Jones, W., Wilson, E., Sooriyabandara, M., & Doufexi, A. (2016). Wireless Network MAC Layer Performance Evaluation with Full-Duplex Capable Nodes. In *Proceedings of the 12th ACM Symposium on QoS and Security for Wireless and Mobile Networks (Q2SWinet '16)* (pp. 111-118). Association for Computing Machinery (ACM). <https://doi.org/10.1145/2988272.2990294>

Peer reviewed version

Link to published version (if available):
[10.1145/2988272.2990294](https://doi.org/10.1145/2988272.2990294)

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

This is the author accepted manuscript (AAM). The final published version (version of record) is available online via ACM at <http://dl.acm.org/citation.cfm?doid=2988272.2990294>. 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/>

Wireless network MAC layer performance evaluation with full-duplex capable nodes

William Jones
Toshiba Research Europe Limited
& University Of Bristol
Bristol, UK
+44(0)117 9060 782
William.Jones@Toshiba-TREL.com

Mahesh Sooriyabandara
Toshiba Research Europe Limited
Bristol, UK
+44(0)117 9060 700
Mahesh.Sooriyabandara@Toshiba-TREL.com

R Eddie Wilson
Dept of Engineering Mathematics
University of Bristol
Bristol, UK
+44 (0) 117 331 5690
RE.Wilson@Bristol.ac.uk

Angela Doufexi
Dept of Electrical and Electronic Engineering
University of Bristol
Bristol, UK
+44 (0) 117 954 5201
A.Doufexi@Bristol.ac.uk

ABSTRACT

This paper demonstrates the potential performance gains that the introduction of full-duplex nodes can provide to a wireless network. The paper focuses on two common issues in current communications networks; bottlenecks and hidden nodes. The technical approach used simulates a simplified MAC protocol and scans over the parameter space for which full MAC simulations would not be computationally tractable. In contrast to most literature which focusses on saturated traffic, this study identifies the capacity region, i.e., vectors of demand that can be met by the system. The study shows that introducing full-duplex access points alone mitigates against the problem of bottlenecks, reduces the impact of hidden nodes and can increase the capacity of a network. When full-duplex access points are able to work with full-duplex clients, the capacity gain is much more significant, however it is shown that much of this capacity gain occurs at uneven demand combinations. When the demand to all nodes is equally high, the introduction of full-duplex capability to clients is shown to increase the number of transmission attempts resulting in a significantly increased number of collisions and reduced network performance. Further we observe that at low traffic levels, a full-duplex access point may improve throughput by simply transmitting a busy tone to silence other transmissions whilst it receives, mitigating against the hidden node problem.

Keywords

Hidden Nodes; Bottlenecks; Full-Duplex; Wireless Networks; Capacity.

1. INTRODUCTION

Through developments in self-interference cancellation technologies, the potential to significantly increase link capacity in wireless networks has emerged by enabling nodes to transmit and receive simultaneously in full-duplex [2,3,6,9]. This development suggests the implementation of full-duplex nodes to be a favorable addition, dramatically increasing performance for next generation communications networks.

The purpose of this paper is twofold; firstly to evaluate the impact on performance of introducing full-duplex nodes into communications networks and secondly to demonstrate a modelling methodology and analysis process for doing so. We base our investigation around two common problems in existing communication networks; the formation of bottlenecks and the classic Hidden Node Problem [7].

When a Wireless Local Area Network (WLAN) network is saturated, a bottleneck occurs at the access point. This occurs because the access point is subject to the same protocol, and thus has the same probability of gaining access to the channel for transmission as each client, while serving more demands. Full-duplex can unblock this bottleneck, because of the ability of a full-duplex access point to transmit simultaneously whenever it is receiving a packet from a client [6]. We propose a method for modeling a simple example of this intervention and evaluate the results. Throughout this paper we refer to this example as the 'bottleneck problem'.

In WLANs the hidden node problem occurs when a half-duplex receiver cannot notify other nodes in the network that it is currently receiving a signal and thus other nodes attempt to transmit to the receiver. Implementing full-duplex has the capability to mitigate against the hidden node problem because of the additional ability of the receiving node to simultaneously transmit. This transmission suppresses nearby nodes thus eliminating the hidden node problem. We propose a methodology for modelling the hidden node problem and evaluate the impact on it of implementing full-duplex [6]. Throughout this paper we refer to this example as the 'hidden node problem'.

The approach presented in this paper differs from many existing works [12–14], who present models and analyze statistical data of large networks. These studies use probabilistic arguments to derive the throughput in saturated conditions for networks of different sizes comparing existing Half-Duplex Medium Access Control (HD-MAC) with proposed new Full-Duplex Medium Access Control (FD-MAC) protocols. Performance statistics, such as the likelihood of a successful packet transmission, are derived based on the concentration of nodes in the network. The

performance benefit of full-duplex is then expressed in terms of throughput per unit area.

This paper supports the existing work in this area [2–4,6,8,9,12–14,16] seeking to identify the performance gain that full-duplex nodes offer communications networks in the MAC layer. Our approach is to present a simple network setup to which we apply a simplified network interference model, like [12] with characteristics evolved from the classic Gupta and Kumar model [5], and apply a reduced complexity MAC protocol. This technical approach allows us to scan over the parameter space and identify performance characteristics across vectors of demand in a way not computationally tractable with normative simulation.

Several works recognize the potential of full-duplex nodes to reduce the effects of the ‘bottleneck problem’ or the ‘hidden node problem’, such as [10,15], and further others begin to analyze the impact [6,11]. The authors of [6] implement a physical test of two client devices hidden from each other attempting to communicate with an access point. They compare the system performance using a half-duplex and full-duplex access point and find significant improvement from the full-duplex access point. We further this analysis using a simulation approach, however limiting our study entirely to the MAC layer. Comparing simulations, using HD-MAC based on analysis in [1] and FD-MAC based on those proposed in [8,14,16], we are able to quantify the resulting performance increase for our simple network setup when full-duplex nodes are introduced, focusing our study around the ‘bottleneck problem’ and ‘hidden node problem’ introduced above.

2. Simple Network Setup

Figure 1. Simple network setup. Each client A and B is connected to the access point by two directional edges $e_1, e_2, \dots, e_4$.

We consider a three node network where two of those nodes are client devices that can communicate directly with the third node, a wireless access point, via a single hop, see fig 1. We consider no routing, each client (A,B) can transmit to the access point directly and hence there are four possible transmissions (edges in graph theory language) that we denote $e_1, e_2, \dots, e_4$. It thus follows we may prescribe demand (number of packets per unit time) on our network in terms of requested flows $d_1, d_2, \dots, d_4$ for each of the six edges respectively. However, to simplify the computational analysis, we restrict our study to the special symmetric case $d_1=d_{CA}, d_2=d_3=d_{AP}$, and $d_4=d_{CB}$, so that we need only scan through the three-dimensional demand space (d_{CA}, d_{AP}, d_{CB}) .

Thus it is assumed each node (including the access point) has a single queue of packets awaiting transmission with Poisson arrival rate: d_{CA}, d_{AP}, d_{CB} respectively. The queue on the access point will consist on average of an equal mix of packets bound for clients A and B. Further we assume there is no active queue management. Hence, regardless of the destination, packets queued at the access point are served in order of arrival.

A further description of our network modelling process is explained in section 3 and our assumptions concerning packet

collisions etc. for the bottleneck and hidden node problems are explained in section 4.

At various stages of our modelling process we consider two forms of full-duplex transmission: unidirectional, where an access point can receive from one client and simultaneously transmit to another, and bi-directional, where an access point can receive from and transmit to the same client simultaneously [14]. These are shown in fig 2, and how we model them is explained in section 3.

Figure 2. Simultaneous transmissions occurring in full-duplex (a) Unidirectional full-duplex (b) Bi-directional full-duplex

Physical layer management in relation to full-duplex is still an open research topic. For example, interference management techniques such as; beam-forming, sectorisation and directional diversity [3,4,13] are under investigation and their implementation in some form may be required to achieve the full-duplex scenarios we describe. We acknowledge that there are two challenges presented in the topologies of fig 2 and these are: firstly self-interference at the base station caused by simultaneous transmission and reception (fig 2 (a) and (b)); secondly interference from the uplink client at the downlink client (fig 2 (a)). In the first case the challenge is how to communicate in the presence of high power interference and in the second the challenge is how to communicate in the presence of a commensurate-power interference [3]. The literature on interference channels and managing interference is extensive, but definitive understanding of managing full-duplex interference is still under development. Acknowledging that both challenges are significant and need to be addressed, in the study that follows we assume a perfect physical layer capable of managing this interference and focus our study on the MAC layer only.

3. Operational Scenarios

We now consider collisions between the various edges within our network and how they might be mitigated by full duplex techniques. We will consider the bottleneck and hidden node problems and three operational scenarios in relation to each of them: half-duplex access point and half-duplex clients (HDAP-HDC), full-duplex access point and half-duplex clients (FDAP-HDC), full-duplex access point and full-duplex clients (FDAP-FDC). In fact, these are the only scenarios of significance. For example, it would not be of use to consider a half-duplex access point and full-duplex clients as the clients would be forced to communicate with the access point via a half-duplex protocol, making this scenario the same as HDAP-HDC.

Throughout we suppose a simple deterministic binary interference model, so that if two incompatible edges transmit simultaneously, there will be a collision with probability one. Thus we use a four-by-four collision matrix $\mathbf{E}$ (see fig 3) to describe each operational

scenario, with entries E_{ij} , $1 \leq i, j \leq 4$. 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 particular example, interference (and hence the collision matrix) is symmetric: that is, if edge i collides with edge j , then edge j also collides with edge i , i.e., $E_{ij}=E_{ji}$.

Scenario 1: Half-Duplex Access Point and Half-Duplex Clients (HDAP-HDC)

In this scenario at most one edge can be active at any specific instant (i.e., at most one transmission can take place at any one instant). Every edge collides with every other edge. Therefore the resulting collision matrix - see fig 3(a) - is full of 1's, in all off-diagonal positions.

Scenario 2: Full-Duplex Access Point and Half-Duplex Clients (FDAP-HDC)

We suppose that Scenario 1 is modified such that the access point is now full-duplex capable. This therefore allows unidirectional full-duplex [14] transmissions as shown in fig 2(a). We suppose that the access point is equipped with a transmitter and receiver that have self-interference cancelation capability. In such a set-up, we suppose that the access point may receive from a client (A say) and simultaneously transmit to the other client (B). Of course, the MAC protocol must be extended in some way to allow this possibility. With the introduction of a full-duplex access point to this scenario certain pairs of simultaneous transmissions are now allowed to occur collision-free. This is reflected in the collision matrix by setting $e_{13}=e_{24}=0$, and $e_{31}=e_{42}=0$, see fig. 3(b). (We suppose it is not possible for the access point to transmit to two clients simultaneously nor to receive from two clients simultaneously.)

Scenario 3: Full-Duplex Access Point and Full-Duplex Clients (FDAP-FDC)

We suppose that Scenario 2 is further modified such that the clients as well as the access point are now full-duplex capable. This therefore additionally allows bidirectional full-duplex [14] transmissions as shown in fig 2(b). We suppose that the clients and access point are equipped with a transmitter and receiver and have self-interference cancelation capability. In such a set-up, we suppose that the access point may receive from a client (A say) and simultaneously transmit to the same client (A). In the collision matrix for this scenario, see fig. 3(c), further pairs of simultaneous transmissions are allowed to occur collision-free. This is implemented in the collision matrix by setting $e_{12}=e_{34}=0$, and $e_{21}=e_{43}=0$.

$$\begin{matrix} \begin{bmatrix} 0 & 1 & 1 & 1 \\ 1 & 0 & 1 & 1 \\ 1 & 1 & 0 & 1 \\ 1 & 1 & 1 & 0 \end{bmatrix} & \begin{bmatrix} 0 & 1 & 0 & 1 \\ 1 & 0 & 1 & 0 \\ 0 & 1 & 0 & 1 \\ 1 & 0 & 1 & 0 \end{bmatrix} & \begin{bmatrix} 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 \\ 1 & 0 & 0 & 0 \end{bmatrix} \\ \text{(a)} & \text{(b)} & \text{(c)} \end{matrix}$$

Figure 3. Edge collision matrix for two nodes transmitting - E.
a) HDAP-HDC b) FDAP-HDC c) FDAP-FDC

4. Modelling the Bottleneck and Hidden Node Problems

To investigate the impact of introducing full-duplex nodes on the bottleneck and hidden node problems, we perform simulations based on the three Scenarios described in section 3.

The MAC model and simulation methodology applied are described in sections 5 and 6 respectively. For Scenario 1 where all nodes are half-duplex we apply a half-duplex MAC protocol and for the other two Scenarios where one or more of the nodes is full-duplex we apply a full-duplex MAC protocol.

4.1 Bottleneck Problem

Scenario 1 represents the bottleneck problem in its simplest form. The clients are both serving just one demand to the access point while the access point is serving demand to each of the two clients. Thus a bottleneck will occur at the access point. Scenarios 2 and 3 are modifications of Scenario 1 introducing full duplex nodes. For this part we assume for each Scenario that the nodes have a global knowledge of the network, i.e. no hidden nodes.

4.2 Hidden Node Problem

The premise of the hidden node problem stems from the issue that edges 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. This can be modeled by introducing a second matrix that we call the 'knowledge of the network' edge matrix $\mathbf{F}$ (see fig 4).

This is a four-by-four edge matrix to describe each edge's awareness of the network around it, with entries F_{ij} , $1 \leq i, j \leq 4$. 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 each other's transmissions). As with the collision matrices in our particular example, the network topology (and hence the matrix $\mathbf{F}$) is symmetric: that is, if edge i has knowledge of edge j , then edge j also has knowledge of edge i i.e., $F_{ij}=F_{ji}$. A lack of knowledge of each other's transmissions can result in the clients transmitting simultaneously, each not realizing the other is transmitting, causing a collision.

Considering the topology depicted in fig 1, the access point is able to hear either of the clients transmitting, similarly both of the clients are able to hear the access point transmitting. The hidden node problem occurs as a result of client A not being able to hear transmissions by client B and vice versa. This is captured in matrices $\mathbf{F}$, which are otherwise the same as the three collision matrices, by setting $f_{14}=0$ and $f_{41}=0$ (see fig 4).

$$\begin{matrix} \begin{bmatrix} 0 & 1 & 1 & 0 \\ 1 & 0 & 1 & 1 \\ 1 & 1 & 0 & 1 \\ 0 & 1 & 1 & 0 \end{bmatrix} & \begin{bmatrix} 0 & 1 & 0 & 0 \\ 1 & 0 & 1 & 0 \\ 0 & 1 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{bmatrix} & \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix} \\ \text{(a)} & \text{(b)} & \text{(c)} \end{matrix}$$

Figure 4. Knowledge of the network edge matrix - F.
a) HDAP-HDC b) FDAP-HDC c) FDAP-FDC

In summary: we model the hidden node problem by implementing the three described Scenarios. We assume for each Scenario that the nodes have knowledge of the network via $\mathbf{F}$ but are subject to collisions via $\mathbf{E}$. For the three simulations $\mathbf{F}$ is modeled using the 'knowledge of the network' edge matrix in figures 4 (a), (b), and (c) respectively, however the collision matrices $\mathbf{E}$ are as in figures 3 (a), (b), and (c) respectively.

5. A Simplified MAC model

Fig 5. State diagram for each node of the simplified MAC model. *Dormant*: node senses channel. *Countdown*: random time delay between sensing channel is free and beginning transmission. *Transmit*: packet is transmitted. *Backoff*: random time delay following collisions.

Our methodology simulates network traffic for a large number of demand scenarios. This is computationally demanding so to reduce the CPU time needed, like [12], our simulation implements a simplification of the MAC protocol. The most significant simplification is in the pre- and post- stages of packet transmission (requests to send, clear to send, acknowledgments). These are not explicitly modeled, instead their influence is captured by the introduction of a *Countdown* state to the protocol - see fig 5. This is heuristically designed to incorporate the relevant complexity mitigated by simplifying other parameters.

A discrete time step is used as opposed to a discrete event methodology for simplicity. Each transmitter manages a queue with discrete Poisson arrivals and cycles through the model protocol states, pending an integer number of time steps in each, as shown in fig. 5.

The cycle begins with the *Dormant* state. If a client or the access point has a packet queued to send it senses the channel. If clear (i.e., if it were to transmit along the edge it wishes to transmit along, to the best of the nodes knowledge this would not cause a collision) it proceeds to the *Countdown* state, spending a randomly selected period there. After this delay, the node continues the cycle and moves to *Transmit* for the duration of the packet (again selected from a distribution).

If the transmission completes collision-free; the queue is decremented, the node returns to *Dormant* and a successful transmission recorded. If at some point during the transmission another node attempts to transmit along an incompatible edge a collision is recorded. Collisions during the transmission are only recognized at the end of the transmission state. In this instance the collided nodes enters a *Backoff* state for a random duration before reassessing the channel in an attempt to resolve the conflict.

When the *Backoff* stage completes, the node listens to the channel, and if clear, enters *Countdown* and proceeds through the cycle again as above; if the channel is busy, then the *Backoff* state repeats with a new random duration.

In the half-duplex protocol used for both problems in Scenario 1, the uniform distributions for the duration of *Backoff* and *Transmit* were developed from the analysis in [1], however additional time steps were included in *Transmit* to encompass the pre- and post-stages of the protocol (request to send, clear to send, acknowledgement) not explicitly modeled. For the full-duplex protocols used in both problems Scenarios 2 and these parameters

were developed from the full-duplex MAC protocols proposed in [7,13,14].

Key to this modelling methodology is the *Countdown* state. This is introduced to artificially represent uncertainty in the MAC protocol arising from the imperfections of the request and clear to transmit stages that have not been explicitly modeled. Without the *Countdown* state, collisions would only arise in our simulations if two nodes were to begin transmission absolutely simultaneously at the same time step. However, if in the *Countdown* state, the node is committed to transmit, a subsequent node may also sense that the channel is clear during this period and likewise enter *Countdown*, consequently resulting in a collision. The *Countdown* length is randomly selected from a uniform distribution and its parameters have been developed so that our simplified MAC model has collision rates representative of the real MAC protocol in the simplest possible network setup of two competing nodes. Broadly speaking, *Countdown* and *Backoff* states are much shorter in duration than a packet *Transmission*.

6. Simulation Methodology

Our investigation simulates each of the Scenarios 1-3 (see section 3) for both the bottleneck problem and the hidden node problem (see section 4) using the simplified MAC model (see section 5). By this, we mean the collision matrix and knowledge of the network edge matrix used to model the network were varied for the respective problem and Scenario being investigated (section 3/4) and the protocol was adjusted between the half-duplex and full-duplex appropriately (as discussed in section 5).

For each Scenario, we perform a large ensemble of simulations, each with a different demand vector (d_{CA}, d_{AP}, d_{CB}) (see section 2). Specifically, we let d_{sat} be the maximum transmission rate of a single half-duplex client or access point transmitting alone in a clear channel. Each of the demands d_{CA} , d_{AP} and d_{CB} is varied independently from 0 to d_{sat} in 40 equal increments, resulting in $41^3 = 68,921$ simulations for each of the two problem's three Scenarios. The duration of each simulation is set at $10,000/d_{sat}$ (i.e., the time needed to transmit 10,000 half-duplex packets without collisions). For more robust numerical results, lengthier simulation time is desirable, however this would be enormously demanding in terms of the CPU requirement.

For each individual simulation, we gather statistics on the total number of packets successfully transmitted, the total number of failed transmissions, the time evolution of queues, and the average latency per packet. The results that follow (section 7) are based on comparisons between Scenarios 1-3 for the bottleneck problem and the hidden node problem.

7. Results

For each of the two problems investigated, the three scenarios' simulation results can be presented in the form of a three-dimensional scatter plot, see fig 6(a), where red markers indicate combinations of demand that are within the capacity region. In this plot $\delta_{CA} := d_{CA}/d_{sat}$, $\delta_{AP} := d_{AP}/d_{sat}$ and $\delta_{CB} := d_{CB}/d_{sat}$ 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 fig 6(a) in order to compare the scenarios.

For the purpose of analysis, firstly we may consider the line $\delta_{CA}=\delta_{AP}=\delta_{CB}$ along which the demands are equal. Secondly, we define aggregate demand intensity $\delta:=\delta_{CA}+\delta_{AP}+\delta_{CB}$. Then $\delta=\text{const.}$ defines triangular cross sections through fig 6(a).

We examine the ratio of successful transmissions to failed transmissions (on the line $\delta_{CA}=\delta_{AP}=\delta_{CB}$) and compare across the two problems and three scenarios. See fig 6(b). From this figure the demand (δ) at which each client and the access point reaches saturation is apparent.

We also examine latency across the triangular sections δ_1 and δ_2 , and we compare across the three scenarios for each of the two problems. See fig 6(c).

Fig 7(a) helps give a clearer picture of the structure of the capacity region. 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 four scenarios: see fig 7(b).

Further, we may measure the proportion of the simulations that are within the capacity region - that is, the proportion of the volume $[0,1]\times[0,1]\times[0,1]$ that is within the capacity region: see fig 7(c).

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

Figure 6. (a) Capacity region plot for Bottleneck Problem HDAP-HDC. Red markers indicate demand combinations within the capacity region. (b) A plot of demand (δ) against ($\epsilon\delta$), where ϵ is the failure rate i.e. the number of failed transmissions per successful transmission, along the line $\delta_{CA}=\delta_{AP}=\delta_{CB}$. (c) Triangular sections $\delta:=\delta_{CA}+\delta_{AP}+\delta_{CB}=\text{const.}$ for constant $\delta_1=0.65$ (top row) and $\delta_2=0.80$ (bottom row). Plots show latency across the corresponding section for the two problems and three Scenarios investigated. The top of each triangle represents the access point and the left and right corners clients A and B respectively.

Figure 7. (a) Triangular sections $\delta=\delta_{CA}+\delta_{AP}+\delta_{CB}=\text{const}$ (see fig 6) for constant $\delta_1=0.65$ and $\delta_2=0.80$ for Bottleneck HDAP-HDC simulation. S denotes the proportion of the triangle covered with red markers. (b) Proportions S of triangular section within the capacity region, as a function of total demand intensity $\delta=\delta_{CA}+\delta_{AP}+\delta_{CB}$, compared across three scenarios for each of the two problems. (c) Table of volume V of the capacity region.

8. Analysis

The findings of our study are summarized in figs 6 and 7.

Fig 6(c) demonstrates the impact of full-duplex for demand intensity $\delta=0.65$ and $\delta=0.80$. The figure indicates how the latency around nodes with full-duplex capability is significantly reduced from half-duplex nodes. The hidden node plots show as expected lower latency around the access point (i.e., for combinations of demands in which the access point has the greatest share), which is able to hear all nodes in the network and therefore is not subject to the problem of hidden nodes as the clients are. Comparing the hidden node plots for $\delta=\delta_1$ Scenario 2 (FDAP-HDC) seems to show a larger area of lower latency around the access point (the top corner of the triangle) in comparison to Scenario 3 (FDAP-FDC). However Scenario 3 shows the area of low latency extending with clearer definition from the access point to the clients down the edges of the triangle. The same is true, but less visibly clear, for $\delta=\delta_2$. This suggests a lower throughput from the access point in Scenario 3 than Scenario 2 but improved throughput from the clients. Comparing the bottleneck problem simulations, the Scenario 1 (HDAP-HDC) plot shows a symmetrical plot with latency equal on all corners of the triangle. This reflects the nodes' awareness of others in the network and equal probability of accessing the channel. Introducing full duplex shows a significant improvement in performance, with the most significant improvement occurring around the access point.

Consider fig 6(b) plotting demand against error rate along the line $\delta_{CA}=\delta_{AP}=\delta_{CB}$. For the bottleneck problem simulations the access points and the clients follow a broadly similar pattern with the error rate increasing as the traffic is increased. Scenarios 2 (FDAP-HDC) and 3 (FDAP-FDC) show an increase in the value of δ at which the network reaches saturation compared to Scenario 1 (HDAP-HDC). In Scenario 1 both the access point and clients follow the same trend and reach saturation at the same value of δ . For Scenario 2 the throughput at which the access point reaches saturation is higher than the clients with roughly similar error rates. For Scenario 3 the client and the access point reach saturation at approximately the same value of δ , however the error rate from clients' transmissions is higher. The access point is

saturated in Scenario 2 at higher δ than Scenario 3 whereas the clients for both Scenarios saturate at similar values of δ .

Let us continue to analyze fig 6(b). For the hidden node problem at low δ , the error rate $\mathcal{E}$ of clients transmissions in Scenarios 2 (FDAP-HDC) and 3 (FDAP-FDC) is shown to be high, significantly higher than those in Scenario 1 (HDAP-HDC). As δ increases for Scenarios 2 and 3, $\mathcal{E}$ for clients decreases and $\mathcal{E}$ for the access point increases. Further, in Scenarios where the traffic at the access point is saturated, or close to saturated, the throughput is shown to be significantly increased and the ratio of successful to failed transmissions is improved using a full-duplex access point relative to a half-duplex access point. It is shown that in Scenario 3 there is a higher error rate (failed packets per successful transmission) than in Scenario 2. Further the saturation of Scenario 2 occurs at a higher value of δ for the clients and a significantly higher value of δ for the access point (more than three times) than Scenario 3.

Consider fig 7(b): a higher value of S for a given δ is better in that it means a higher proportion of demand scenarios for a given total demand δ are within the capacity region. Therefore a right shift of the curves indicates an improvement in capacity. Here for both the bottleneck and hidden node problems we see an increase from Scenario 1 (HDAP-HDC) to Scenario 2 (FDAP-HDC) and a further increase to Scenario 3 (FDAP-FDC). This is summarized in fig 7(c), which shows the capacity regions' volumes V . This result differs from our findings above considering the line $\delta_{CA}=\delta_{AP}=\delta_{CB}$ where Scenario 2 performed better than Scenario 3. This indicates that much of the growth in V is in uneven demand combinations – for example where (say) δ_{AP} and δ_{CA} are high, and δ_{CB} is low.

8. Discussion

Consider the hidden node problem: for all scenarios simulated the knowledge of the network edge matrices differs from the collision matrix which indicates a risk that certain incompatible edges may attempt to transmit simultaneously (i.e. client A and B attempt to transmit simultaneously to the access point) regardless of the applied protocol. Due to the ability of a full-duplex access point

to transmit whilst receiving, it can dictate to the clients which other edges can transmit simultaneously in full-duplex. Further, the busy signal from the access point silences other clients that may potentially cause a collision with the transmitting client. As a result, the problem of hidden nodes is, to an extent, mitigated by full-duplex on the access point.

It is clear from fig 7 that regardless of the protocol or the number of full-duplex capable nodes, the presence of hidden nodes significantly decreases the network capacity. For both the bottleneck and hidden node simulations the introduction of full-duplex nodes shows an increase in the capacity region thus contributing (although in some Scenarios only slightly) to addressing the two problems. As the findings of [13] show, we anticipate that in a network with a greater number of clients, the gain achieved from full-duplex would be less due to additional aggregate interference.

Fig 6(b) shows that in hidden node Scenarios, as traffic is increased, the error rate of the clients decreases (up to the point of saturation). It is apparent from the results that in a situation where the demand at the access point is low, the throughput of the network could be improved by the access point simply transmitting noise (i.e., a busy tone) to silence other clients.

Both the bottleneck and hidden node problem simulations have shown Scenario 2 (FDAP-HDC) to reach saturation at a higher value of δ than Scenario 3 (FDAP-FDC), and have a lower error rate at saturation when considering equal demand to all nodes. Further for both problems the clients of Scenario 2 and 3 have reached saturation at similar values of δ , however, in both cases the error rate for Scenario 2 has been lower. This result is analogous with observations made by, among others, Bianchi [1] in his analysis of 802.11 networks. It is commonly known that several random access schemes exhibit an unstable behavior. As the load on the network increases, the throughput rises up to a maximum value. Further increasing the load on the network can lead eventually to a decrease in system performance. In comparison to Scenario 2, for an equivalent value of δ , the number of attempted transmissions from the clients in Scenario 3 (all clients full-duplex) is significantly higher. However, this increase in the number of transmission attempts is found to have a negative effect on the overall network performance, resulting in an increased number of collisions, reduced performance of the access point and lower throughput.

The results show that at saturated (or sufficiently high) traffic the introduction of full-duplex nodes can significantly reduce the effect of hidden nodes. The improvement in performance from all nodes will allow for increased throughput in networks with full-duplex access points. Further the greater improved performance of the access points over the clients with the introduction of full-duplex will enable networks to better cope with bottlenecks formerly caused by one access point serving multiple clients.

Of course, the results presented here could be refined by increasing the number of demand combinations (i.e., use more than 40 increments from zero to d_{sat}) and the sampling error (apparent as irregularity / noise in figs 6(c) and 7(b)) could be reduced by increasing the run time for each individual simulation. However both such refinements would significantly increase the total CPU time required for this study. In any case the saturation throughputs shown in fig 6(b) and the aggregate capacity region measure V are relatively stable to misclassification of individual

simulations, and so the present study seems safe in its overall findings.

Further, normative testing of our conclusions should be carried out using more detailed simulation models (e.g., with NS3 [17]) of the MAC protocol and the physical layer. However the computational cost of such normative simulations is very high and will result in impractical run times if a full scan of demand space (in the manner of the present study) were to be attempted. However, in parallel to this publication we are producing further work to demonstrate by probabilistic analysis that our MAC models is a suitable substitute for normative testing.

9. Conclusion

This paper has investigated the performance gain available from the introduction of full-duplex nodes to wireless networks. The study has considered two recognised problems in communications networks: bottlenecks and hidden nodes. We have assumed a highly simplified approach to modelling the physical layer (deterministic binary collision model) and rather focussed our efforts in understanding performance at the MAC level.

We have modelled the two problems for three different operational scenarios: one with two half-duplex clients and a half-duplex access point, one with two half-duplex clients and a full-duplex access point and one with two full-duplex clients and a full-duplex access point.

Our simulation results have shown that the introduction of full-duplex access points alone can improve the network capacity, reduce bottlenecks and the undesirable effect of hidden nodes. The addition of full-duplex clients further improves the capacity of the overall network, however much of this gain is at uneven demand combinations. For even demand combinations, the addition of full-duplex clients increased the number of attempted transmissions however, also increased the error rate such that the overall network performance decreased. Statistics we provide in relation to the capacity region enable improved understanding of potential network performance gains compared to simple measurement of saturated flows. To mitigate against hidden nodes at low traffic levels, a full-duplex access point may improve throughput by simply transmitting a busy tone to silence other transmissions whilst it receives.

Further work should consider how our conclusions extend to larger and more realistic network examples with more nodes. A further complication is that an exhaustive search of the (very high dimensional) demand space will not be computationally tractable, and more refined search techniques will need to be adopted.

10. ACKNOWLEDGMENTS

This research is supported by EPSRC, Toshiba Research Europe Limited and the University of Bristol Systems Centre.

11. REFERENCES

- [1] Giuseppe Bianchi. 2000. Performance analysis of the IEEE 802.11 distributed coordination function. *IEEE Journal on Selected Areas in Communications* 18, 3: 535–547. <http://doi.org/10.1109/49.840210>
- [2] Jung Il Choi, Mayank Jain, Kannan Srinivasan, Phil Levis, and Sachin Katti. 2010. Achieving single channel, full duplex wireless communication. *Proceedings of the sixteenth*

- annual international conference on Mobile computing and networking - MobiCom '10, ACM Press, 1.
<http://doi.org/10.1145/1859995.1859997>
- [3] Evan Everett, Melissa Duarte, Chris Dick, and Ashutosh Sabharwal. 2011. Empowering full-duplex wireless communication by exploiting directional diversity. 2011 Conference Record of the Forty Fifth Asilomar Conference on Signals, Systems and Computers (ASILOMAR), IEEE, 2002–2006. <http://doi.org/10.1109/ACSSC.2011.6190376>
 - [4] Sanjay Goyal, Pei Liu, Shivendra S Panwar, Robert A Difazio, Rui Yang, and Erdem Bala. 2015. Full duplex cellular systems: will doubling interference prevent doubling capacity? IEEE Communications Magazine 53, 5: 121–127. <http://doi.org/10.1109/MCOM.2015.7105650>
 - [5] Piyush Gupta and P.R. Kumar. 2000. The capacity of wireless networks. IEEE Transactions on Information Theory 46, 2: 388–404. <http://doi.org/10.1109/18.825799>
 - [6] Mayank Jain, Jung Il Choi, Taemin Kim, et al. 2011. Practical, real-time, full duplex wireless. Proceedings of the 17th annual international conference on Mobile computing and networking - MobiCom '11, ACM Press, 301. <http://doi.org/10.1145/2030613.2030647>
 - [7] Jangkeun Jeong, Hyuntai Kim, Sangtae Lee, and Jitae Shin. 2010. An analysis of hidden node problem in IEEE 802.11 multihop networks. Sixth International Conference on Networked Computing and Advanced Information Management (NCM), 282–285. Retrieved from http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=5573151
 - [8] Ruizhi Liao, Boris Bellalta, and Miquel Oliver. 2015. Modelling and Enhancing Full-Duplex MAC for Single-Hop 802.11 Wireless Networks. IEEE Wireless Communications Letters 4, 4: 349–352. <http://doi.org/10.1109/LWC.2015.2417166>
 - [9] Bozidar Radunovic, Dinan Gunawardena, Peter Key, et al. 2010. Rethinking Indoor Wireless Mesh Design: Low Power, Low Frequency, Full-Duplex. 2010 Fifth IEEE Workshop on Wireless Mesh Networks, IEEE, 1–6. <http://doi.org/10.1109/WIMESH.2010.5507905>
 - [10] Ashutosh Sabharwal, Philip Schniter, Dongning Guo, Daniel W Bliss, Sampath Rangarajan, and Risto Wichman. 2014. In-Band Full-Duplex Wireless: Challenges and Opportunities. IEEE Journal on Selected Areas in Communications 32, 9: 1637–1652. <http://doi.org/10.1109/JSAC.2014.2330193>
 - [11] Achaleshwar Sahai, Gaurav Patel, and Ashutosh Sabharwal. 2011. Pushing the limits of Full-duplex: Design and Real-time Implementation. Retrieved from <http://arxiv.org/abs/1107.0607>
 - [12] Shu Wang, Vignesh Venkateswaran, and Xinyu Zhang. 2015. Exploring full-duplex gains in multi-cell wireless networks: A spatial stochastic framework. 2015 IEEE Conference on Computer Communications (INFOCOM), IEEE, 855–863. <http://doi.org/10.1109/INFOCOM.2015.7218456>
 - [13] Xudong Wang, Huaiyu Huang, and Taewon Hwang. 2015. On the Capacity Gain from Full Duplex Communications in A Large Scale Wireless Network. IEEE Transactions on Mobile Computing 1233, c: 1–1. <http://doi.org/10.1109/TMC.2015.2492544>
 - [14] Wenchi Cheng, Xi Zhang, and Hailin Zhang. 2013. RTS/FCTS mechanism based full-duplex MAC protocol for wireless networks. 2013 IEEE Global Communications Conference (GLOBECOM), IEEE, 5017–5022. <http://doi.org/10.1109/GLOCOMW.2013.6855746>
 - [15] Xiufeng Xie and Xinyu Zhang. 2014. Does full-duplex double the capacity of wireless networks? Proceedings - IEEE INFOCOM, 978: 253–261. <http://doi.org/10.1109/INFOCOM.2014.6847946>
 - [16] Xi Zhang, Wenchi Cheng, and Hailin Zhang. 2015. Full-duplex transmission in phy and mac layers for 5G mobile wireless networks. IEEE Wireless Communications 22, 5: 112–121. <http://doi.org/10.1109/MWC.2015.7306545>
 - [17] NS-3. Retrieved from <https://www.nsnam.org/>