Investigating Environmental Causes of TCP Retransmission and Flags in Wireless Networks

Stuart Cunningham, Nigel Houlden, John Davies, Vic Grout, and Richard Picking
Creative and Applied Research for the Digital Society (CARDS)
Glyndŵr University
Wrexham, Wales, UK
{s.cunningham | n.houlden | j.n.davies | v.grout | r.picking}@glyndwr.ac.uk

***Abstract*—Wireless computer networks are pervasive, but recent work has identified that there remain performance issues that prevent users from utilizing the full capacity of wireless equipment. Much research has taken place to compare the performance of different configurations of wireless networks to one another, but there has been little done recently to determine the effect that the physical environment has upon wireless performance. In this study, we compare TCP retransmission and flag proportions of captured network traffic over a wired and two wireless connection conditions, to determine what extent the physical environment is impacting upon performance in contemporary, real-world networks. Though limited by the practicalities of conducting our experimental work, our analysis indicates that physical environment is not playing a substantial role in the underutilization of conventional wireless networking capacity.**

Keywords—reliability; retransmission, wi-fi; wireless; 802.11.

I. Introduction

This paper attempts to determine the extent to which the physical environment around a real world wireless network affects its performance. We propose that an effective strategy to determine the influence of these factors is to perform a series of experiments that compare undertaking a large data transfer between a wired and wireless network configuration. In fact, we consider the comparison of the wired network performance to two wireless situations, neither of which permits line of sight between the receiving computer and access point. The aim of this work then is to compare results for these three conditions by considering packet retransmission and flagged packet ratios.

The remainder of this paper is organized as follows: in section 2 we provide background and our motivation on the subject. In section 3 environmental causes of wireless networking errors are outlined. Section 4 discusses the experimental setup that we employed and the results obtained. Finally, in section 5 we provide discussion on the outcomes of the research, limitations therein, and directions of future work.

II. Background

Recently published work sought to investigate the phenomenon of impeded throughput in 802.11 networks, with a particular focus upon determining causal factors or conditions. A particular strength in this recent article is that experimental data is captured in real-world networks, thereby acknowledging and incorporating the range of environmental and usage situations that are likely to inhibit throughput in the day-to-day user experience of 802.11 networks [1].

As such, the research detailed in [1] shares aspects of rationale and intention with the work that is presented here. These are twofold: First, the desire to conduct research into real-world wireless networks as opposed to those in a controlled, laboratory setting. Second, to determine, since a period of over 10 years has passed since initial work into throughput in wireless networks was conducted, showing throughput of 41% and 55% [2]; 28% of data transmissions being retransmitted data [3]; and that link reliability can be indicated by measuring packet retransmission [4], if the performance of contemporary 802.11 networking equipment still suffers with the same degree of problems caused by the physical environment and conditions. As such, the work presented here builds, and seeks to expand upon, aspects highlighted in the work of Murray *et al.*, specifically by attempting to determine if physical and environmental factors where the network is deployed are contributory to the level of retransmission encountered in 802.11 networks.

Ensuring optimum performance of wireless networks is more and more important given their pervasive nature and the increasing number of devices to be connected, due to the growth in mobile devices, sensor devices, and the broader Internet of Things (IoT). As such, performance and fairness of access are paramount [5].

One limitation in the current literature is that evaluation of wireless network performance predominantly is concerned with comparing variations in wireless configuration and standards, rather than evaluating performance due to physical factors by benchmarking results against a wired network. However, one notable work that does employ wired and wireless comparison, of peer-to-peer software performance, is that of Quan, Lee and Pinkston [6].

Another example, work by Sallabi Abu Odeh, and Shuaib is commendable for experimental work in a real world environment, although their analysis is oriented around the use of overlapping and non-overlapping channels [7]. An earlier work by Doefexi, Armour, Lee, Nix, and Bull also examined wireless network performance in an office scenario, but only considered a comparison between 802.11a and 802.11g [8].

III. Investigating Causes of Wireless Packet Retransmission

A. *Scope for Investigation*

Based on the literature discussed in the previous section, it seems that retransmission rates for packets in wireless networks is a topic that has not received significant attention, especially in determining the cause of this phenomenon with respect to the physical environment where the network is installed. The work that has been done on the topic is largely descriptive and explorative. To this extent, there is scope for future research in two main categories: first, to gather larger volumes of data on networks of varying scale and size in order to *validate* that these retransmission rates are representative of the wider phenomenon and, second, to take a more controlled and deductive approach in determining the *causes* of the observed retransmission rates.

The theme of validation has been discussed, so far as is possible, in the previous literature review. It is not the main intention of this study to attempt to recreate or extend any of these previous studies to determine the reliability and transferability of these figures. This is not to dismiss such an activity as being without value, but it is of size and scope that is presently impractical. For researchers interested in this avenue, we suggest several replications will be of use as well as extending the work to undertake similar studies in wireless networks of variable size, hardware configuration, and traffic loading.

The focus of this paper, therefore, is in undertaking an initial probe as to the cause of the reportedly high retransmission rates of previous studies. We begin this task by considering the technological and physical properties of wireless transmission and its environment, which separates it from its wired counterpart. Following this, we undertake an initial pilot study to illustrate how the problem might be approached on a larger scale in future.

B. *Environmental Causes of Packet Retransmission*

All computer networks are prone to errors in transmission and this is an acknowledged impediment in data transfer, regardless of whether the network is wired or wireless. In the focus of this research, we are primarily interested in determining the characteristics of wireless networks that introduce scope for errors, and hence retransmission, over and above those encountered in wired equivalents.

Given that errors in data networks are generally caused by attenuation and noise [9] it is upon these lines that we outline possible causes of error in wireless scenarios.

In terms of attenuation, wireless networks are particularly prone. Data exchanges are affected by natural attenuation of the signal as a function of distance between a node and access point. In practical situations attenuation is also impacted by the fact that there is rarely line of sight between node and access point. This is especially the case since wireless networks are designed to allow users to be mobile and therefore the intermittent, possibly persistent, obstruction or occlusion of the wireless signal will occur. As such, attenuation beyond distance will be hard to guarantee or predict, since various obstructions will be made from different density materials, and therefore absorption coefficients, reflective characteristics, and absorptive surface areas will vary [10].

In terms of noise in wireless networks, contemporary standards are generally robust, although naturally still impaired by electromagnetic interference and impulse noise cause by other microwave frequency devices in the vicinity.

In this work, issues of noise from causes such as electromagnetic interference are considered beyond the scope and control of the experimental work conducted here. It is intended that these factors will appear only as background noise given the scale of the intended study.

IV. Experimental Study

A. *Aims*

The main intention of the experimental work documented here is to identify the extent to which physical conditions, such as reflection and absorption of the wireless signal, significantly contribute to performance differences when compared to wired networks. The null hypothesis of the work is that differences experienced in performance between wired and wireless networks are down to other factors, whilst the alternate hypothesis is that physical conditions around the network access point and receiving node are likely to cause performance differences between wireless and wireless conditions.

Ideally, conducting several repetitions of a pairwise experiment in a highly controlled environment would be used to investigate these hypotheses. For example, the ideal research design would be within a microwave anechoic chamber, which would allow for one experimental configuration with 100% absorption in the room and a clear line of sight connection between an access point (AP) and computer as the control condition. The test configuration would then take place by installing a highly reflective box within the anechoic chamber whilst still maintaining line of sight between the AP and the computer. Under both conditions, a series of data transfers would then be initiated, over a fixed period of time, and the number of packet errors logged. Analysis of this data would therefore indicate if any significant difference in the number of packet errors were attributable to the reflections caused inside the test condition.

However, this ideal situation is, at the present time, not practicable. Furthermore, it does not represent the experience of wireless network users in real world scenarios. To this extent, we conduct experimental work under conditions of convenience in a real world setting to evaluate any differences between a wired connection; a wireless connection without line of sight; and a wireless connection without line of sight and with no unobstructed signal path.

B. *Methodology*

Data was captured on five days, selected by the researchers at convenience, over a two-week period. Samples were captured during working days and hours of the University, again at the convenience of the researchers. A MacBook Pro laptop computer was located on a desk, at a height of 73 cm above the floor, in an office at the University and the Wireshark 1.12.7 software was used to capture network packets. The built-in Ethernet and wireless adapters of the computer were used.

There was a wireless AP located in the corridor outside the office, the model of which was a Cisco Aironet. The AP was located on a wall at a height of 250 cm above the floor. The wall dividing the office from the main corridor is made of plastered brick and has a thickness of approximately 45 cm. The office door is 4 cm thick, wood with a large glass panel occupying around 60% of the door area. In the corridor outside there is a double fire door of 4.3 cm thickness with a small glass panel in each door. The physical environment around the office, computer, and AP are illustrated in **Fig. 1**. From this information, the straight-line distance between the AP and computer is calculated as being approximately 906 cm.

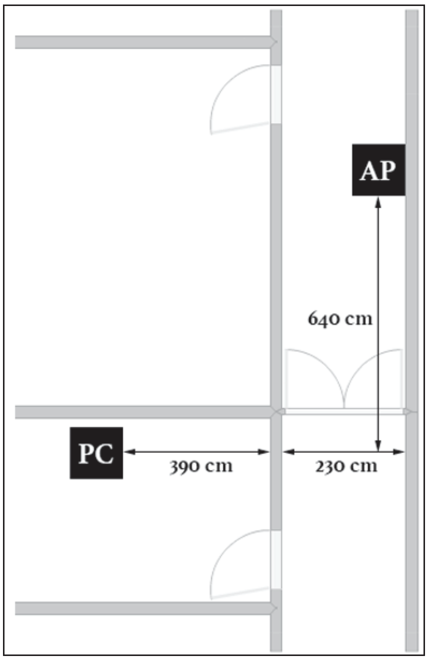

Fig. 1: Experiment Physical Setup

The overarching aim of the experiment was to capture substantial amounts of network traffic over three connection conditions: wired Ethernet; wireless with the office door open; and wireless with the office door closed. To generate a predictable amount of traffic during the captures and to keep the majority of the data under experimental control, a large video was loaded into the Firefox web browser that had duration of 61 minutes and 18 seconds (the video used is at https://www.youtube.com/watch?v=dVfmS-dG38U) and was viewed at Full HD (1080p) resolution. Wireshark was set to capture for a fixed time period of 65 minutes on each occasion. This allowed the researchers to begin the capture, load Firefox and the YouTube video, skip any pre-viewing adverts, and set the video resolution to Full HD. The YouTube video was chosen since it is a non real-time source and hence uses the TCP protocol to deliver the video data. These parameters aside, no other controls were exerted over the experiment, such as the state of the fire door in the corridor, people in the office or corridor outside, or the overall usage of the wired and wireless networks of the University when captures were made. It is intended that these extraneous factors will be normalized out of the data by the size of the packet captures.

It is important to stress that each triplet of connection condition captures were not recorded simultaneously. Instead, they were recorded in a linear sequence, chosen at random.

C. *Results*

The five capture sessions resulted in a total of 16798416 packets. These are broken down into the separate five capture sessions represented, in chronological order, in the results and analysis as labels S1 through S5. Although time was fixed at 65 minutes for each capture, the number of packets varied due to network conditions. The mean number of packets per capture condition is 1119894. **TABLE I** provides an overview of the data captured.

TABLE I. SUMMARY OF PACKETS CAPTURED BY CONNECTION CONDITION

Sess.	**Experiment Condition**	**Number of Packets**			
		Total	**TCP**	**TCP Retransmission**	**Bad TCP**
1	Wired	1216691	1203401	773	26720
	Wireless DoorOpen	1090689	1090208	758	29837
	Wireless DoorClosed	1156434	1155994	3659	57015
2	Wired	1144501	1131097	1426	39970
	Wireless DoorOpen	1247702	1247071	4940	83872
	Wireless DoorClosed	1064452	1063877	331	22033
3	Wired	1174117	1154817	15544	111902
	Wireless DoorOpen	1081932	1081462	623	28180
	Wireless DoorClosed	1100937	1100574	10882	95085
4	Wired	1123591	1106750	565	23218
	Wireless DoorOpen	1202512	1201998	5061	106416
	Wireless DoorClosed	1066336	1065822	7023	76111
5	Wired	937030	914674	7432	91214
	Wireless DoorOpen	1273043	1272461	5061	95189
	Wireless DoorClosed	918449	917897	19290	121338

The work of this paper is primarily concerned with TCP packets and the number of these that have reported issues. Hence, the analysis focuses on the presence of 'Bad TCP' packets, as defined by Wireshark as being all TCP packets that have been flagged, with the exception of updates to the sender's TCP buffer. As another explicit measure in this experiment the number of TCP packets that have been retransmitted are also shown, which is a subset of the Bad TCP packets.

D. Analysis

1) TCP Retransmission

The initial analysis of the data captured is concerned with the proportions of TCP packet retransmission rates (percentage of all TCP packets), over the five sessions and this is illustrated in **Fig. 2**. In terms of TCP retransmission rates over all five of the sessions, the wired condition experienced 0.467%; the wireless with the door closed 0.776%; and wireless with the door open 0.279%.

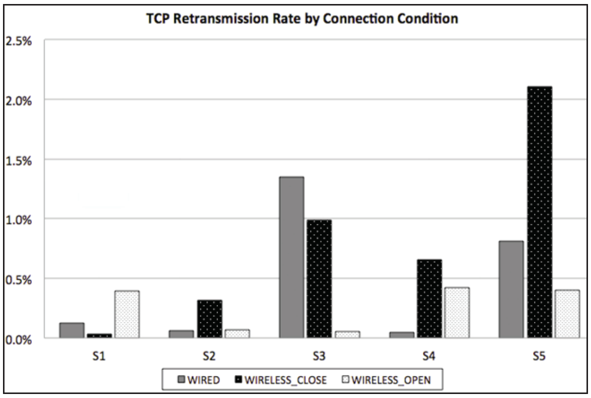

Fig. 2: Proportion of TCP Packets Retransmitted by Connection Condition

From inspection it is seen that there is a degree of fluctuation and variation within the results and across the five sessions. Generally, both wireless connections have higher rates of retransmission, although it is noteworthy that in S3 and S5 the retransmission rates in the wired connections have increased. Indeed, in condition S3 it is unusual that the wired connection condition has a higher retransmission rate than both wireless conditions. Aside from that anomaly, the results broadly fit what would be expected: that the wireless connection with the office door closed, results in a larger number of retransmissions than those where the door is open, suggesting that the presence of the door means the signal is suffering from more attenuation and negative reflection.

A repeated-measures Analysis Of Variance (ANOVA) with a Greenhouse-Geisser correction was employed to objectively analyze the TCP retransmission rates. Prior to this analysis, the percentage values were transformed using the arcsin function, which changes the binomially distributed percentage scale to one that is normally distributed so as to allow for the ANOVA test to be conducted. The arcsin transformation function in degrees of a percentage value x represented as a proportion (i.e. 25% is recorded as 0.25) is, as follows

$$\arcsin(x) = \left(\sin^{-1}(\sqrt{x})\right)\frac{180}{\pi} \tag{1}$$

The test concluded that there was not a statistically significant difference between the connection conditions ($F(1.734, 6.937) = 1.153$, $p > 0.05$).

2) Bad TCP Packets

Similar analysis is then conducted regarding the Bad TCP proportions that have been recorded during the data capture experiment, to determine if this broader set of information can also indicate if reduction in performance is related to the connection condition. An illustration of Bad TCP (as percentage of all TCP packets) is shown in **Fig. 3**. In terms of Bad TCP packet rates over all five of the sessions, the wired condition experienced 5.317%; the wireless with the door closed 7.005%; and wireless with the door open 5.829%.

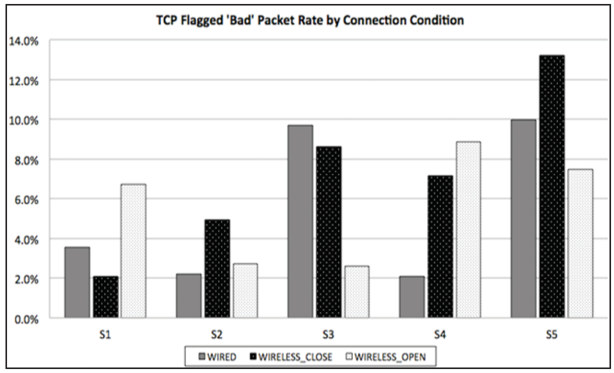

Fig. 3: Proportion of Bad TCP Packets by Connection Condition

This analysis shows an unexpected degree of variation in the results. There are no easily identifiable trends, when looking across all of the capture sessions, especially once the TCP retransmissions, as accounted for in the previous sub-section, are factored out of the analysis. Whilst the total figures are broadly inline with what would be expected, on a session-by-session basis it is hard to identify any particular outcome.

A repeated-measures ANOVA with a Greenhouse-Geisser correction was again employed to objectively analyze the Bad TCP rates upon the percentage values transformed using the arcsin function. The test concluded that there was not a statistically significant difference between the connection conditions ($F(1.539, 6.158) = 0.389$, $p > 0.05$).

3) TCP Packet Lengths

As a final piece of analysis, the packet lengths used across each of the three connection conditions was investigated, since the presence of any major problems on the wireless network conditions would show as a reduction in packet lengths. This is illustrated in **Fig. 4**.

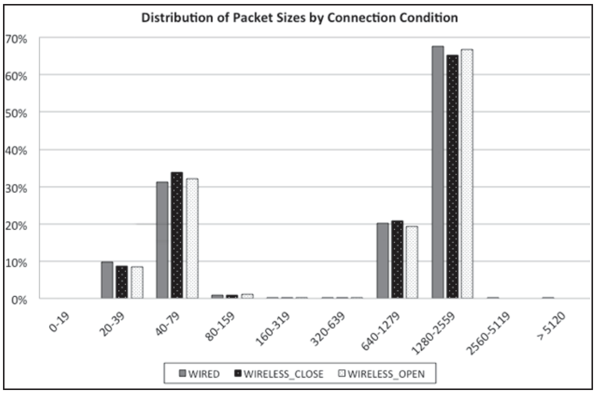

Fig. 4: Packet Sizes by Experiment Connection Condition

Distribution of packet lengths does not appear to have been impacted by the connection conditions evaluated in the study, indicating that there have not ben any exceptional problems in the delivery of the video data in the experiment.

V. Discussion and Future Work

The work presented here shows lower than anticipated presence of issues with the two wireless connection conditions being evaluated. Further, the analysis of experimental results concludes that there are no significant differences between the two wireless conditions and the wired counterpart. As such, this evidence would suggest that issues relating to and overall link quality and TCP issues are not related to the physical environment. These results, as measured in this paper, seem to be out of step with the general experience of others in the literature. However, there are a number of methodological limitations that may be contributing to these results and that warrant future investigation.

One issue stems from the non-concurrent capture of the data that has been analyzed. Whilst each session was undertaken on the same day these captures were, by design, taken at different times of day, albeit randomized over all conditions. To an extent, this partly negates the use of repeated measures ANOVA, although it was the closest to the ideal condition as could be practically obtained at this time.

Another factor to be considered is the distance between the AP and computer used in the experiments. Placed at a straight-line distance of 906 cm, the receiving computer, it is suggested, was working with a good level of signal strength and quality, albeit attenuated by the thick walls. A more effective analysis might have implemented the experimental conditions at greater distances and by measuring Signal-to-Noise Ratio (SNR) at each point.

Thus, future work proposed should attempt to take multiple measures of performance, in parallel data capture situations, so as to allow for an unquestionable triplet-wise comparison of the results. Other modifications, such as: altering distances, presence of obstructions and occlusion between the AP and computer, and measuring SNR to determine if this correlates with packet loss and/or retransmission proportions, should also be pursued. Ultimately, it is hoped that work can be undertaken in a microwave anechoic chamber, as indicated in section 4, which will serve to provide a more conclusive, though not real world, test bed for future work to use as a reliable benchmark.

References

[1] D. Murray, T. Koziniec, M. Dixon, K. Lee, "Measuring the reliability of 802.11 WiFi networks," in *Internet Technologies and Applications (ITA), 2015*, pp.233-238, 8-11 Sept. 2015

[2] J. Jun; P. Peddabachagari, M. Sichitiu, "Theoretical maximum throughput of IEEE 802.11 and its applications," in *Network Computing and Applications, 2003. NCA 2003. Second IEEE International Symposium on* , vol., no., pp.249-256, 18-18 April 2003.

[3] M. Rodrig, C. Reis, R. Mahajan, D. Wetherall, J. Zahorjan, “Measurement-based characterization of 802.11 in a hotspot setting,” in *Proceedings of the 2005 ACM SIGCOMM workshop on Experimental approaches to wireless network design and analysis 2005* Aug 22 (pp. 5-10). ACM.

[4] A.P. Jardosh, K.N. Ramachandran, K.C. Almeroth, E.M. Belding-Royer, “Understanding link-layer behavior in highly congested IEEE 802.11 b wireless networks,” in *Proceedings of the 2005 ACM SIGCOMM workshop on Experimental approaches to wireless network design and analysis 2005* Aug 22 (pp. 11-16). ACM.

[5] H. Shi, R.V. Prasad, E. Onur, I.G., Niemegeers, “Fairness in wireless networks: Issues, measures and challenges,” *Communications Surveys & Tutorials*, IEEE. 2014 May;16(1):5-24.

[6] L. Quan, K.G. Lee, T.M. Pinkston, “Performance analysis of unstructured peer-to-peer schemes in integrated wired and wireless network environments,” In Parallel and Distributed Systems, 2005. Proceedings. 11th International Conference on 2005 Jul 20 (Vol. 1, pp. 419-425). IEEE.

[7] F.M. Sallabi A.O. Abu Odeh, K. Shuaib, “Performance evaluation of deploying wireless sensor networks in a highly dynamic WLAN and a highly populated indoor environment,” In Wireless Communications and Mobile Computing Conference (IWCMC), 2011 7th International 2011 Jul 4 (pp. 1371-1376). IEEE.

[8] A. Doefexi, S. Armour, B.S. Lee, A. Nix, D. Bull, “An evaluation of the performance of IEEE 802.11 a and 802.11 g wireless local area networks in a corporate office environment,” In Communications, 2003. ICC'03. IEEE International Conference on 2003 May 11 (Vol. 2, pp. 1196-1200). IEEE.

[9] W. Stallings, Data and computer communications, International ed., Pearson Education, 2014.

[10] J.E. McDonnell, “Characteristics of the indoor wireless propagation environment at microwave and millimetre frequencies,” In Radio Communications at Microwave and Millimetre Wave Frequencies (Digest No. 1996/239), IEE Colloquium on 1996 Dec 16 (pp. 13-1). IET.