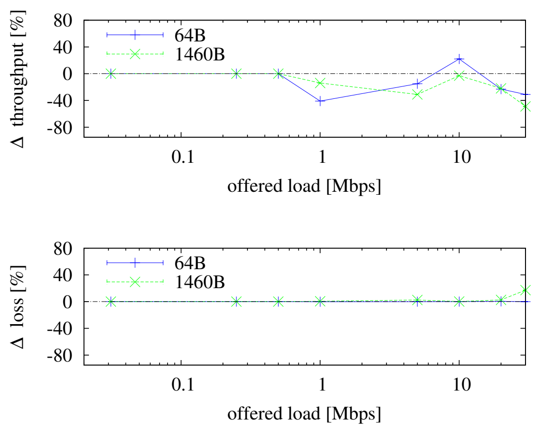

Fig. 3. Differences in throughput and loss of IEEE 802.11g between and IEEE 802.11n with low (∼-85dBm) RSSI. Horizontal zero line is a visual aid. Positive values indicate where 11g has higher values.

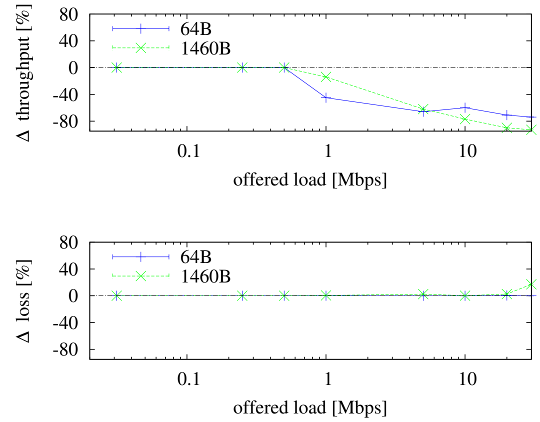

Fig. 4. Differences in throughput and loss of IEEE 802.11g between low (∼-85dBm) and high (∼-65dBm) RSSI. Horizontal zero line is a visual aid. Positive values indicate where low RSSI has higher values.

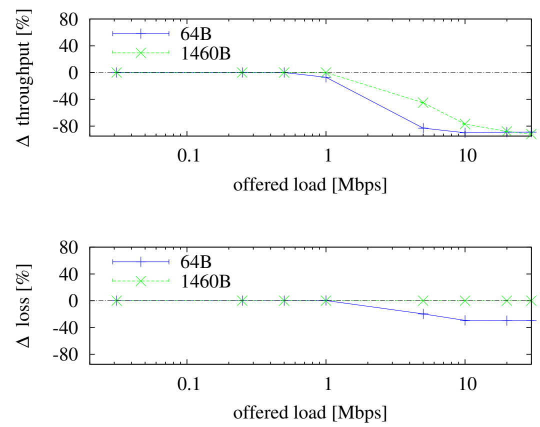

Fig. 5. Differences in throughput and loss of IEEE 802.11n between low (∼-85dBm) and high (∼-65dBm) RSSI. Horizontal zero line is a visual aid. Positive values indicate where low RSSI has higher values.

under good conditions, 802.11n allows higher throughput (up to ∼30Mbps) compared to 802.11g (up to ∼20Mbps), but with higher loss. (To allow comparison between 802.11g and 802.11n, we have constrained our upper limit of offered load to the expected maximum throughput achievable with 802.11g – 30Mbps.) Under poor RSSI conditions, hardly any difference exists up to ∼2.5Mbps between the two 802.11 variants. The lower loss of 802.11g, however, may mean that, under poor RSSI conditions, it is better for loss sensitive applications, e.g. VoIP and ViIP. Meanwhile, 802.11n appears to be more suitable for applications which require high data rates and which are, for example, able to compensate for loss by caching (e.g. streamed video or bulk data transfers can use caching and/or retransmission). We have observed similar results when studying the effect of low RSSI on other variants of IEEE 802.11 (excluded due to space constraints).

Increased use of the 5GHz in future WLAN access scenarios will, quite likely, result in the over-crowding as we now see for the 2.4GHz band. While cost remains a key factor for equipment, 2.4GHz and two antenna will remain a popular configuration. Currently, however, application-specific requirements may need to be considered when deciding which technology should be used in a specific use-case, as planning purchases, deploying, configuring and managing new equipment may not lead to additional benefit. A practical application of our results could be that future deployments may wish to exploit this through a dual/parallel-mode deployment of 802.11g and 802.11n and allocate these to different applications. An alternative solution may be to allow commodity applications to operate in more frequency bands or with spectrum agility, e.g. white space technologies [24].

VI. Future Work

Due to space constraints, we do not report on all available variants of 802.11 and leave this for future work. Also, other scenarios, e.g. with multiple clients, and different traffic loads would be useful to consider. This also includes experimental conditions with various degrees of attenuation and configurations such as various MIMO settings and configurations (e.g. 3 or 4 antenna, rather than the 2 antenna used in this study), auto rate selection variants and the impact of link-layer

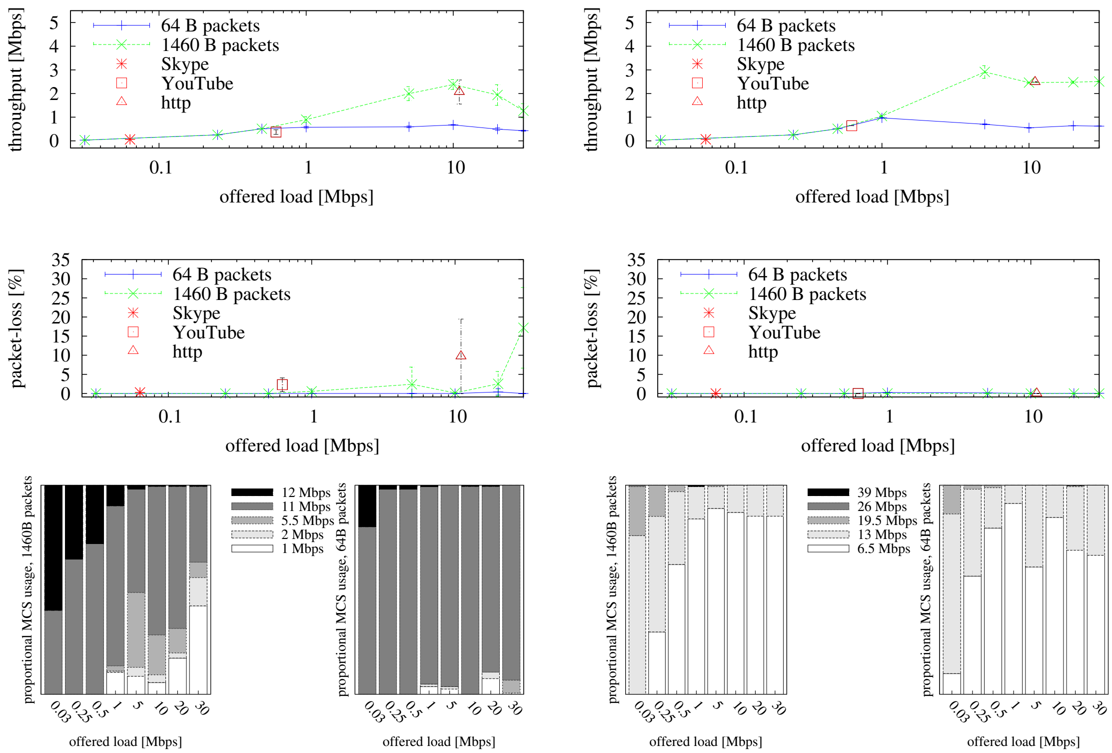

Fig. 6. IEEE 802.11g with avg. RSSI 84.55±0.02 dBm (left column) and 802.11n 20 MHz Channel with avg. RSSI -87.89±0.04 dBm (right column).

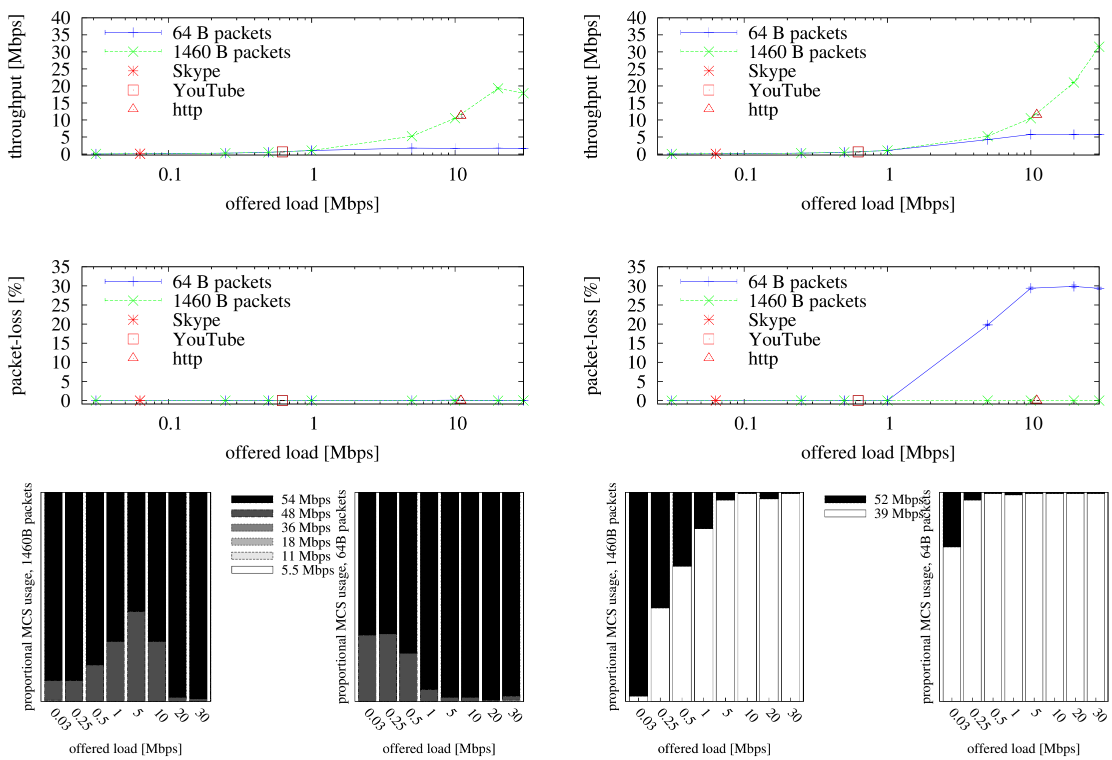

Fig. 7. IEEE 802.11g with avg. RSSI -64.89 ±0.03 dBm (left column) and 802.11n 20 MHz Channel with avg. RSSI -63.37±0.02 dBm (right column).

security. We hope to present analyses of these issues in future studies in order to profile various 802.11 variants in more use cases. Also, from an application point of view, examining the performance of TCP variants [18] and other protocols such as DCCP [25], [26].

This will contribute to a larger body of work in which we have already identified other use cases and scenarios in which the application determines which mode or which 802.11 variant is best suited in the specific situation (see [27]).

Our overall aim is to provide the foundations for the development of protocols and applications which are able to adapt themselves to a changing environment and, for instance, choose a more appropriate operational mode depending on the use case. This is not only true for performance but also for energy efficiency in WLAN for which we have investigated a number of scenarios and discuss dynamic intervention strategies in [7], [8]. This is, again, in support of our overall aim to make applications able to adapt them self in the presence of a changing operational environment.

ACKNOWLEDGEMENTS

This work has been supported mainly by the EPSRC-funded project IU-ATC (http://www.iu-atc.com/).

REFERENCES

[1] Cisco Systems. (2010, February) Cisco visual networking index: Global mobile data traffic forecast update, 2009-2014. [Online]. Available: http://tinyurl.com/cisco-traffic-forecast

[2] G. Bernardi, P. Buneman, and M. K. Marina, "Tegola Tiered Mesh Network Testbed in Rural Scotland," in *Proceedings of ACM Mobicom 2008, Workshop on "Wireless Networks and Systems for Developing Regions" (WiNS-DR)*, 2008.

[3] C. Reis, R. Mahajan, M. Rodrig, D. Wetherall, and J. Zahorjan, "Measurement-based models of delivery and interference in static wireless networks," *SIGCOMM Comput. Commun. Rev.*, vol. 36, pp. 51–62, August 2006.

[4] G. Hardin, "Tragedy of the Commons," *Science*, vol. 162, no. 3859, pp. 1243–1248, Dec 1968. [Online]. Available: http://goo.gl/2NsnG

[5] C.-Y. Wang and H.-Y. Wei, "Ieee 802.11n mac enhancement and performance evaluation," *Mob. Netw. Appl.*, vol. 14, no. 6, pp. 760–771, Dec. 2009. [Online]. Available: http://dx.doi.org/10.1007/s11036-008-0129-2

[6] S. N. Bhatti and G. Knight, "Enabling QoS adaptation decisions for Internet applications," *Computer Networks*, vol. 31, pp. 669–692, 1999.

[7] M. Tauber, S. N. Bhatti, and Y. Yu, "Application Level Energy and Performance Measurements in a Wireless LAN," in *Proc. GreenCom2011: IEEE/ACM Intl. Conference on Green Computing and Communications*, Aug 2011.

[8] M. Tauber and S. N. Bhatti and Y. Yu, "Towards Energy-Awareness in Managing Wireless LAN Applications," in *Proc. NOMS 2012: IEEE/IFIP Network Operations and Management Symposium*, Apr 2012.

[9] R. Punnoose, R. Tseng, and D. Stancil, "Experimental results for interference between Bluetooth and IEEE 802.11b DSSS systems," in *Vehicular Technology Conference, 2001. VTC 2001 Fall. IEEE VTS 54th*, vol. 1, 2001, pp. 67 – 71 vol.1.

[10] S. Fiehe, J. Riihijärvi, and P. Mähönen, "Experimental Study on Performance of IEEE 802.11n and Impact of Interferers on the 2.4 GHz ISM Band," in *Proce. 6th International Wireless Communications and Mobile Computing Conference*, ser. IWCMC '10. New York, NY, USA: ACM, 2010, pp. 47–51.

[11] V. Shrivastava, S. Rayanchu, J. Yoonj, and S. Banerjee, "802.11n under the microscope," in *Proceedings of the 8th ACM SIGCOMM conference on Internet measurement*, ser. IMC '08. New York, NY, USA: ACM, 2008, pp. 105–110.

[12] S. Sendra, P. Fernandez, C. Turro, and J. Lloret, "IEEE 802.11a/b/g/n Indoor Coverage and Performance Comparison," in *6th International Conference on Wireless and Mobile Communications (ICWMC)*, 2010.

[13] D. J. Leith and D. Malone, "Field Measurements of 802.11 Collision, Noise and Hidden-Node Loss Rates," in *WiOpt'10: Modeling and Optimization in Mobile, Ad Hoc, and Wireless Networks*, Avignon, France, 2010, pp. 521–526.

[14] L. Verma, D. Sim, C.-Y. H. Yang, K. Cho, and S. S. Lee, "On Measurement Campaign with IEEE 802.11n Devices," in *Proc. IEEE ICOIN'10*, 2010.

[15] K. Pelechrinis, T. Salonidis, H. Lundgren, and N. Vaidya, "Experimental characterization of 802.11n link quality at high rates," New York, NY, USA, pp. 39–46, 2010. [Online]. Available: http://doi.acm.org/10.1145/1860079.1860087

[16] K. Pelechrinis, I. Broustis, T. Salonidis, S. Krishnamurthy, and P. Mohapatra, "Design and Deployment Considerations for High Performance MIMO Testbeds," in *Wireless Internet Conference, WICON 2008*, 2008.

[17] G. Kassem, I. Ahmad, F. Hameed, and A. Zakariyya, "Tcp variants: An overview," *Computational Intelligence, Modelling and Simulation, International Conference on*, vol. 0, pp. 536–540, 2010.

[18] S. Bhatti and M. Bateman, "Transport Protocol Throughput Fairness," *Journal of Networks*, vol. 4, no. 9, Nov 2009. [Online]. Available: http://dx.doi.org/10.4304/jnw.4.9.881-894

[19] K.-T. Chen, C.-Y. Huang, P. Huang, and C.-L. Lei, "Quantifying Skype User Satisfaction," in *Proceedings of SIGCOMM 06*, August 2006.

[20] D. Bonfiglio, M. Mellia, and M. Meo, "Revealing Skype Traffic: When Randomness Plays with You," in *Proceedings of SIGCOMM 07*, 2007.

[21] M. Zink, K. Suh, Y. Gu, and J. Kurose, "Watch Global, Cache Local: YouTube Network Traffic at a Campus Network - Measurements and Implications," in *Proceedings of IEEE MMCN 08*, 2008.

[22] P. Gill, M. Arlitt, Z. Li, and A. Mahanti, "Youtube traffic characterization: a view from the edge," in *Proceedings of the 7th ACM SIGCOMM conference on Internet measurement (IMC '07)*, 2007.

[23] D. L. Mills, "RFC 958: Network time protocol (NTP)," Sep. 1985. [Online]. Available: http://www.ietf.org/rfc/rfc958.txt

[24] P. Bahl, R. Chandra, T. Moscibroda, R. Murty, and M. Welsh, "White space networking with wi-fi like connectivity," *SIGCOMM Comput. Commun. Rev.*, vol. 39, pp. 27–38, Aug. 2009.

[25] E. Kohler, M. Handley, and S. Floyd, "RFC 4340: Datagram congestion control protocol (DCCP)," Mar. 2006. [Online]. Available: http://www.ietf.org/rfc/rfc4340.txt

[26] S. Bhatti, M. Bateman, and D. Miras, "A comparative performance evaluation of DCCP," in *Performance Evaluation of Computer and Telecommunication Systems, 2008. SPECTS 2008. International Symposium on*, June 2008, pp. 433 –440.

[27] M. Tauber, S. N. Bhatti, N. Melnikov, and J. Schönwälder, "The Case for Heterogeneous WLAN Environments for Converged Networks," in *Proc. ICNC2013 - International Conference on Computing, Networking & Communications*, Jan 2013.