Jones, W., Wilson, R. E., Doufexi, A., & Sooryabandara, M. (2020). A Pragmatic Approach to Clear Channel Assessment Threshold Adaptation and Transmission Power Control For Performance Gain in CSMA/CA WLANs. *IEEE Transactions on Mobile Computing*, 19(2), 262-275. <https://doi.org/10.1109/TMC.2019.2892713>

Peer reviewed version

Link to published version (if available):
[10.1109/TMC.2019.2892713](https://doi.org/10.1109/TMC.2019.2892713)

[Link to publication record in Explore Bristol Research](#)
PDF-document

This is the accepted author manuscript (AAM). The final published version (version of record) is available online via Institute of Electrical and Electronics Engineers (IEEE) at <https://ieeexplore.ieee.org/document/8611215> . Please refer to any applicable terms of use of the publisher.

University of Bristol - Explore Bristol Research

General rights

This document is made available in accordance with publisher policies. Please cite only the published version using the reference above. Full terms of use are available:
<http://www.bristol.ac.uk/red/research-policy/pure/user-guides/ebr-terms/>

A Pragmatic Approach to Clear Channel Assessment Threshold Adaptation and Transmission Power Control For Performance Gain in CSMA/CA WLANs

William Jones, *Student Member, IEEE*, R Eddie Wilson, Angela Doufexi, *Member, IEEE*, Mahesh Sooriyabandara, *Member, IEEE*

Abstract—We propose practical sets of rules for adapting Clear Channel Assessment (CCA) and Transmit Power (TP) parameters by simulating ensembles of randomly generated wireless networks and collecting throughput statistics. The rules' performances depend strongly on network topology, with increases in throughput in many cases. However, networks with a high clustering coefficient are often adversely effected by the adaptations. But simulations of small-scale networks show that apparently adverse adaptations may still yield benefits for uneven demand combinations. Finally, we have found that throughput is not usually correlated with the number of hidden or exposed nodes in any non-trivial network set-up.

Index Terms—Network, Capacity Region, Clear Channel Assessment Threshold, Transmission Power Control.

1 INTRODUCTION

THE densification of carrier sense multiple access with collision avoidance (CSMA/CA) wireless local area networks (WLANs) has helped improve coverage, however, it has also increased interference and led to poor spatial reuse [1]. One idea, to avoid simultaneous transmissions in neighboring cells being suppressed unnecessarily, is to adapt the clear channel assessment (CCA) threshold.

The opportunity is best explained by the simple example of Fig. 1 consisting of four nodes, two Clients ($C_{1,2}$) and two Access Points ($AP_{1,2}$). Cell 1 (C_1 and AP_1) is adjacent to Cell 2 (C_2 and AP_2). C_1 and C_2 sense each others' transmissions and will not transmit simultaneously even though their signals as received at the APs would not interfere with each other (exposed nodes [2]). If the Clients were to adjust their CCA threshold sensitivity from the legacy [3] (i.e., fixed original value before any adaptation of the parameter may take place - black solid circles in Fig. 1) to a higher threshold (less sensitive value, green dashed circles in Fig. 1) they would be deaf to each others transmissions and thus could transmit simultaneously.

Alternatively, parallel transmissions can be achieved by adapting the transmission power (TP) of the clients (maintaining legacy CCA threshold). Again considering Fig. 1, if C_1 and C_2 modify their TP from the legacy (now represented by the black solid circles) to a reduced value (the green dashed circles) the transmissions are too quiet for the clients to sense each other but the APs can still receive them.

William Jones is a PhD student supported by the University of Bristol EPSRC funded Industrial Doctorate Centre in Systems (Grant EP/G037353/1) in collaboration with Toshiba Research Europe Limited (wj14134@bristol.ac.uk). R Eddie Wilson is Professor of Intelligent Transport Systems at the University of Bristol.

Angela Doufexi is Professor in Wireless Networks at the University of Bristol. Mahesh Sooriyabandara is Associate Managing Director of Toshiba Research Europe Limited in Bristol, UK.

Fig. 1. Basic CCA threshold and TP adaptation. Large black solid circles surrounding each node represent legacy CCA threshold or TP. Blue and green dashed circles represent reduced CCA sensitivity/TP for APs and Clients respectively. The diagram illustrates the opportunity to prevent suppression of transmissions in adjacent cells by reducing CCA sensitivity or TP.

An extensive survey by Thorpe [4] summarizes the CCA literature. Similarly, Chiang's book [5] provides a comprehensive review of TP control research. Much of the work discussed in [4] and [5] focuses on finding the best possible solution for one particular fixed network topology. Uncoordinated optimization of the CCA and TP parameters across a high number of nodes is unlikely to ever reach a true system optimum throughput: with all nodes aiming to maximize their own throughput, a select number will dominate the channel at the expense of others. Rather, we define and investigate a set of heuristic rules for adapting CCA threshold and TP and we evaluate their efficiency by collecting throughput statistics.

This paper is organized around a simulation study in two parts as follows. In Part A we propose a set of rules

to adapt the CCA threshold, the TP or a combination of the two, and we study an ensemble of randomly generated networks at legacy CCA and TP compared with each of the rules applied, at saturated traffic conditions (Sec. 2). We analyze the results (Sec. 2.1) and show that adapting TP alone typically results in the most significant throughput gain and often improves the fairness [6] of the network. Further, we observe that networks with a low clustering coefficient are most likely to achieve a throughput gain (Sec. 2.2). We show that the average throughput gain, for some rules, is higher when only applied to networks with a low clustering coefficient [7], [8]. To better understand these observations and the impact of the proposed rules, in Part B we study two simple networks of five nodes (Sec. 3) applying similar CCA and TP rules as proposed in Part A (outlined in Sec. 3.1). Due to the reduced number of nodes we are able to explore a larger parameter space: instead of focusing on equal saturated demands, we investigate vectors of demand that can be met by the system (Sec. 3.2), and present our results in terms of the capacity region [9] (Sec. 3.3). This metric allows us to analyze our findings (Sec. 3.4) and identify gains in performance at asymmetric demand combinations. Further, it provides an improved understanding of the relationship between throughput performance and fairness, and how the rules impact these.

The key outputs of this paper (discussed in Sec. 4) are an evaluation of the performance of the proposed rules, and an improved understanding of the capability of CCA and TP adaptation, in terms of their ability to reduce cross cell interference, improve spatial reuse and increase system throughput. Our work highlights that adapting CCA threshold or TP can impact nodes' ability to carrier sense each other, thus improving spatial reuse. By adapting TP, simultaneous transmissions are enabled which would previously not have been possible due to interference at receivers. In networks with large numbers of nodes, uncoordinated adaptation of the CCA and TP parameters is unlikely to reach a true system optimal throughput. However, we have shown that simple heuristic rules can bring some performance gains. We discuss how the approach may be implemented in dynamic networks.

2 PART A: RANDOMLY GENERATED NETWORKS

We consider a physical set-up where the nodes (clients and access points) are positioned within a square subset $[0, L] \times [0, L]$ of the (x, y) plane. Firstly, clients C_i , with $i = 1, 2, \dots, n$, are uniformly randomly distributed across the square, which is divided into cells organized around four access points AP_j , $j = 1, 2, 3, 4$, positioned at $(1/2, 1/2) \pm (1/4, 0) \pm (0, 1/4) + \mathbf{e}_j$, where the $\mathbf{e}_j$ are independently drawn random variables from the uniform distribution on $[-L/10, +L/10] \times [-L/10, +L/10]$ (see Fig. 2). The idea is that the APs are (roughly) symmetrically positioned, but the clients might be spread quite unevenly.

Clients i and APs j are then allocated default values for their CCA and TP. The simplified approach that we follow (see [10], [11]) is to model CCA sensitivity and TP as separate distance thresholds, rather than as a single power density threshold experienced by the receiver. Specifically, we suppose that a node N_A (client or AP) is assumed able

Fig. 2. Example of a randomly generated network, the subset $[0, L] \times [0, L]$ of the (x, y) horizontal plane (large black box). The rectangle is divided into four cells organized around $AP_{1,2,3,4}$ (blue triangles) at $[L/4, L/4] \pm \mathbf{e}$, $[3L/4, L/4] \pm \mathbf{e}$, $[L/4, 3L/4] \pm \mathbf{e}$, $[3L/4, 3L/4] \pm \mathbf{e}$, where $\mathbf{e}$ is in each case an independently drawn random variable from the uniform distribution on $[-L/10, +L/10] \times [-L/10, +L/10]$ (small dashed boxes). $C_{1,2,\dots,5}$ (small red circles) are randomly distributed across the rectangle and four cells. Each client and AP is at the center of a circle representing its CCA threshold and TP. To preserve clarity the legacy CCA and TP of the APs (blue dashed and blue solid circles respectively centering at each AP) and the legacy CCA and TP of just C_1 (large red dashed and large red solid circles respectively centering at C_1) are shown. It is clear C_1 could be heard by both AP_1 and AP_2 at legacy conditions.

to transmit to, or interfere with, another node (client or AP) N_B if

$$\|\mathbf{x}_{N_A} - \mathbf{x}_{N_B}\| < TP_{N_A}, CCA_{N_B} \quad (1)$$

i.e., the receiver must be within 'transmission range' of the transmitter, and the transmitter must be within an 'audible range' of the receiver. In fact, (1) simplifies the true state of affairs, which in these terms is

$$\|\mathbf{x}_{N_A} - \mathbf{x}_{N_B}\|^2 < TP_{N_A} CCA_{N_B}. \quad (2)$$

This is because power density at the receiver scales with transmit power TP_{N_A} and the inverse square of the separation $\|\mathbf{x}_{N_A} - \mathbf{x}_{N_B}\|$ of the transmitter and receiver. Note further that CCA sensitivity measured as a distance should scale reciprocally with CCA measured as a power density threshold. The difficulty is that (1) and (2) are not equivalent in situations where both TP_{N_A} and CCA_{N_B} are varied, because (1) implies (2) but not vice versa. However, the approximate formulation (1) enables an intuitive graphical understanding of range that (2) does not, that proves especially useful in Sec. 3. In addition, we have tested the accuracy of this simplification in terms of the similarity of the binary collision matrices they generate for our systems and found that over 99% of the matrices agree. Our simplified approach is thus to proceed with (1) as an exact model.

For simplicity, we consider a fixed legacy CCA sensitivity range of $4L/10$ for all nodes and we suppose the legacy TP value is drawn from the uniform distribution on $[3L/10, 4L/10]$.

Client C_i then associates to the closest access point AP_j with

$$j = \arg \min_k \|\mathbf{x}_{AP_k} - \mathbf{x}_{C_i}\| \quad (3)$$

such that

$$\|\mathbf{x}_{AP_k} - \mathbf{x}_{C_i}\| < TP_{C_i}, CCA_{C_i}, TP_{AP_k}, CCA_{AP_k}, \quad (4)$$

so that the access point and client can hear each others' transmissions bi-directionally. In the case where there is no access point in range, the client is randomly re-positioned and the association considered again. We speak of a 'link' or an 'edge' (in graph theory language) being formed between C_i and AP_j along which transmissions occur.

The study area dimensions and CCA/TP values are selected to model a four cell network, where cell borders overlap and hence are likely to be subject to the detrimental effects of cross cell interference. For example, by considering the example network of five clients shown in Fig. 2, it is clear that C_1 will be audible to both AP_1 and AP_2 at legacy CCA and TP.

If two clients attempt to transmit to the same AP simultaneously, a collision will occur and one or both transmissions will fail. Further, if the received interference from a transmitter in a neighboring cell is too great at an AP, a transmission may fail. Thus we use a n by n collision matrix E of edges to describe each operational scenario, with entries $E_{ij}, 1 \leq i, j \leq n$. If $E_{ij} = 1$, then edge i is incompatible (cannot transmit simultaneously) with j , whereas if $E_{ij} = 0$ then edge i is compatible (can transmit simultaneously) with j . In our examples, interference (and hence the collision matrix) is not always symmetric: that is, if edge i collides with edge j , in some cases only one edge transmission may fail, i.e., $E_{ij} = 1$ but $E_{ji} = 0$.

If C_1 (say) could hear C_2 (say) transmitting, C_1 would wait for C_2 to finish before transmitting itself. C_1 could hear C_2 if $\|\mathbf{x}_{C_1} - \mathbf{x}_{C_2}\| < TP_{C_2}, CCA_{C_2}$. Transmitters do not have a global knowledge of the network and therefore are not necessarily aware of transmissions on other edges with which they may collide (hidden node problem [2]) or edge transmissions which suppress them from transmitting, with which they could successfully transmit in parallel (exposed nodes [2]). The 'knowledge of the network' edge matrix F is an n by n edge matrix that we use to describe each edge's awareness of the network around it, with entries $F_{ij}, 1 \leq i, j \leq n$. If $F_{ij} = 1$, then edge i is aware of j and incompatible (cannot transmit simultaneously), whereas if $F_{ij} = 0$ then edge i is either compatible (can transmit simultaneously) or unaware of j (cannot hear edge j 's transmissions). As with the collision matrices in our examples, the network topologies (and hence the matrix F) are not symmetric: that is, if edge i has knowledge of edge j , then edge j does not necessarily have knowledge of edge i i.e., $F_{ij} = 1$ but $F_{ji} = 0$.

We propose five different adaptation rules as follows which we apply in turn:

- R1 All nodes of all cells have legacy CCA threshold value and legacy TP.
- R2 All nodes of all cells reduce their CCA sensitivity (increasing threshold) to the minimum value at which they can still carrier sense all nodes of their

cell that they were able to at legacy CCA value. A 5% margin is added. Legacy TP for all nodes.

- R3 All nodes of all cells reduce their TP to the minimum value at which they can still transmit to, and be carrier sensed by, all nodes of their cell they were able to at legacy TP. A 5% margin is added. Legacy CCA for all nodes.
- R4 All nodes of all cells reduce their CCA sensitivity by half the potential difference to the minimum value at which they can still carrier sense all nodes of their cell which they were able to at legacy CCA value. Simultaneously, all nodes of all cells reduce their TP by half the potential difference to the minimum value at which they can still transmit to, and be carrier sensed by, all nodes of their cell which they were able to at legacy TP. A 5% margin is added to both parameters.
- R5 At random, cells choose to reduce their CCA sensitivity (as R2) or reduce their TP (as R3).

We generate 10,000 networks of n clients with legacy TP and CCA threshold values (i.e., R1). We consider $n = 5, 10, \dots, 50$, hence ten network densities and therefore 100,000 networks each defining a pair of matrices E and F . Using the simplified CSMA/CA Medium Access Control (MAC) model described and used in our previous work [12], we simulate data traffic to each of these networks and collect throughput statistics. Further developing each of the legacy networks, we, in turn, apply the described CCA threshold and TP adaptations (R2-R5) generating four further E and F matrices pairs to capture them. Again, we simulate data traffic to each of these. In total this produces 500,000 pairs of matrices E and F (i.e., 500,000 individual simulations). Matrices E and F are fixed for each individual simulations.

The CSMA/CA MAC protocol [13] requires completion of a two-way handshake in order for one transmission to be successful. We are interested in understanding the extent of hidden node [2] issues when multiple clients unable to carrier sense each other attempt to transmit to the same AP. Similarly, we seek to gain insight into the impact of exposed nodes [2] when clients in neighboring cells suppress each other's transmissions. In our particular topological setup where AP positions are restricted, down-link transmissions are less prone to these issues. Hence we focus our study and consider only up-link demand, that is APs do not transmit data to clients, just clarification that the client is clear to send (CTS) and acknowledgement (ACK) of received packets. The general findings are however applicable to both up and down-link. We suppose that there is no routing and that each packet has a journey that consists of a single hop. It follows that we may prescribe demand (number of packets per unit time) on our network in terms of requested flows $d_1, d_2, \dots, d_n$ for each edge respectively. Thus each client is assumed to have a single queue of packets awaiting transmission, with Poisson arrivals. For each simulation, demand to all transmitters is set at d_{sat} , i.e., the maximum transmission rate of a single transmitter transmitting alone in a clear channel. The duration of each simulation is set at $100,000/d_{\text{sat}}$ (i.e., the time needed to transmit 100,000 packets without collisions).

Each network and the five rules under which that net-

work is simulated are compared relative to each other. We measure the total number of packets successfully transmitted and the fairness (based on Jain's fairness index [6]) under each of those five rules for the same (saturated) traffic demand and time period. The results that follow (Sec. 2.1) are based on comparisons between rules R1-R5 applied in turn to each of the 10,000 networks for n clients.

2.1 Randomly Generated Networks Results and Analysis

For each network, we explore the impact on throughput of each adaptation relative to the legacy condition (i.e., R2-R5 relative to R1 described earlier). Fig. 3 illustrates this, plotting four bars for each n clients indicating the percentage of adaptations to the 10,000 legacy networks simulated that resulted in an increased or decreased throughput change. The colored section of the bar indicates throughput change greater than 10%.

Considering Fig. 3, CCA threshold adaptation (as R2) is shown to have an approximately even probability of generating an increase or decrease in throughput and only a small number of the increases or decreases achieve greater than 10% change. Adapting TP (as R3) showed a higher probability of a throughput gain than loss in all cases, most significantly at five & ten clients. A correlation for the proportion of TP results showing an improvement in throughput decreasing with increasing network density can clearly be seen. A similar correlation can be seen when combining CCA and TP as (R4). This combination only realized higher probability of a throughput gain for five and ten clients, whereas greater numbers of clients saw a higher probability of a loss. Applying CCA and TP adaptation randomly to different cells (as R5) caused an increase in throughput in more cases than it caused a decrease relative to the legacy for all network densities tested, however, the improvement was only marginal. The most significant improvements occurred at five clients and between thirty to forty-five clients.

Fig 4 plots the change in fairness [6] and throughput for each of the adaptations (R2-R5) relative to the 10,000 legacy (R1) networks of thirty clients. The example is typical for any number of clients. When the applied adaptations produce a throughput performance improvement, fairness often decreased, and when fairness increased the network throughput often decreased. The implication is that the adaptation is helping some clients achieve a better throughput, potentially at the expense of others.

The range of impact of each particular adaptation relative to the legacy is visible in Fig. 4. The CCA threshold adaptation results are concentrated, showing the smallest range of throughput and fairness change. The result of CCA and TP adaptation randomly applied to different cells show an increased range from CCA alone, but both the TP and CCA/TP combined adaptations show much greater variation, with the potential for significant gains in throughput ($\approx 1400\%$ increase is the most extreme example). Fig. 4 reflects the indications of Fig. 3 which shows CCA adaptation leading to an increase in throughput on an almost equal number of occasions as it led to a decrease. The figure shows TP adaptation to achieve the biggest change in throughput and the same can be shown for all numbers of clients.

Fig. 5 illustrates for each number of clients the percentage of adaptations to the 10,000 legacy networks simulated that resulted in an increased or decreased fairness change (R2-R5 relative to R1). The colored section of the bar indicates a change greater than 10%. As with throughput, CCA adaptation alone was shown to have an approximately equal probability of achieving a fairness gain or loss. All other adaptations typically led to a fairness improvement on more occasions than a loss, with TP achieving improved fairness most frequently up to thirty-five clients. For forty and forty-five clients TP and combined CCA/TP achieve a fairness improvement on an approximately equal number of occasions. At fifty clients combined CCA/TP most frequently achieved a fairness improvement.

A small number of simulations show very extreme changes from the legacy when an adaptation is applied (see Fig. 4). Fig. 6 shows the 80th percentile range of throughput changes as a result of the applied adaptations and indicates the mean and median. CCA adaptation shows a small range ($\approx 15\%$) of variations equally distributed positively and negatively either side of zero performance change, reflecting the indications of Fig. 3. The TP adaptation shows the potentially highest increase in throughput in Fig. 6 for all numbers of clients. The two methodologies that combine CCA and TP adaptation have a similar range of impacts. However, Fig. 3 illustrates how differently they perform with combined CCA/TP in each cell achieving a throughput loss on more occasions than that of randomly applying CCA or TP adaptation to different cells. TP is shown to be the most beneficial adaptation for all networks with the most significant throughput improvement at lower density.

Fig. 7 shows the total reduction in CCA and TP as a result of the applied adaptations. This figure should be considered in parallel with Figs. 3, 5 and 6. The percentage reduction equates to the cumulative amount by which the radius of the two circles surrounding each node in the model, defining their CCA sensitivity and TP, are reduced from the legacy values. The average change shows a decrease for each adaptation, tending to zero with increasing client density. The figure illustrates the limited change these adaptations are able to achieve as network density is increased, however Figs. 3 and 6 indicate that this small adaptation can make a significant impact.

2.2 Analysis of Matrix Characteristics

The matrices (E and F) underlying each simulation, for all numbers of clients, were analyzed to identify network characteristics contributing to the performance impact of the applied adaptations (R2-R5 relative to R1 in Sec. 2). The matrices (represented in Fig. 8) were ordered by their associated AP. The four square clusters visible along the diagonal represent the clients associated with the four access points. The size of each cluster indicates the number of clients associated with that access point. Marked points away from the diagonal indicate clients on the border of two cells. Asymmetry, number of hidden nodes, number of suppressed transmitters, clustering coefficient [7], [8] and other measures were considered.

For all E matrices, prior to any adaptation being applied (legacy conditions (R1)), the mean local clustering coefficient

Fig. 3. Performance impact of applied modifications. Bars indicate the number of occasions (%) that the modification led to an increase or decrease in throughput. The colored sections of the bars indicate a change greater than 10%. For all numbers of clients, TP adaptation is shown to be most likely to give a performance improvement. Combined CCA/TP within a cell is shown to become increasingly likely to result in a performance loss with increasing numbers of clients. Other scenarios show roughly equal chance of a gain or loss for all client numbers.

Fig. 4. Impact of adaptation on throughput and fairness relative to legacy for 30 Clients

was calculated [7], [8]. For each of the four adaptations (R2-R5) the throughput change achieved after the adaptation was ordered from highest to lowest. These were then divided into ten equal-sized groups i.e., the top 10% of performance improvements to the lowest 10% in intervals of 10%. For each of these ten groups the mean cluster coefficient (of the legacy networks) and mean throughput performance improvement were plotted against each other (see Figs. 2.2 and 2.2). A correlation shows the percentage change in throughput decreases with increasing clustering coefficient. We can approximately identify a clustering coefficient value for each adaptation and network density above which the

application of each adaptation (R2-R5) has a negative effect on throughput from the legacy (R1). This same trend can be identified by comparison of throughput change with E or F. For reference we further show the relationship between density of matrix E, i.e., proportion of the collision matrix filled with 1's, against throughput change (Figs. 2.2 and 9d) to demonstrate there is no clear trend, confirming the clustering coefficient is not simply a factor of the collision matrix density. The same correlation we show, for twenty-five and forty-five clients, can be shown similarly in the other size networks investigated. Similar processes were conducted comparing the number of hidden or exposed nodes in each network with the throughput improvement, however, no trend could be identified.

The identified relationship between clustering coefficient and throughput change (Fig. 9) now allows us to better benefit from TP and CCA adaptations. We re-analyze our set of 100,000 randomly generated legacy networks (R1) (generated as described in Sec. 2) only implementing the adaptation (R2-R5) for a particular number of clients if the cluster coefficient is below a defined threshold, estimated from Figs. 2.2 and 2.2 (and equivalent figures for other numbers of clients) as the clustering coefficient where the plotted (assumed linear correlation) change in throughput values approximately equals zero. Above that threshold, we maintain legacy conditions and therefore the number of simulations to which the adaptation was applied was reduced. Of those to which the adaptation was applied, for twenty to fifty clients an improvement in accuracy was observed (i.e., the proportion of times the implementation led to an improvement in throughput relative to the legacy when

Fig. 5. Fairness impact of applied modifications. Bars indicate the number of occasions (%) that the modification led to an increase or decrease in fairness. The colored sections of the bars indicate a change greater than 10%

Fig. 6. Performance improvement/reduction range for the 80_{th} percentile, for all adaptation types. For each adaptation type, the group of plots shows 10 colored pairs of plots for each number of clients, $n = 5, 10, \dots, 50$ from left to right. For each range, the circle shows the median value, the black cross shows the mean. The left plot of each pair shows the initial result. The right plot of each pair shows the revised range, mean and median when the adaptation is applied only to networks with clustering coefficient below a threshold value (see Sec. 2.2). The potential improvement using the clustering coefficient is visible for 20 to 50 clients.

Fig. 7. Mean power reduction for all networks for all adaptation types, normalized for the number of nodes. The plot shows the combined percentage by which both circles surrounding each node defining CCA sensitivity and TP are reduced in the model when each adaptation is applied.

Fig. 8. An example plot of $\mathbf{F}$ and $\mathbf{E}$ matrix for 50 edge topology. Blue dots represent a one, no dot represents a zero.

implementing was improved from that shown in Fig. 3). No improvement could be shown for five, ten or fifteen clients. For twenty to fifty clients, the mean performance change of TP and the two combined CCA and TP adaptations were improved, most significantly at higher numbers of clients. The revised performance change range, mean and median are added to Fig. 6, and the improvement is clearly visible with the number of wrong implementations reduced and the average performance changes improved.

3 PART B: SIMPLE TOPOLOGY NETWORKS

To better understand the effects observed in some of the large-scale randomly generated networks, the impact of CCA threshold and TP adaptation are further investigated on two highly simplified networks (see Figs. 10 and 11). The two networks are similar to those considered in [1] consisting of two APs and three Clients. The only difference between the two is as follows. In Fig. 11, C_1 is positioned slightly further away from AP_1 and C_2 than in Fig. 10. This makes no difference to the clients association with APs, but the change impacts on C_1 and AP_1 's ability to adapt their CCA threshold or TP. For both networks (Figs. 10 and 11) we investigate, as with the larger network examples, the application of CCA threshold and TP adaptation and a

combination of the two. The study that follows highlights the topological constraints of the interference management techniques and helps us to understand the results presented in the previous sections.

We make some basic assumptions about both networks:

- C_1 and C_2 are associated with AP_1 . These three nodes make up Cell 1. C_3 is associated with AP_2 . These two nodes make up Cell 2.
- As for the random network examples, we consider only up-link demand from Clients to APs. APs do not transmit data to clients, just CTS and ACK. Thus, the network topologies can be captured by a series of three edges; connecting C_1 to AP_1 , C_2 to AP_1 and C_3 to AP_2 . Again we consider no routing and we suppose that each packet has a journey that consists of a single hop. Thus, we may prescribe demand d_A , d_B , d_C for each of the three edges respectively. Hence, as for the random network example, we use a collision matrix $\mathbf{E}$ (see Fig. 12) and knowledge of the network matrix $\mathbf{F}$ (see Fig. 13) to describe each operational scenario, and these are now just three-by-three in size.
- At legacy CCA threshold and TP values, C_1 can carrier sense transmissions from AP_1 and C_2 ; C_2 can carrier sense transmissions from C_1 and C_3 and AP_1 and AP_2 ; C_3 can carrier sense transmissions from C_2 , AP_1 and AP_2 ; AP_1 can carrier sense transmissions from C_1 , C_2 and C_3 ; AP_2 can carrier sense transmissions from C_2 and C_3 .
- C_1 and C_2 transmitting simultaneously, at legacy conditions, will result in a collision causing both transmissions to fail. A collision will occur as both clients have the same destination AP which is only capable of receiving one transmission at any one time. Due to the received signal strength from both clients at the AP being similar, both mutually cause interference too great to allow the other transmission to be heard.
- C_2 and C_3 transmitting simultaneously, at legacy conditions, will result in a collision causing both transmissions to fail. Despite the destinations of C_2 and C_3 transmissions being different (AP_1 and AP_2 respectively) the interference received at each AP (say AP_1) from the client whose destination is not that AP (C_2) is roughly similar to the received signal from the client whose intended destination is that AP (C_1) resulting in a signal to interference noise ratio (SINR) of the transmission intended for that AP with too high an interference level for the AP to interpret the message.
- Prior to any optimization, due to the position of C_1 and C_3 on opposite sides of AP_1 out of carrier sensing range of each other, a hidden node issue exists [14] at AP_1 . C_1 and C_3 transmitting simultaneously, at legacy conditions, will result in a collision causing the transmission from C_1 only to fail. Despite the destination of C_1 and C_3 transmissions being different (AP_1 and AP_2 respectively) the interference received at AP_1 from C_3 is roughly similar to the received signal from the C_1 . This results in a SINR

Fig. 9. The mean clustering coefficient and density of $\mathbf{E}$ matrices compared to mean throughput performance improvement for the top 10% of performing networks to the lowest 10% of performing networks in intervals of 10%.

ratio at AP_1 where the interference of C_3 is too high for C_1 's transmission to be received. AP_2 cannot hear the transmissions of C_1 and so the transmissions of C_3 are not interrupted. This asymmetry is reflected in matrix $\mathbf{E}$ in Fig. 12a, i.e., $E_{13} = 1$ but $E_{31} = 0$. The hidden node effect is captured by the difference in the matrices $\mathbf{E}$ and $\mathbf{F}$ (legacy CCA values) of Figs. 12a and 13a respectively showing that edges do not have awareness of a potential collision, i.e., $E_{13} = 1$ but $F_{13} = 0$.

- Clients and APs would not knowingly attempt to reduce their CCA sensitivity to become deaf to other nodes in their cell as this would create a hidden node issue [14] which is known to have a negative impact on performance. Similarly, when optimizing transmission power, nodes would not reduce their transmission power such that others in their cell could not hear them, as this too would lead to a hidden node issue.

We apply the same adaptation techniques we applied to the randomly generated networks to the simple smaller networks. The operational scenarios are described in table 3.1 with matrices $\mathbf{E}$ and $\mathbf{F}$ show in Figs. 12 and 13 respectively.

Fig. 10. Simple topology 1 with CCA threshold adaptation. Large black solid circles surrounding each node represent legacy CCA threshold or TP. Blue and green dashed circles represent APs and clients respectively.

3.1 Simple Topology Operational Scenarios

We identify all possible optimizations using CCA threshold, TP adaptation or a combination of the two capable of impacting the link structure of our simple networks (Figs.

Fig. 11. Simple topology 2 with CCA threshold adaptation. Large black solid circles surrounding each node represent legacy CCA threshold or TP. Blue and green dashed circles represent reduced CCA/TP for APs and clients respectively. C_1 is slightly further away from AP_1 than in Fig. 10, therefore reducing Cell 1's ability to adapt CCA/TP.

10 and 11), i.e., adaptations that will change the edge interference and/or edge carrier sensing ability from those of the legacy CCA threshold and TP. Table 3.1 summarizes the legacy and then eight possible combinations of optimization using CCA threshold and TP adaptation. The table explains if and how these can be captured in terms of matrices **E** and **F** shown in Figs. 12 and 13 respectively.

When applying adaptation techniques (CCA, TP or combine), in some instances the resultant **E** and **F** matrix are the same for both of the simple network topologies (i.e., No.1, No.4 and No.7 in table 3.1). For others, it is only possible to apply the optimization to one topology (i.e., No.3, No.6, No.8 and No.9 in table 3.1). In other cases, applying different optimization methods result in the same **E** and **F** (i.e., No.4 topology 1 & 2 and No.2 topology 2, No.2 topology 1 and No.5 topology 1, No.5 topology 2 and No.7 topology 1 & 2 in table 3.1). Only eight possible unique combinations of **E** and **F** (in Figs. 12 and 13) exist, capturing all possible optimization combinations described in Sec. 3 and summarized in table 3.1. These are labeled (S1-S8) in table 3.1.

3.2 Simple Topology Simulation Methodology

Our investigation simulates each of the Scenarios S1-S8 (see table 3.1) using the simplified MAC model as Sec. 2 and our previous work [12]. Due to the simple examples' significantly reduced parameter space, for each three-edge scenario, rather than focus on fixed equal demands, we perform a large ensemble of simulations, each with a different demand vector (d_A, d_B, d_C) . Specifically, as before, we let d_{sat} be the maximum transmission rate of a single client transmitting alone in a clear channel. Each of the demands d_A, d_B and d_C is varied independently from 0 to d_{sat} in 40 equal increments, resulting in $41^3 = 68,921$ simulations for each of the eight scenarios. The duration of each simulation is set at $10,000/d_{\text{sat}}$ (i.e., the time needed to transmit 10,000 packets without collisions).

For each individual simulation, we gather statistics on the total number of packets successfully transmitted, the time evolution of queues, and the average latency per packet. These statistics are post-processed with various heuristics to decide whether each simulation is within the capacity region, meaning the network and protocol can meet the prescribed demand; or outside the capacity region, meaning that in the large time limit, latency and at least one queue grows without bound.

The results that follow (Sec. 3.3) are based on comparisons between Scenarios S1-S8. The proportion of simulated demand vectors that are within the capacity region and asymmetry of those vectors provide interesting insight.

3.3 Simple Topology Results

For each of the eight scenarios, the simulation results can be presented in the form of a three-dimensional scatter plot, see Fig. 14, where red markers indicate combinations of demand that are within the capacity region. In this plot $\delta_A := d_A/d_{\text{sat}}$, $\delta_B := d_B/d_{\text{sat}}$ and $\delta_C := d_C/d_{\text{sat}}$ denote non-dimensional demand intensities that range from 0 to 1. The concavity of the capacity region (the volume covered by red dots) is apparent from such plots and represents the loss of efficiency in the channel due to competition between transmitters and the resulting collisions. However, we require numerical measures that can be derived from these plots in order to compare the scenarios.

Firstly, we define total demand intensity $\delta := \delta_A + \delta_B + \delta_C$. Then sections $\delta = \text{const.}$ describe triangular cross sections through Fig. 14. One may then count (as a function of δ) the proportion of the corresponding triangular area that is within the capacity region, and we denote this quantity S . The dependence of S on δ may then be studied and compared across the eight scenarios: see Fig. 15a.

Secondly, we may consider the line $\delta_A = \delta_B = \delta_C$ along which the demands are equal. We may then identify the maximum value δ_{cap} (of $\delta_A = \delta_B = \delta_C$) which is within the capacity region, and compare across scenarios. See the caption under each sub figure in Fig. 14.

Further, we may measure the proportion of the simulations that are within the capacity region — that is, the proportion of the volume $(V) [0, 1] \times [0, 1] \times [0, 1]$ that is within the capacity region — and compare across scenarios. See the caption under each sub figure in Fig. 14.

We consider the demand combination $\delta_A = \delta_B = \delta_C = \delta_{\text{sat}}$ and identify the proportion of that demand met by each of the three respective edges δ_{met} . See the caption under each sub figure in Fig. 14.

For selected scenarios, we examine the ratio of satisfied demand to failed transmissions (on the line $\delta_A = \delta_B = \delta_C$). See Fig. 15b. From this figure the demand (δ) at which each edge reaches saturation is apparent. Further insights into the competition for limited channel access between the edges can be observed.

The point of the various measures is that they allow one to distinguish whether capacity is added by enhancing throughput in symmetric demand situations, or by allowing fresh combinations of highly asymmetric flows.

$$\begin{aligned}
 &\begin{pmatrix} 0 & 1 & 0 \\ 1 & 0 & 1 \\ 1 & 1 & 0 \end{pmatrix} & \begin{pmatrix} 0 & 1 & 0 \\ 1 & 0 & 1 \\ 0 & 0 & 0 \end{pmatrix} & \begin{pmatrix} 0 & 1 & 0 \\ 1 & 0 & 0 \\ 1 & 1 & 0 \end{pmatrix} \\
 &\text{(a)} & \text{(b)} & \text{(c)}
 \end{aligned}$$

Fig. 12. Collision matrices **E**

$$\begin{aligned}
 &\begin{pmatrix} 0 & 1 & 0 \\ 1 & 0 & 1 \\ 0 & 1 & 0 \end{pmatrix} & \begin{pmatrix} 0 & 1 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix} & \begin{pmatrix} 0 & 1 & 0 \\ 1 & 0 & 0 \\ 0 & 1 & 0 \end{pmatrix} & \begin{pmatrix} 0 & 1 & 0 \\ 1 & 0 & 1 \\ 0 & 0 & 0 \end{pmatrix} \\
 &\text{(a)} & \text{(b)} & \text{(c)} & \text{(d)}
 \end{aligned}$$

Fig. 13. Knowledge of the network matrices **F**

Table 3.1 Operational Scenarios

No.	Description	Topology 1 (Fig. 10)	Topology 2 (Fig. 11)
1	Legacy CCA Legacy Transmission Power	E = Fig. 12a F = Fig. 13a (S1)	E = Fig. 12a F = Fig. 13a (S1)
2	Max allowable Reduced CCA Sensitivity all nodes Legacy Transmission Power	E = Fig. 12a F = Fig. 13b (S2)	E = Fig. 12a F = Fig. 13c (S4)
3	Reduced CCA Sensitivity C_1 , C_2 and AP_1 only Legacy Transmission Power	E = Fig. 12a F = Fig. 13d (S3)	C_2/AP_1 cannot reduce CCA and still hear all nodes in its cell
4	Reduced CCA Sensitivity C_3 and AP_2 only Legacy Transmission Power	E = Fig. 12a F = Fig. 13c (S4)	E = Fig. 12a F = Fig. 13c (S4)
5	Legacy CCA Max allowable Reduced Transmission Power all nodes	E = Fig. 12a F = Fig. 13b (S2)	E = Fig. 12b F = Fig. 13d (S6)
6	Legacy CCA Reduced Transmission Power C_1 , C_2 and AP_1 only	E = Fig. 12c F = Fig. 13c (S5)	C_2 cannot reduce transmission power and still be heard by all nodes in its cell
7	Legacy CCA Reduced Transmission Power C_3 and AP_2 only	E = Fig. 12b F = Fig. 13d (S6)	E = Fig. 12b F = Fig. 13d (S6)
8	Reduced CCA Sensitivity C_1 , C_2 and AP_1 only Reduced Transmission Power C_3 and AP_2 only	E = Fig. 12b F = Fig. 13b (S7)	C_2/AP_1 cannot reduce CCA and still here all nodes in its cell
9	Reduced CCA Sensitivity C_3 and AP_2 only Reduced Transmission Power C_1 , C_2 and AP_1 only	E = Fig. 12c F = Fig. 13d (S8)	C_2 cannot reduce transmission power and still be heard by all nodes in its cell

3.4 Simple Topology Analysis

S5 and S8 were the only scenarios simulated to show any gain in capacity region volume (V) from S1 with legacy CCA threshold and TP. Both of these scenarios are defined by the same **E** matrix in Fig. 12c and their **F** matrices are shown in Figs. 13c and 13d respectively, both containing a hidden node. Due to the adaptation to the TP and CCA threshold in both S5 and S8, the edge connecting AP_2 with C_3 has a reduced risk of failure due to collision, because the TP in the neighboring cell has been reduced such that a transmission from C_3 to AP_2 (or vice-versa) will always be heard over other interference from the network. Scenario S5 showed the greatest increase in V improving from $V = 0.182$ for S1 to $V = 0.318$ for S5, narrowly ahead of $V = 0.313$ for S8. Notably for S5, δ_{cap} showed a slight decrease from S1's $\delta_{\text{cap}} = 0.34$ to $\delta_{\text{cap}} = 0.32$ whereas S8 was the same as S1.

CCA and TP adaptation can be beneficial and allow simultaneous transmissions that previously would have been prevented as illustrated in Fig. 1. In the networks we investigate in (Figs. 10 and 11) a further problem becomes apparent that adjusting nodes' ability to hear each other can increase hidden node effects. It is expected that hidden nodes have a significant negative impact to a networks' throughput [2], [14]. We have explored the impact of hidden nodes on simple networks in our previous work [12]. The asymmetry of the network and the number of hidden nodes (i.e., where the corresponding position in matrix **E** shows

a '1' but **F** shows a '0') however does not directly impact on the volume of the capacity region or δ_{met} , rather the position and asymmetry of the hidden node within the network has a far greater significance. Scenarios S2, S3, S4, S6 and S7 all showed a reduction in V and δ_{cap} . In S5 and S8 the reduced TP to Cell 1 creates a hidden node effect in that Cell 2 can now not hear the transmissions of Cell 1 but Cell 1 cannot hear Cell 2. This, however, increases the overall capacity region volume, by eliminating any competition for channel access for the edge in Cell 2, facilitating it in meeting its demand vector at certain asymmetric demand combinations. There is no correlation between the number of hidden nodes or asymmetry of the network and total capacity achieved.

It is apparent that δ_{met} does not necessarily increase with V (see Fig. 14). Most notably, although showing similar gains in V , scenarios S5 and S8 achieve very different δ_{met} with the total of S8 approximately half that of S5 and a significant reduction from S1. S5 was the only scenario to show an increased total δ_{met} from S1. Fig. 15b provides insight into the competition between edges considering equal demands. For the three scenarios displayed, the legacy (S1) and two scenarios that showed an increase in V (S5 and S8); the plot shows how the edges behave for equal demands. S8 (red markers Fig. 15b) shows the error rate increases with increased demand to the three edges. The clustering of three plotted red shapes representing the three edges shows how

Fig. 14. Capacity region plots for S1-S8. Red dots indicate a point is inside the capacity region. Here V is the volume of the capacity region and δ_{cap} is its extent along the line $\delta_A = \delta_B = \delta_C$. Further δ_{met} identifies the proportion of the demand combination $\delta_A = \delta_B = \delta_C = \delta_{\text{sat}}$ met by each edge respectively. ΔE and ΔF are the reduction in the total number of unit entries relative to legacy S1 in matrices **E** and **F** respectively i.e., colliding edge combinations and suppressed transmissions respectively. S5 & S1 result in maximum achievable V for topologies 1 & 2 respectively.

all edges saturate at similar ratios of demand to error. S1 and S5 show a different trend. Initially, as S8, the error rate of S1 and S5 edges increases with demand, however, both reach a point where as further demand is added, the error rates for edges A and C decrease. Beyond the same point, the ability for edge B to meet its demand decreases and further its error rate also decreases. This indicates that the edge has a lower probability of accessing the channel, losing out in competition for the limited resource to the other two edges. This starvation effect [2], [15] can be explained by edge B's position between the other two edges in the topology. The difference between S5 and S8 can be identified as due to the difference between their knowledge of the network, i.e., Fig. 13c and 13d respectively. The matrices contain the same number of 1s, however, the positioning indicates nodes' different knowledge of the network surrounding them and this results in a significantly higher total δ_{met} for S5.

4 DISCUSSION

The results of this study relating to CCA adaptation support existing work and are comparable in terms of the magnitude of performance gain. Paper [16] found CCA adaptation could achieve up to $\approx 20\%$ throughput improvement if implemented appropriately and further could improve fairness. As with our study [16] recognizes the potential improvement to be highly dependent on topology. The TP adaptation results are comparable to a study [17] carried out using the OPNET [18] simulator. The authors of the study were able to show significant improvements in throughput of several hundred percent in certain situations.

Through further examining modified CCA thresholds in simple networks, we were able to identify specific changes to performance stemming from the underlying 'knowledge of the network' matrix **F** in our model. Considering the topology in Fig. 10, Cell 1 nodes (C_1 and C_2 along with

Fig. 15. (a) Proportions S of triangular cross section within the capacity region, as a function of total demand intensity $\delta = \delta_A + \delta_B + \delta_C$, compared across eight scenarios. (b) A plot of demand (δ) per client against $(\epsilon\delta)$ for scenarios S1 (legacy CCA and TP), S5 (legacy CCA and reduced TP to C_1, C_2 and AP_1 only) and S8 (reduced CCA to C_3 and AP_2 only and reduced TP to C_1, C_2 and AP_1 only). Here ϵ is the failure rate i.e., the number of failed transmissions per successful transmission, along the line $\delta_A = \delta_B = \delta_C$. The triangle, square and star shapes represent edges A, B and C respectively. The colors indicate the scenario (S1) black, (S5) blue (S8) red.

AP_1) can all reduce their CCA sensitivity such that they can only hear each other and are deaf to the transmissions of Cell 2 nodes (C_3 and AP_2). Similarly Cell 2 nodes can do the same making themselves deaf to the transmissions of C_1, C_2 and AP_1 . Consider the topology in Fig. 11; because of the slight change of position of C_1 , in order for Cell 1 nodes to all hear each others' transmissions, they cannot adapt their CCA threshold by a significant margin. This simultaneously has the effect of preventing AP_1 and C_2 reducing CCA sensitivity to become deaf to C_3 . Cell 2's nodes are unaffected by the position of C_1 and so can reduce their CCA sensitivity to become deaf to the other three nodes in the network. This results in the matrix F becoming asymmetric when CCA sensitivity is reduced, such that C_2 and AP_1 can hear nodes in their neighboring cell (i.e., AP_2 and C_3) but C_3 and AP_2 are deaf to their neighbors.

Adapting the TP for any number of clients showed the largest range of performance change, potentially leading to the greatest throughput gain or the most significant loss. It was the adaptation that most consistently led to an increase in throughput performance. Whereas modifying the CCA threshold can only enable pairs of edges to simultaneous transmit that would previously have been silenced (i.e., change matrix F), modifying the TP additionally has the potential to enable pairs of edges that would have previously collided to transmit simultaneously (i.e., changing matrix E also). When all clients reduced their TP by the same proportion, the SINR at the access points remains the same and therefore matrix E stayed the same. In this case, reducing the volume of transmission from all nodes has much the same effect as reducing the CCA sensitivity at all nodes changing matrix F . More significant changes occurred when the transmission power is adjusted asymmetrically.

Adaptation of TP is subject to similar topological constraints as discussed with CCA adaptation. In the first simple network (Fig. 10) it is possible for Cell 1 to reduce TP where Cell 2 cannot hear (but still the nodes in Cell 1 can hear each other). However, in the second (Fig. 11), C_2 cannot reduce TP to a level where Cell 2 nodes cannot hear whilst still being heard by all nodes in its cell. When

C_2 reduces its TP, the SINR ratio at AP_2 changes such that when C_2 and C_3 transmit simultaneously the louder transmission from C_3 is always heard over C_2 (i.e., $E_{32} = 0$). There is still of course interference at AP_1 , such that in the case of C_2 and C_3 transmitting simultaneously, AP_1 cannot interpret its signal from C_2 and the transmission fails. The case where Cell 2 only reduces TP to the minimum value required for AP_2 to hear C_3 and vice-versa removed the risk of a potential collision between C_3 and C_1 at AP_1 (i.e., $E_{13} = 0$). Further, now the SINR ratio at AP_1 , when C_2 and C_3 transmit simultaneously, is modified such that the louder transmission from C_2 is heard over C_3 (i.e., $E_{23} = 0$). Of course interference at AP_2 occurred such that when C_2 and C_3 transmit simultaneously, AP_2 was unable to receive from C_3 .

Reducing the TP at one node effectively generated the same impact as reducing the CCA sensitivity of surrounding nodes in terms of allowing simultaneous transmission that previously would have been silenced (i.e., by changing F). When C_3 reduced its TP such that only AP_2 can just hear him, AP_1 and C_2 could not hear him with their fixed CCA legacy threshold. Hence matrix F changed to reflect this (Fig. 13d). Likewise when C_2 reduced its TP (first topology only Fig. 10) such that only AP_1 and C_1 could just hear him, AP_2 and C_3 could no longer hear him with their fixed CCA legacy threshold. Hence matrix F changed to reflect this (Fig. 13c). Because of its capability to impact both matrices E and F , adapting TP has a far more significant impact than adapting CCA threshold.

We tested randomly applying CCA or TP adaptation to cells in our randomly generated network or combining both within each cell. Random deployment is representative of networks where there is no co-ordination between separately deployed cells. As the results of the randomly generated networks demonstrated (Fig. 3), this generally provided a performance improvement, however only slight. In the simple examples, both methods impacted on the cells' ability to hear each other, changing matrix F , and collisions, changing matrix E . The results in Fig. 14g and 14h, however, show very different impacts on the capacity of the network

depending on which cell reduced CCA sensitivity or TP. This highlights the sensitivity of the adaptation's impact to topology.

In our simple networks (Fig. 10 and 11) the options for combining the CCA and TP adaptation are very limited. Simultaneously reducing CCA sensitivity of the receiver and TP of the sender risks breaking the connection between the two nodes. Of course in some situations, a balance could be reached where both transmitter and receiver are able to reduce the TP and CCA sensitivity by a margin and still maintain a connection. One might expect this may minimize their aggregate interference on the surrounding network. Our simulation of this on the randomly generated networks (i.e., R4) showed such an adaptation to increase mean throughput performance, however, the change led to a loss on a higher number of occasions than a gain. This indicates the mean is raised by a few outlier networks even though for the majority of networks, the adaptation results in a slight throughput loss.

Each node, when adjusting parameters in accordance with our prescribed rules, sets its TP and/or CCA value such that they can still carrier sense, or be carrier sensed by, the same nodes of their cell as at legacy values. In practice, this adjustment might be achieved by each node monitoring the received signal strength indication (RSSI) of all of its associated nodes, and by choosing the RSSI of the farthest and setting CCA and TP accordingly [3]. However, In a dynamically changing network, the challenge is more difficult. If a new client joins and associates with an AP, other clients in the same cell, with reduced CCA, might not carrier sense the new client and the new client may not carrier sense those existing clients with reduced TP. All clients associated with a particular cell receives a beacon frame from their respective AP at every beacon interval (typically $\approx 100\text{ms}$) [16]. Beacons inform of configuration changes notifying existing clients of new arrivals to the cell. Hence, if a client who has reduced CCA cannot carrier sense the new client, it may return to legacy settings and re-apply the rule. Alternatively, if a client of reduced TP is made aware of a new client it may alter its TP to incorporate the new arrival in accordance with the rule.

Other authors have explored individual nodes dynamically optimizing CCA and TP parameters [3], [4], [5], [16] observing 5-20% throughput improvements. In any non-trivial network it is unlikely a true system optimal throughput will be reached by such an approach. Even if nodes cooperated, searching the vast parameter space for an optimal solution will likely not be computationally tractable. Further, these approaches often result in poor fairness. The advantage of the simple rules we propose is that they restrict the parameter space, reducing the required computation, and hence might be applicable to dynamic networks. Additionally, we have identified a relationship between the clustering coefficient and the likely throughput: the potential performance improvement that adaptations can provide decreases with increasing clustering coefficient. This relationship can be explained as follows. A higher clustering coefficient indicates that the clients are grouped together tightly around the APs with fewer clients falling on the border between cells. When there are few or no clients on the border between cells, cross cell interference

cannot be improved by adapting the CCA threshold or TP. We used the clustering coefficient to determine when we should apply an adaptation and showed this to reduce the number of implementations leading to a loss in throughput, and thus to an increase in the mean throughput change. If a network's topology can be monitored and the clustering coefficient computed in real time, rules could be implemented accordingly to achieve maximum gain. Exploring different combinations of parameters from the five rules we propose and better use of the clustering coefficient and other network characteristics as a measure to predict if and when to apply them, are topics for future research.

As many other studies in related areas do, we used Jain's Fairness Index [6] which rates the fairness of a set of n nodes' throughput values where x_i is the throughput for the i^{th} connection. The metric ranges from a minimum of $1/n$ to a maximum of 1 (when all users receive the same allocation). This index is k/n when k users equally share the resource, and the other $n - k$ users receive zero allocation. The fairness impact from our simulations on randomly generated networks was reported in Fig. 5. In our simple examples, the shape of the capacity regions (Figs. 14) indicates how some highly asymmetric demand combinations, and therefore asymmetric (i.e., unfair) throughput combinations, can be met by the system where even demand combinations, of equivalent total throughput, cannot. The change in result from scenarios S1 to S5 and S8, which both showed an increase in overall volume of the capacity region, show a significant change in fairness. In these two scenarios, the edge connecting C_3 and AP_2 has a higher chance of transmitting, after the adaptation, than the other two edges. This change in fairness, however, is slightly misleading as only the ability to meet the capacity of one edge is changed. The probability of the edges connecting the two clients to AP_1 transmitting successfully remains the same for all three scenarios, the third edge's probability of successful transmission is increased without impacting on these. The two scenarios S5 and S8 that show growth in the capacity region to the legacy S1 do so by enabling additional asymmetric demand combinations to be met by the system. Comparing S1, S5 and S8 at δ_{sat} , S8 achieves the fairest throughput with similar values of δ_{met} to all clients, however it also achieves the lowest total δ_{met} . Fairness, although still a useful measure, must be considered carefully. If a network of three transmitting edges, each achieving equal channel access (i.e., fair) is adapted, via TP or CCA adaptation, such that two of those edges achieve the same throughput as before but one achieves greater, the system is now less fair although the only impact to throughput has been positive.

5 CONCLUSION

We explored CCA threshold and TP adaptation considering practical methods by which they could be implemented. The study contributes to knowledge by (1) presenting a modeling methodology for investigating CCA/TP adaptation; (2) demonstrating CCA and TP adaptation methods that result in performance improvements; (3) identifying that network throughput does not correlate with the number of hidden or exposed nodes when the network contains more than one cell and explaining how asymmetrically suppressing part of

a network can remove competition for channel access and lead to overall network throughput improvement; (4) identifying a relationship between throughput improvement and network clustering coefficient and thus network conditions under which CCA and TP adaptations are most likely to realize throughput improvements.

CCA, TP and combined CCA/TP adaptation are all capable of generating a performance improvement in WLANs, however, the impact of the adaptation is highly dependent on topology. TP adaptation provides the best chance of a performance improvement. When applied only to networks with low clustering coefficient, the probability of achieving a performance gain from TP adaptation is increased.

REFERENCES

- [1] Z. Zhong, P. Kulkarni, F. Cao, Z. Fan, and S. Armour, "Issues and challenges in dense WiFi networks," in *2015 Int. Wirel. Commun. Mob. Comput. Conf.* IEEE, aug 2015, pp. 947–951. [Online]. Available: <http://ieeexplore.ieee.org/document/7289210/>
- [2] L. B. Jiang and S. C. Liew, "Improving throughput and fairness by reducing exposed and hidden nodes in 802.11 networks," *IEEE Trans. Mob. Comput.*, vol. 7, pp. 34–49, 2008.
- [3] P. Kulkarni and F. Cao, "Dynamic Sensitivity Control to improve Spatial Reuse in Dense Wireless LANs," in *Proc. 19th ACM Int. Conf. Model. Anal. Simul. Wirel. Mob. Syst. - MSWiM '16*. New York, New York, USA: ACM Press, 2016, pp. 323–329. [Online]. Available: <http://dl.acm.org/citation.cfm?doid=2988287.2989138>
- [4] C. Thorpe and L. Murphy, "A Survey of Adaptive Carrier Sensing Mechanisms for IEEE 802.11 Wireless Networks," vol. 16, no. 3, pp. 1266–1293, 2014.
- [5] M. Chiang, P. Hande, T. Lan, and C. W. Tan, *Power Control in Wireless Cellular Networks*, 2007, vol. 2, no. 4.
- [6] R. Jain, A. Dursesi, and G. Babic, "Throughput fairness index: An explanation," Tech. Rep., 1999.
- [7] D. J. Watts and S. H. Strogatz, "Collective dynamics of small world networks," *Nature*, vol. 393, no. 6684, pp. 440–442, jun 1998. [Online]. Available: <http://dx.doi.org/10.1038/30918>
- [8] A. Barrat, M. Barthélemy, R. Pastor-Satorras, and A. Vespignani, "The architecture of complex weighted networks," *Proc. Natl. Acad. Sci. U. S. A.*, vol. 101, no. 11, pp. 3747–3752, 2004.
- [9] R. Laufer and L. Kleinrock, "The Capacity of Wireless CSMA/CA Networks," *IEEE/ACM Trans. Netw.*, vol. 24, no. 3, pp. 1518–1532, jun 2016. [Online]. Available: <http://ieeexplore.ieee.org/document/7121036/>
- [10] T.-S. Kim, J. C. Hou, and H. Lim, "Improving spatial reuse through tuning transmit power, carrier sense threshold, and data rate in multihop wireless networks," *Proc. 12th Annu. Int. Conf. Mob. Comput. Netw. - MobiCom '06*, p. 366, 2006. [Online]. Available: <http://portal.acm.org/citation.cfm?doid=1161089.1161131>
- [11] J. D. J. Deng, B. L. B. Liang, and P. Varshney, "Tuning the carrier sensing range of IEEE 802.11 MAC," *IEEE Glob. Telecommun. Conf. 2004. GLOBECOM '04.*, vol. 5, pp. 2987–2991, 2004.
- [12] W. Jones, R. E. Willson, M. Sooriyabandara, and A. Doufexi, "Wireless Network MAC Layer Performance Evaluation with Full-Duplex Capable Nodes," in *Proc. 12th ACM Symp. QoS Secur. Wirel. Mob. Networks - Q2SWinet '16*. New York, New York, USA: ACM Press, 2016, pp. 111–118. [Online]. Available: <http://dl.acm.org/citation.cfm?doid=2988272.2990294>
- [13] G. Bianchi, "Performance analysis of the IEEE 802.11 distributed coordination function," *IEEE J. Sel. Areas Commun.*, vol. 18, no. 3, pp. 535–547, mar 2000. [Online]. Available: <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=840210>
- [14] J. Jeong, H. Kim, S. Lee, and J. Shin, "An analysis of hidden node problem in IEEE 802.11 multihop networks," in *Sixth Int. Conf. Networked Comput. Adv. Inf. Manag.*, 2010, pp. 282–285. [Online]. Available: http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=5573151
- [15] M. Garetto, T. Salonidis, and E. W. Knightly, "Modeling per-flow throughput and capturing starvation in CSMA multi-hop wireless networks," *IEEE/ACM Trans. Netw.*, vol. 16, no. 4, pp. 864–877, 2008.
- [16] M. S. Afaqui, E. Garcia-Villegas, E. Lopez-Aguilera, G. Smith, and D. Camps, "Evaluation of dynamic sensitivity control algorithm for IEEE 802.11ax," in *2015 IEEE Wirel. Commun. Netw. Conf.*, no. Wcnc. IEEE, mar 2015, pp. 1060–1065. [Online]. Available: <http://ieeexplore.ieee.org/document/7127616/>
- [17] V. P. Mhatre, K. Papagiannaki, and F. Baccelli, "Interference Mitigation Through Power Control in High Density 802.11 WLANs," in *IEEE INFOCOM 2007 - 26th IEEE Int. Conf. Comput. Commun.* IEEE, 2007, pp. 535–543. [Online]. Available: <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=4215651> <http://ieeexplore.ieee.org/document/4215651/>
- [18] OPNET, "OPNET Technologies - Network Simulator — Riverbed." [Online]. Available: <http://www.riverbed.com/gb/products/steelcentral/opnet.html?redirect=opnet>

William Jones received the BSc. Eng in Mechanical Engineering from Cardiff University, UK, in 2013; and is now a Research Engineer for Toshiba Research Europe Limited undertaking a Ph.D. degree from the University of Bristol, UK. His main areas of research interest are future wireless networks, 5G, network traffic modeling. He is an IEEE student member.

R Eddie Willson MA DPhil (Oxon) is Professor of Intelligent Transport Systems at the University of Bristol. Prior to this (2010 - 2012) he was Professor of Modelling and Simulation in the Transportation Research Group at the University of Southampton. He is a Fellow of the Institute of Mathematics and its Applications and his research interests span the application of Mathematical Modelling across a wide range of Engineering topics, but with a particular focus on the Transport sector. He is best known for his work on stop-and-go waves and the stability of highway traffic.

Angela Doufexi Received the BSc degree in Physics from the University of Athens, Greece, in 1996; the MSc degree in electronic engineering from Cardiff University, UK, in 1998; and the Ph.D. degree from the University of Bristol, UK, in 2002 where she is now a Professor in wireless networks. Her research interests include OFDM; multiuser diversity & resource allocation, WLANs, vehicular communications; multiple antenna systems; LTE and 5G systems; mmWave; SDN; and multimedia transmission. She has authored over 150 journal and conference papers in these areas.

Mahesh Sooriyabandara Received the BSc Eng from the University of Peradeniya, Sri Lanka, and the Ph.D. degree from the University of Aberdeen in the UK. In 2004, he joined Toshiba Research Europe Limited, where he is currently the Associate Managing Director. His research interests are wireless networks, Internet engineering, smart grid communications, and machine-to-machine communications. He is a Senior Member of IEEE and ACM and an honorary visiting professor at the Cardiff University.