# RCBurst: a Mechanism to Mitigate the Impact of Hidden Terminals in Home WLANs

Mustafa Al-Bado, Cormac J. Sreenan, Kenneth N. Brown
CTVR, Department of Computer Science, University College Cork, Ireland
Emails: mustafa.albado@insight-centre.org, cjs@cs.ucc.ie, k.brown@cs.ucc.ie

***Abstract*—In dense wireless deployments, such as Enterprise WLANs (EWLANs) and home WLANs, interference may occur because of neighbouring WLANs sharing the same unlicensed spectrum. Mechanisms to centrally manage WLAN deployments cannot effectively mitigate the interference caused by hidden terminals (HTs) in WLANs that belong to different organisations. Furthermore, the impact of interference is amplified if it is combined with long-lived TCP traffic flows, which are becoming increasingly commonplace. In this paper, we focus on mitigating the impact of HTs on long-lived TCP flows in home WLANs. In particular, we study the effect of five key factors on long-lived TCP flows under the impact of HTs: packet bursting, back-off mechanisms, maximum number of RTS attempts, capture affect and the number of associated clients with the same Access Point (AP). Extensive simulation results show that a combination between RTS/CTS messages and bursting increases the throughput up to 8x in the presence of HTs. Therefore, we develop a mechanism called joint RTS/CTS with Bursting (RCBurst) that leverages RTS/CTS messages and packet bursting to mitigate the impact of HTs. The simulation results show that RCBurst achieves an improvement of up to 0.3 in Jain's fairness index over the conventional CSMA/CA, without reducing the overall throughput.**

## I. Introduction

Wireless LANS (WLANs) based on 802.11/WiFi are widely deployed in residences and enterprises. The unplanned deployment of WLANs, together with the shared-medium nature of wireless communications, may cause performance deterioration in congested scenarios [14], a key factor being the impact of hidden terminals (HTs). To mitigate the impact of HTs in home WLANs and Enterprise WLANs (EWLANs), several research studies focus on having centralised network-wide solutions, in the form of controllers or traffic schedulers [6], [7], [13]. These efforts can help to mitigate the impact of HTs that are internal to the network, but so-called rogues HTs which exist outside the network, are not considered. In [14] the authors report that the effect of rogues HT causes up to 50% packet loss at the MAC layer during rush hours in a well-deployed and managed EWLAN [13],[14], [15], [10]. In this paper we focus on home WLAN deployments and dealing with the interference caused by rogue HTs[1]. Virtually all home WLANs today are deployed in an uncoordinated manner, and often using default channel settings. In urban settings it can be expected that such deployments will be dense, and thus the impact of rogue HTs will be highly significant. Meanwhile the popular use of video streaming (e.g., Netflix and Youtube) implies that interference periods are unlikely to be transient. In the presence of rogue nodes, the resulting interference cannot be detected or resolved using only a centralised network controller; additional mechanisms and feedback from Access Points (APs) is necessary, thus motivating our study.

In this paper we make two key contributions. Firstly, we present an extensive simulation study to measure the effect of several critical AP mechanisms on mitigating HT interference. Secondly, we propose RCBurst to mitigate the interference in dense and saturated WLANs. using NS-3, RCBurst is evaluated in four different dense home WLAN topologies. The results show that using RTS/CTS messages together with packet bursting and a fixed back-off improves the throughput by up to 8x for a link under HT interference. Also, RCBurst achieves up to 0.30 improvement in Jain's fairness index over the conventional CSMA/CA in dense topologies without reducing the overall throughput.

The remainder of the paper is organised as follows. Section II presents the related work. In section III, the aims and methodology are introduced. In IV, we investigate the impact of key AP mechanisms on mitigating the interference. In section V, we propose and evaluate RCBurst. Finally, we conclude the paper in Section VI.

## II. Related Work

We start this Section with a review of measurement studies that report interference regardless of the deployment and type of the WLANs. Then we summarise published work on interference mitigation, followed by recent network-wide proposals to manage interference.

**Measurement:** User traffic behaviour and wireless performance in EWLANs and home WLANs have been widely studied [8], [14], [15], [10]. One finding of these studies is that interference is the main source of packet loss. They also show that APs tend to use high transmit data-rates unless the link is subject to interference; the latter thus leads to a severe decline in the transmit data-rate. In [14], the authors conduct their measurement study on a well-planned EWLAN. They conclude that the impact of rogue nodes is relatively high in rush hours and that the external interference causes up to 50% packet loss at MAC layer.

**Interference mitigation:** The authors of [12] propose a hybrid CSMA/CA-TDMA scheme to mitigate the impact of both congestion and interference on the uplink. The proposed

[1]Throughout the paper, when we refer to interference we specifically mean interference caused by HTs.

978-1-5386-1629-1/17/$31.00 ©2017 IEEE

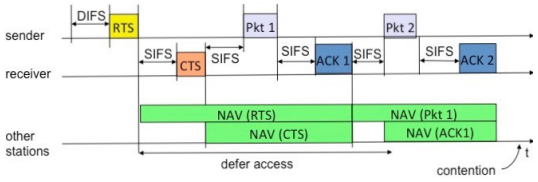


Fig. 1: An example for a combination of RTS/CTS exchange with burst of two packets.

scheme relies on the client's buffer status to decide whether CSMA/CA or TDMA is more efficient to use. Our focus is on downlink traffic because it is far more relevant to home WLANs [15], [10]. Adaptive RTS/CTS is used to mitigate interference in several papers [5], [16], [3], [4], [11]. In all cases they use adaptive RTS/CTS, where RTS/CTS is turned on if congestion is suspected. However, these studies do not consider using RTS/CTS messages to detect the interference, and neither do they take into account the benefit of trying to reduce employing these messages.

**Enterprise and residential Network Management**: In [13], the authors propose CENTAUR to mitigate the impact of both hidden and exposed terminals. CENTAUR assumes that the traffic in a EWLAN comes from a single gateway. Therefore, the traffic is scheduled from the gateway to avoid internal interference. In [6] and [7], they adopt mechanisms using time slotting to schedule the traffic among internal HTs for home WLANs. These do not consider the impact of rogue nodes, which cannot be mitigated purely by packet scheduling. Moreover, the mechanisms in [6] and [7] do not consider the implications of having multiple-client associated with the same AP.

The outcome of this paper is a new mechanism, RCBurst, that handles strong interference regardless of its type (internal or external), and yields demonstrated improvements in overall performance.

## III. Aims and Methodology

The first goal of our research is to identify the influence of key factors on the performance of long-lived TCP flows under HT interference. Understanding the effect of each factor allows us to develop RCBurst in Section V to mitigate the interference. Therefore, we follow a step-by-step approach and evaluate:

1) **The impact of burst length:** Packet bursting is a well known technique in wireless networks. For instance, it is employed in 802.11n to increase the offered throughput [9]. In 802.11n, multiple frames are aggregated in a single frame then sent and acknowledged by a BlockAck [9]. Frame aggregation makes transmitted frames be more vulnerable to interference due to the length of the aggregated frame. In [9], Pefkianakis et al. propose a solution that relies on using RTS/CTS messages in case of interference suspicion. This mechanism may cause an unnecessary limitation for medium access for neighbouring APs and clients in case of missing RTS messages. In addition, in [9], it is reported the effect of interference and used data-rate on reducing/cancelling the gain of frame aggregation. Moreover, frame aggregation does not work properly in strong and saturated interference environments. In our work, packet bursting is employed in a different context, which works for both legacy 802.11 and advanced versions (e.g., 802.11n). Each packet in the burst reserves the medium for the following packet using a Network Allocation Vector (NAV). Here, the NAV duration field in the MAC header of the current packet is filled in with the required time to transmit the following packet. Packets are sent back-to-back after a successful exchange of RTS/CTS messages. This approach forces the interfering nodes to postpone their transmissions if they decode either the transmitted packet or its ACK. As a result, clients that suffer from strong interference can receive their packets in a relatively interference-free environment. Packet bursting reduces the cost of using RTS/CTS messages because only one RTS/CTS is required per-burst (i.e., multiple-packet) instead of per-packet. Fig 1 illustrates an example for combination of RTS/CTS exchange with burst of packets with length 2. To the best of our knowledge, we are the first to measure the effect of using the combination of RTS/CTS messages and packet bursting with NAV.
2) **The impact of the maximum number of RTS attempts:** We study the impact of using different maximum number of RTS attempts (MNRA). Certain types of traffic, such as TCP, are severely affected by packet losses, leading to lengthy timeouts and costly retransmissions. Reducing the effect of packet losses can be achieved by either using a large maximum number of packet retries or a large MNRA. In an environment with high levels of interference, RTS/CTS messages cause less collision probability than the other approach due to their relatively small packet sizes. Therefore, increasing the MNRA, instead of the maximum number of packet retries, reduces packet losses and improves goodput.
3) **The impact of rate control algorithms in capture effect environments:** Rate control algorithms (RCAs) are widely used to tune the transmission rate to optimise throughput based on the wireless channel conditions. We study the ability of RCAs in tolerating weak interference and accommodating concurrent transmissions. We use Minstrel [1] as a RCA because it is widely known and offers excellent performance.

There are two additional factors to consider:

- **The impact of the number of associated clients with an AP:** We vary the number of clients associated with an AP that serves a link suffering from strong interference. The aforementioned factors directly affect the traffic of each client, and this altered traffic can directly influence the performance of the other client(s) and we explore the implications of this interaction.
- **The impact of the back-off mechanisms:** Two mech-

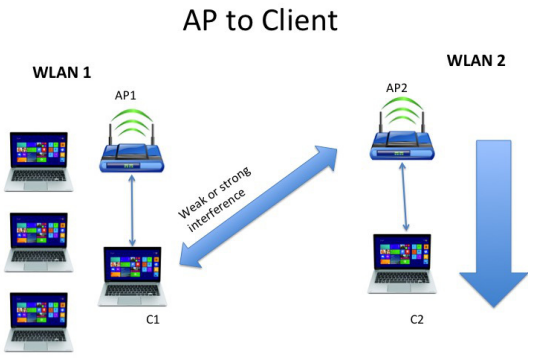


Fig. 2: The scenario represents the interference from a neighbouring WLAN (i.e., WLAN2) where client C1 is suffering interference from AP2.

anisms are considered when reacting to the loss of RTS messages: Fixed Back-Off (FBO) and Exponential Back-Off (EBO). For all other packet types, EBO is used as normal. Using a large MNRA may lead to a relatively long back-off time. This back-off time influences all clients that associated with an AP dealing with a link suffering from strong interference. To mitigate this problem we investigate the impact of using FBO, which is shown in [13] to perform well in handling the exposed terminal scenario. We have adapted FBO for use in strong interference scenarios; to handle RTS message losses by immediately employing the minimum contention window.

The NS-3 simulator version 3.23 is used for the experimental evaluation. We use the model in [2] to create network topologies that represent the wireless link errors observed in a real testbed, with Received Signal Strength (RSS) profiles for the links. Each RSS profile consists of a RSS ECDF distribution, where the difference between the minimum and maximum values of the distribution is $4dB$. In our setup, three profiles are used: None, Weak and Strong. The None RSS profile is used to represent no interference/connectivity between nodes where the RSS values are much lower than the noise level. The Weak profile supports data-rates up to $12Mpbs$ in the absence of interference. The Strong profile supports data-rates up to $54Mbps$ in the absence of interference. In the simulations we send TCP unicast traffic on long-lived flows, using the NS-3 OnOffApplication as a traffic generator with $1024B$ packet size. The simulation results are presented with their $95\%$ confidence intervals and based on $12-50$ simulation runs as shown. Each run lasts $120s$.

## IV. AP's Actions to Mitigate the Impact of Uncontrolled Interference

In this section, we follow the approach outlined in Section III. The scenario in Fig. 2 exhibits the considered topology for running the characterisation experiments. The explored actions are implemented only on the downlink from AP1 to C1. We assume that APs run optimisation algorithms for channel selection and power control, so the contention among APs is at the lowest level. This assumption is reported in [14] and [10] for both EWLAN and home WLAN, respectively. In Section V-C, we use more complex dense topologies that consist of four simulated home WLAN scenarios, where APs may serve multiple clients that suffer from strong interference.

### A. The impact of burst length

To alleviate the collision probability due to the packet size, the MNRA is set to 24 (the impact of MNRA is investigated in Section IV-B). We had previously conducted simulations for UDP traffic which showed that when the burst length is greater than 13, the medium access tends to be unfair towards WLAN2. However, in this experiment, we vary the burst length between $1-36$ to measure the effect of burst length on TCP traffic. Note that when the burst length is one, RTS/CTS messages are sent per-packet. The type of interference from AP2 to C1 is strong. In case of multiple clients, a separate queue is assigned for the client that suffers from strong interference. When the transmission turn comes to the first packet in the separate queue, it will be transmitted back-to-back with the following $m$ packets from the same queue, where $m \leq l-1$ and $l$ is the burst length. The AP does not send the full burst length if the separate queue runs out of packets.

Fig. 3 (left) presents the throughput of the client that suffers from strong interference (i.e. C1 in fig. 2). From the figure we can see that exchanging RTS/CTS messages per-burst provides significant throughput gains over exchanging them per-packet (e.g., from $0.5Mpbs$ to over $3Mbps$ for $N=4$). This is expected because the longer the burst means that less RTS/CTS messages are needed and the less collisions occur. When C1 is the only associated client with AP1 (i.e., $N=1$), C1's throughput converges from the burst length of 16 packets. One explanation for this convergence is that the downlink traffic of WLAN2 limits the uplink traffic of C1. As a result, the TCP server of C1 cannot send more packets to C1 exceeding the limit of its congestion window. Moreover, the results show that C1's throughput declines significantly (to 30%) for every new active client added to WLAN1.

Fig. 3 (middle) demonstrates the overall throughput of WLAN1. It shows that the throughput is in direct proportion to the burst length of C1 when the number of clients is above 1. The figure shows that the burst length has a considerable impact on the EBO performance. For instance, increasing the burst length from 1 to 36 leads to throughput gains up to 6.5x for $N=3$. This can be explained because using longer bursts implies employing less RTS messages, which save the AP from exponential waiting for every RTS loss. Fig. 3 (right) presents the throughput of C2 (i.e., the client in WLAN2). It shows that C2's throughput declines with the increase of C1's burst length. However, the decline becomes smaller with increasing the number of active clients in WLAN1.

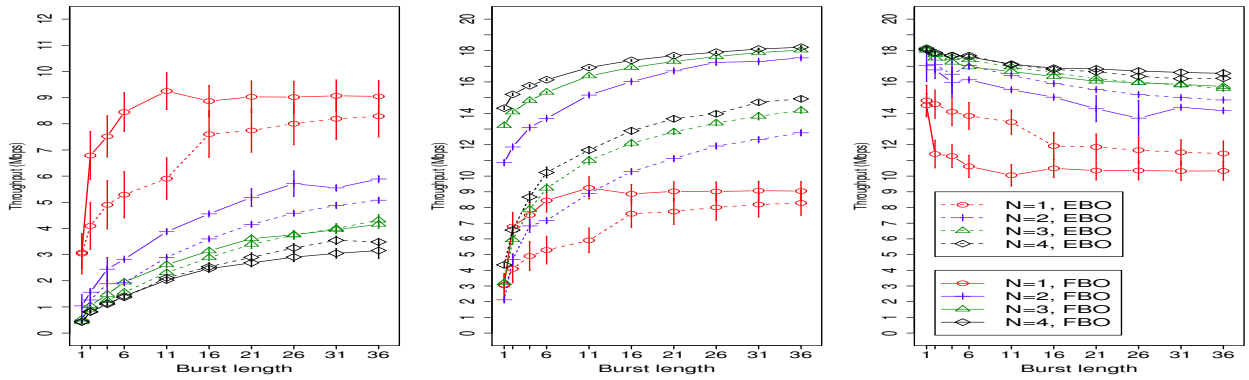


Fig. 3: Burst length Vs. throughput. The left figure presents the throughput of the client under a strong interference. The middle figure shows the overall throughput of WLAN1. The right figure presents the overall throughput of WLAN2 (i.e. client C2). N indicates to the number of clients in WLAN1

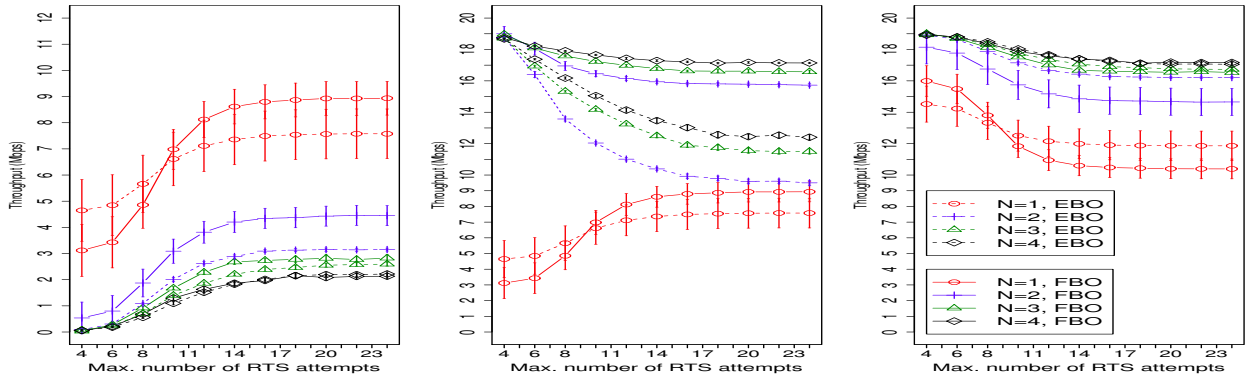


Fig. 4: MNRA Vs. throughput. The left figure presents the throughput of the client under a strong interference. The middle figure shows the overall throughput of WLAN1. The right figure presents the overall throughput of WLAN2 (i.e. client C2). N indicates to the number of clients in WLAN1

### *B. The impact of the MNRA*

Here, we focus on the impact of the MNRA joint with a fix burst length (13 packets, see Section IV-A). The MNRA is varied between $4-24$. The interference from AP2 to C1 is strong.

Fig. 4 (left) presents the throughput of C1. It shows that when C1 is the only client of AP1, the throughput converges after the MNRA of 14. Also, we can see that increasing the MNRA from 4 to 16 leads to throughput gains up to 3x and 1.6 for FBO and EBO respectively.

When two clients associated with AP1, the C1's throughput gain for increasing the MNRA from 4 to 24 is up to 8x for FBO. In the EBO's case, the throughput gain goes from nearly $0.0Mbps$ to $3Mbps$. When the number of associated clients in WLAN1 is above 2, the C1's throughput gains for increasing the MNRA from 4 to 16 go from nearly $0.0Mbps$ to around $2.5Mbps$ and $2Mbps$ for $N=3$ and $N=4$ respectively. Also as mentioned in Section IV-A, increasing the number of clients in WLAN1 declines significantly (to 50%) the throughput of C1 for every new active client.

Fig. 4 (middle) illustrates the overall throughput of WLAN1. It shows that the throughput declines to 1.7x when EBO is used compared to FBO. The results suggest that using a large MNRA causes a significant increase in C1's throughput. However, it leads to a decline in the overall WLAN1's throughput relative to the amount of throughput gained by C1. For instance, when MNRA is 24 the WLAN1's throughput declines around 0.67x and 2.5x of C1's gained throughput for FBO and EBO respectively. We believe that the cost can be reduced if the transmissions are scheduled based on their priority. Fig. 4 (right) shows the throughput of WLAN2. It shows that increasing the number of clients in WLAN1 leads to reduce the impact on WLAN2.

### *C. The impact of RCAs in capture effect environments*

In this step, we study the behaviour of RCAs (in this case Minstrel) in weak interference environments. The experiment is run using four fixed data-rates in addition to Minstrel. The fixed data-rates are $24, 36, 54$ and $54Mbps$ with RTS/CTS messages. The MNRA is set to 7 and EBO is only used. Packet bursting is not used in this experiment as we wish to isolate the

RCA effect. Fig. 5 shows that Minstrel achieves the highest throughput for C1. In fig. 6 we can see that Minstrel reaches the second highest throughput for the remaining associated clients with AP1. The $54Mbps$ data-rate achieves the highest throughput for the remaining clients but has almost zero throughput for C1 because the SINR value is not enough for C1 to perform a successful capture affect in that data-rate. Minstrel and $54Mbps$ achieve the highest throughput for C2, but the results is not presented because of the space limitations. The previous results suggest that RCAs tolerate weak interference and can perform concurrent transmissions without the need for synchronisation among APs for accessing the medium.

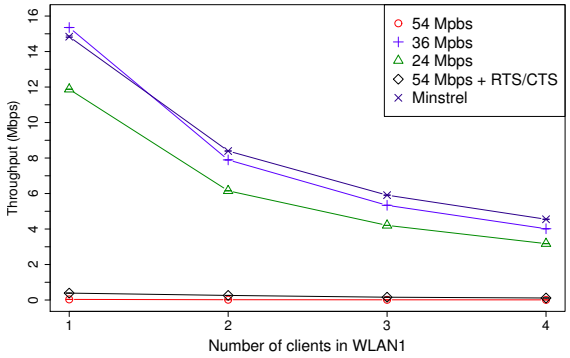


Fig. 5: The throughput of C1 from Fig. 2 when the interference is weak from AP2 to C1

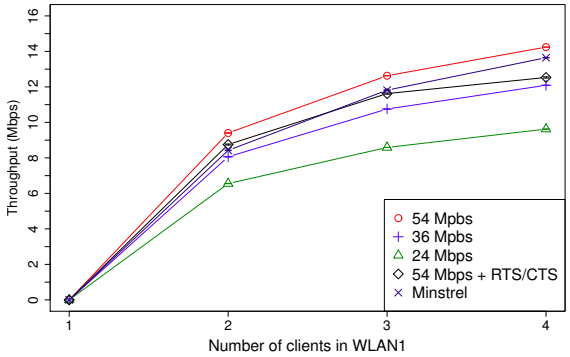


Fig. 6: The throughput of clients in WLAN1 excluding C1 when the interference is weak from AP2 to C1

## V. RCBurst

In this section, we propose RCBurst to detect and mitigate strong interference. RCBurst can either operate in a stand-alone manner to manage interference or collaborate with a network controller by reporting to the controller when a rogue HT effect is present. Moreover, RCBurst can provide the controller with a list of APs where clients have used CTS messages recently. The controller can leverage this list to manage the internal interference by narrowing down the interference sources.

The following properties are specified in the design of RCBurst: 1) to be compatible with 802.11, 2) to run only at APs, 3) to enforce fairness among clients, 4) to allow concurrent transmissions if the interference is tolerable. RCBurst consists of two parts: strong interference detection and opportunistic burst scheduling. We elaborate about these parts in the following two subsections.

### A. Strong interference detection

The link is declared to suffer from strong interference in two situations. Firstly, when the previous lost packet is transmitted with the basic data-rate. Here, we assume that the packet losses are nearly zero when the basic data-rate is used in the absence of interference [14], [15], [10]. Secondly, when the packet losses reach a certain threshold ($RTS_{thresh}$). Declaring strong interference triggers enabling RTS messages for a fixed number of rounds ($K$). Enabling RTS/CTS messages allows APs to measure the medium in the absence of strong interference. Missing a CTS message is considered as an indication of strong interference. Therefore, the number of rounds before disabling the mechanism (i.e., $RTS_{round}$) is incremented by one for every missing CTS, when $RTS_{round} < K$. The link resumes using the conventional CSMA/CA when the RTS rounds are zero.

### B. Opportunistic Burst Scheduling

To manage multiple clients, the scheduling technique in Section IV-A is followed. We assume that the number of clients suffering from strong interference is less than the number of available queues at their respective AP. Packet bursting is used only if RTS/CTS messages are enabled by the mechanism described in Section V-A. Only the first packet in the burst or the packet retransmission requires RTS messages. Therefore, a single round of the mechanism in Section V-A becomes a burst of packets starting with an RTS/CTS exchange, using the bursting technique in Section III. Therefore, the NAV value forces the interfering stations to postpone their transmissions if they decode the ACK. Otherwise, the interference is considered weak and the interferers resume transmission. As a result, concurrent transmissions can be accommodated if the interference is acceptable.

### C. Evaluation

In this section, we evaluate RCBurst in four dense simulated home WLAN scenarios. The scenarios are classified based on the number of neighbouring home WLANs, which are varied between 2-5. Each home WLAN consists of 5 clients and an AP. The application data-rate is $3.5Mbps$ and the RSS profile from an AP to its clients and visa versa is strong. There are probabilities of 10% and 15% that a client receives strong interference from other home WLANs' clients or APs, respectively, and the same probabilities for weak interference. APs do not exist in the carrier sense range of each other. Each scenario is represented by 40 simulated topologies. The

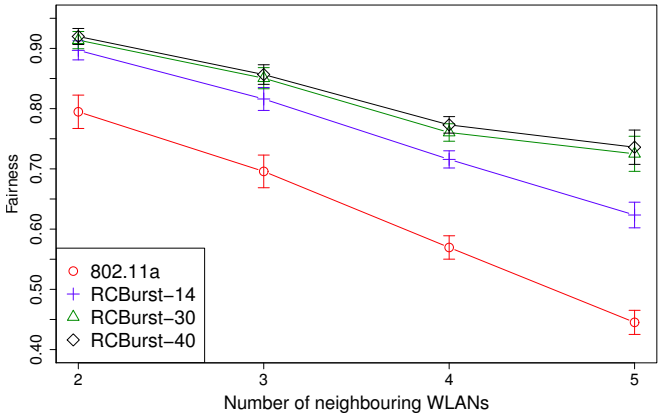


Fig. 7: The summary of Jain's fairness index.

TABLE I: The average throughput for the four scenarios in Mbps

| # of neighbouring WLANs | 802.11a | RCBurst-14 | RCBurst-30 | RCBurst-40 |
|---|---|---|---|---|
| 2 | 22.2 | 22.9 | 23 | 23.1 |
| 3 | 28.1 | 28.7 | 28.7 | 28.7 |
| 4 | 29.6 | 31 | 30.8 | 30.1 |
| 5 | 28.6 | 31.6 | 30.5 | 30.1 |

interference map is independent for each topology and it is based on the above-mentioned properties. We evaluate both the conventional 802.11a and RCBurst. In RCBurst's case, we vary the MNRA among three values; 14, 30, and 40 to measure the impact of MNRA in dense deployments. Minstrel is used as a RCA for the 802.11a and RCBurst.

Fig. 7 shows the overall fairness for the four different scenarios. The results illustrate that RCBurst improves the Jain's fairness index up to 0.30 over the 802.11a. Also, we can see that the gap between RCBurst-14 and RCBurst-{30,40} becomes wider when increasing the number of WLANs into the scenario. However, the gap between RCBurst-30 and RCBurst-40 is small for all scenarios. These observations indicate that increasing MNRA leads to a significant impact by improving fairness. However, this impact becomes trivial between large numbers of MNRA (i.e., RCBurst-30 vs. RCBurst-40). Table I presents the average throughput for the different scenarios. The throughput values are close among 802.11a and RCBurst, however, when the number of neighbouring WLANs is 5, the improvement of RCBurst-14 over 802.11 is around 10%. This improvement in throughput comes with a cost of reducing the fairness gains.

## VI. Conclusion

In this paper we explored whether it is useful to leverage a range of possible mechanisms at APs to mitigate strong interference from hidden terminals in home WLANs. Our extensive simulation study shows that a combination of many RTS attempts, packet bursting and fixed backoff improves the TCP throughput up to 8x for a link suffers from strong interference. On this basis we proposed and evaluated a mechanism called RCBurst to detect and mitigate strong interference in home WLANs. This resulted in a 0.3 improvement based on Jain's fairness index.

## VII. Acknowledgement

This publication has emanated from research conducted with the financial support of Science Foundation Ireland (SFI) and is co-funded under the European Regional Development Fund under Grant Number 13/RC/2077.

## References

[1] minstrel specification. https://sourceforge.net/p/madwifi/svn/HEAD/tree/madwifi/trunk/ath_rate/minstrel/minstrel.txt

[2] M. Al-Bado, C. Sengul, and R. Merz. What details are needed for wireless simulations? A study of a site-specific indoor wireless model. In *Proceedings of the INFOCOM*, pages 289–297, Mar. 2012.

[3] X. Chen, P. Gangwal, and D. Qiao. RAM: rate adaptation in mobile environments. *IEEE Transactions on Mobile Computing*, 11(3):464–477, March 2012.

[4] H.-J. Ju, I. Rubin, and Y.-C. Kuan. An adaptive RTS/CTS control mechanism for IEEE 802.11 MAC protocol. In *Vehicular Technology Conference, 2003. VTC 2003-Spring. The 57th IEEE Semiannual*, volume 2, pages 1469–1473 vol.2, April 2003.

[5] J. Kim, S. Kim, S. Choi, and D. Qiao. CARA: Collision-aware rate adaptation for IEEE 802.11 WLANs. In *Proceedings of the IEEE INFOCOM*, pages 1–11, April 2006.

[6] J. Manweiler, P. Franklin, and R. Choudhury. RxIP: Monitoring the health of home wireless networks. In *Proceedings of the IEEE INFOCOM*, pages 558–566, March 2012.

[7] A. Patro and S. Banerjee. Outsourcing coordination and management of home wireless access points through an open API. In *Computer Communications (INFOCOM), 2015 IEEE Conference on*, pages 1454–1462, April 2015.

[8] A. Patro, S. Govindan, and S. Banerjee. Observing home wireless experience through WiFi APs. In *Proceedings of the 19th Annual International Conference on Mobile Computing & Networking*, MobiCom '13, pages 339–350, New York, NY, USA, 2013. ACM.

[9] I. Pefkianakis, Y. Hu, S. H. Wong, H. Yang, and S. Lu. Mimo rate adaptation in 802.11n wireless networks. In *Proceedings of the Sixteenth Annual International Conference on Mobile Computing and Networking*, MobiCom '10, pages 257–268, New York, NY, USA, 2010. ACM.

[10] I. Pefkianakis, H. Lundgren, A. Soule, J. Chandrashekar, P. Le Guyadec, C. Diot, M. May, K. Van Doorselaer, and K. Van Oost. Characterizing home wireless performance: The gateway view. In *Computer Communications (INFOCOM), 2015 IEEE Conference on*.

[11] T. Shigeyasu, M. Akimoto, and H. Matsuno. Throughput improvement of IEEE802.11DCF with adaptive RTS/CTS control on the basis of existence of hidden terminals. In *Complex, Intelligent and Software Intensive Systems (CISIS), 2011 International Conference on*.

[12] B. Shrestha, E. Hossain, and K. W. Choi. Distributed and centralized hybrid CSMA/CA-TDMA schemes for single-hop wireless networks. *IEEE Transactions on Wireless Communications*, 13(7):4050–4065, July 2014.

[13] V. Shrivastava, N. Ahmed, S. Rayanchu, S. Banerjee, S. Keshav, K. Papagiannaki, and A. Mishra. CENTAUR: Realizing the full potential of centralized wlans through a hybrid data path. In *Proceedings of the 15th Annual International Conference on Mobile Computing and Networking*, MobiCom '09, pages 297–308, New York, NY, USA, 2009. ACM.

[14] K. Sui, Y. Zhao, D. Pei, and L. Zimu. How bad are the rogues' impact on enterprise 802.11 network performance? In *Computer Communications (INFOCOM), 2015 IEEE Conference on*, pages 361–369, April 2015.

[15] S. Sundaresan, N. Feamster, and R. Teixeira. Measuring the performance of user traffic in home wireless networks. In *Passive and Active Measurement - 16th International Conference, PAM 2015, New York, NY, USA, March 19-20, 2015, Proceedings*, pages 305–317, 2015.

[16] S. H. Y. Wong, H. Yang, S. Lu, and V. Bharghavan. Robust rate adaptation for 802.11 wireless networks. In *Proceedings of the 12th Annual International Conference on Mobile Computing and Networking (MobiCom)*, pages 146–157, 2006.