When Is the Right Time to Transmit in Multi-hop White-Fi?

N. Facchi*, F. Gringoli*, D. Malone†, and P. Patras‡

*CNIT / University of Brescia, Italy
†Hamilton Institute, Maynooth University, Ireland
‡School of Informatics, The University of Edinburgh, UK

Abstract—**While Western societies are becoming increasingly connected, many developing regions lack basic Internet connectivity, primarily due to the high costs associated with infrastructure deployment and maintenance. Potential exists for the TV white-space (TVWS) wireless technology to bridge this digital divide, though efficient channel access mechanisms suited to multi-hop networks that operate in sub-gigahertz bands are yet to be developed. Using a small test bed, we demonstrate a prototype implementation of a medium access protocol that learns appropriate transmission opportunities in such settings, achieving pseudo-scheduled behaviour *ex tempore* and providing substantial gains over the de facto IEEE 802.11af protocol.**

I. INTRODUCTION

Regulators worldwide recently mandated license exempt access to ultra-high frequencies (UHF), as part of a broader effort to enable global access to the Internet for all (GAIA) through wireless technologies, and curb illiteracy in underdeveloped countries. This spectrum is particularly attractive for affordable data communications, as signals propagate through obstacles more easily and coverage increases substantially in the 500–800MHz frequency range [1]. In areas where wired communications infrastructure is sporadic, building multi-hop White-Fi networks to provide end users with digital connectivity only seems natural. On the other hand, due to long communications ranges specific to these bands, the de facto IEEE 802.11af protocol [2] suffers severe throughput degradation as stations independently employ carrier-sensing to schedule their transmissions [3].

To overcome this problem, we have implemented a decentralised medium access protocol that does not rely on carrier sensing, but instead uses a simple yet effective learning technique to identify transmit opportunities that do not collide with those of neighbouring stations.

We note that learning medium access control (MAC) protocols have been explored in the past for deterministic access to single-hop networks (see e.g. [4], [5]). However, these approaches do not work in multi-hop settings, since all stations must overhear each other, the wireless channel is slotted, and slots need to be aligned. In contrast, our solution works on commodity hardware without requiring information about the neighbourhood size, whether the channel is idle/busy, or precise synchronisation among stations. Our

technique converges quickly to pseudo-scheduled operation and attains substantial throughput and fairness improvements in multi-hop settings.

We demonstrate the performance of our prototype through a simple yet pathological set-up. Precisely, we use four wireless stations to simultaneously create both hidden and exposed terminal circumstances, while establishing traffic flows between pairs. For experimentation we rely on Linux-based embedded PCs with off-the-shelf Broadcom wireless cards that, although operating in the 2.4GHz band (due to limited availability of commercial UHF devices), can replicate the behaviour at lower frequencies. When all stations transmit backlogged UDP traffic, we show that our learning-driven MAC attains up 50% more throughput than the standard 802.11af, improves fairness, and eliminates frame loss almost entirely.

In what follows we first give a brief overview of IEEE 802.11af operation, then describe our prototype implementation, outline the proposed demonstration, and end with concluding remarks.

II. IEEE 802.11AF OVERVIEW

The 802.11af amendment [2] addresses the unlicensed access of popular Wi-Fi devices to unused TV bands (e.g. 470–790MHz in Europe), commonly referred to as white-spaces. This technology is also known as White-Fi and operates on 6/7/8 MHz wide channels, which ensures backward compatibility with existing TV broadcasters. For high data rates, up to four contiguous and non-contiguous channels can be aggregated. To avoid interfering with incumbents, white-space devices (WSD) access a geolocation database to identify vacant channels at a specific location and time.

The WSDs' access to the channel is regulated by a distributed coordination function (DCF) which implements a Carrier Sense Multiple Access with Collision Avoidance (CSMA/CA) protocol. Through this, as in standard 802.11, stations observe a slotted channel and listen to the medium before attempting to transmit. Specifically, if the medium is idle for a DIFS time (DCF inter-frame space), a contender starts a back-off procedure whereby it selects a random counter, uniformly distributed in the zero – contention window range, decrements this every slot and transmits when the

counter reaches zero. A successful transmission is acknowledged by the intended receiver, whereas upon collision the contention window is doubled and the procedure is repeated.

In multi-hop settings where some stations are not within carrier sensing range of each other and their slots are no longer aligned, this mechanism suffers severe throughput degradation, which our prototype tackles as we discuss next. We note that *802.11 and 802.11af MAC operation is similar, which allows us to use commodity Wi-Fi hardware for demonstration purposes, without compromising our goals.*

III. A LEARNING DRIVEN MAC PARADIGM

Undoubtedly, carrier sensing is ineffective in multi-hop topologies and especially problematic when stations operate on TV white spaces (TVWS) [6]. This mechanism is however useful to certain extent, to avoid (some) collisions when triggering transmissions. In the absence of carrier sensing, as well as of a central entity that has complete knowledge of the network topology and can coordinate the access of all stations, the question we seek to answer in this demo is: "*When is the right time to transmit in multi-hop White-Fi?*".

Protocol Operation: Instead of employing carrier sensing, we propose regarding the channel as divided into small slots (e.g. equal to the duration of an idle slot as defined by the 802.11 standard). Contenders observe cycles comprising a sufficient number of such slots, in order to accommodate their transmission (i.e data frame and acknowledgement) and those of their inferred neighbours. However, they need not explicitly agree on the start of their individual cycles or slots.

We note that neighbouring stations that observe different schedule lengths may periodically collide if these are not multiples of each other. To overcome this issue, a station operates with a schedule length that contains a number of slots equal to the power of two value nearest to its number of one- and two- hop neighbours, n, multiplied by the average packet duration, T. Formally,

$$\mathcal{S} = 2^{\lceil \log_2 n \rceil} \cdot \frac{T}{t},$$

where t is the duration of a mini slot. We illustrate this with a simple chain topology example in Fig. 1.

Stations running our approach use a simple learning technique to individually identify a sets of slots in which to
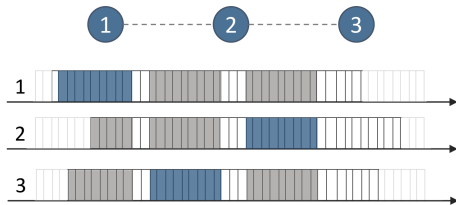

Figure 1: Simple multi-hop network example with three stations. Stations 1 and 3 are within transmission range of node 2, and two-hop neighbours to each other. Channel divided in mini slots of size t. The schedule observed by all contenders is $\mathcal{S} = 4T/t$. Individual schedules are not aligned. The set of slots chosen by each station for transmission shown with dark shade.

transmit to direct neighbours without collision. Specifically, the learning driven access protocol updates the probabilities of attempting transmission at different slots and works according to the following logic:

- Initialise all slot probabilities with $1/\mathcal{S}$ and choose one slot randomly;
- Upon failure, decrease the probabilities of the same and adjacent slots, and increase those of farther ones, i.e.

$$p_k = (1 - \eta)p_k + \eta \frac{2^{\|k-j\|}}{3(2^{\mathcal{S}/2} - 1)};$$

then choose a new slot probabilistically;
- Upon success, continue using the same slots in the following cycle.

In the above, $\eta \in (0, 1)$ is the learning rate and $\|k - j\|$ is the modulo $\mathcal{S}$ distance between slots j and k. Through this procedure the multi-hop network converges to pseudo-scheduled operation.

Implementation: We prototype our learning driven MAC for multi-hop White-Fi using Broadcom BCM4318 cards. Though these do not work in TVWS, they are appropriate for a proof-of-concept demonstration of the practicality of our approach. Importantly, they are inexpensive and support a well-documented highly-programmable open-source firmware [7], [8]. Particularly important to our implementation is the ability to interrupt ongoing reception and have fine control over transmission start, which is feasible with these devices and firmware.

Since the CPU capabilities are limited to elementary operations such as sum and subtraction, we only work with power of two values and use the following workaround to be able to update slot probabilities:

$$2^{\|k-j\|}/3 \approx 2^{\|k-j\|-16} \cdot \text{0x5555},$$

which no longer involves the division by 3 operation.

Finally, to trace the chosen slots and frame transmission times, we introduce small modifications that enable embedding this information into the Logical Link Control (LLC) protocol headers.

IV. DEMONSTRATION SET-UP

To demonstrate the operation of our prototype and the performance improvements achievable by the learning driven protocol over the standard 802.11af DCF, we will deploy a small test bed as shown in Fig. 2. This consists of four Linux-based PC Engines Alix 2D2 boards equipped with Broadcom BCM4318 cards and a laptop computer we use for experiment control and visualisation purposes. To confine the demonstration to a small physical space while creating pathological hidden and exposed terminal settings, we interconnect the wireless adapters of the stations using RF cables and attenuators. This also shields our experiments from other demonstrations we expect to take place at the conference venue.
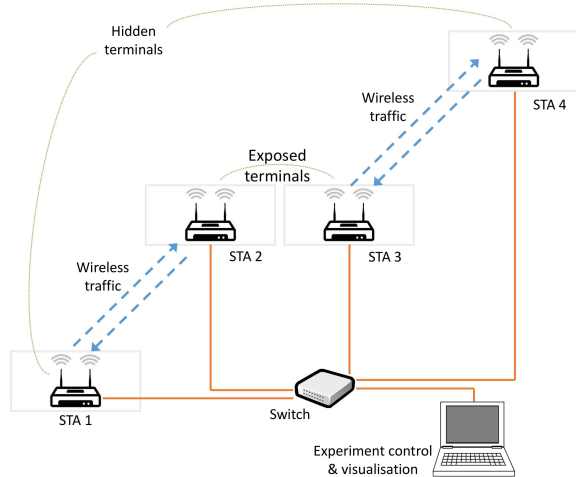
Figure 2: Demonstration set-up. Stations 1 and 3 send UDP traffic to 2 and respectively 4, and vice-versa. Nodes can employ the 802.11 DCF protocol and our learning driven MAC. A laptop PC is used for controlling the experiments and visualising the performance of the two approaches.
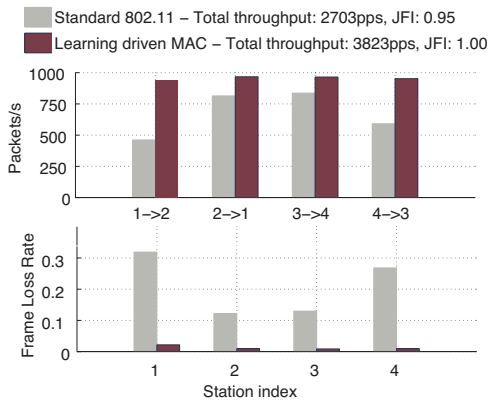

Figure 3: Example performance analysis with a network set-up similar to that shown in Fig. 2. Throughput (above) and frame loss (below) comparison with standard DCF.

All nodes will simultaneously transmits backlogged UDP flows to their direct neighbours using the `iperf` tool, namely station 1 to 2, station 3 to 4, and vice-versa. For comparison, *we will demonstrate the performance of our solution and that of 802.11af DCF, in terms of total throughput, throughput fairness, and packet loss.*

Convergence Analysis: To verify that stations running our learning driven channel access protocol reach a steady state where they transmit deterministically without overlapping, we capture all transmitted frames using `tcpdump` and examine their LLC headers, where recall that the firmware inserts the slot number and time chosen for the start of transmission. We use simple scripts to combine all traces captured and plot the time evolution of the slot selected by each station. Since we do not rely on any synchronisation between nodes, we do expect to observe certain clock drift, but will demonstrate the system converges to pseudo-scheduled operation quickly.

Performance Gains: We show that the standard DCF scheme faces hidden and exposed terminal circumstances, as it relies on carrier sensing to schedule transmissions. This results in poor overall throughput performance, unfair distribution of the individual throughputs, and substantial frame loss rates. In contrast, when stations run our learning driven approach, Jain's Fairness Index (JFI) is ~ 1, the total throughput is $\sim 50\%$ higher, and frame loss is almost completely mitigated, as exemplified in Fig. 3.

We believe these experiments will demonstrate the practical feasibility of our solution and its potential to achieve efficient scheduling in multi-hop White-Fi networks.

V. CONCLUDING REMARKS

Being able to deploy an efficient decentralised medium access paradigm onto inexpensive multi-hop 802.11af devices is arguably key to providing connectivity to underdeveloped regions. In this demo we showcase a proof-of-concept implementation of a learning driven scheme on off-the-shelf wireless hardware, that enables stations to learn independently appropriate transmit opportunities. Our solution does not require explicit synchronisation, converges quickly to pseudo-scheduled operation, and substantially outperforms the MAC protocol defined by the 802.11 standard.

REFERENCES

[1] A. Giannoulis, P. Patras, and E. Knightly, "Mobile access of wide-spectrum networks: Design, deployment and experimental evaluation," in *Proc. IEEE INFOCOM*, Turin, Italy, Apr. 2013.

[2] IEEE Std 802.11af, *Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications. Amendment 5: Television White Spaces (TVWS) Operation*, Dec. 2013.

[3] B. Radunović, R. Chandra, and D. Gunawardena, "Weeble: enabling low-power nodes to coexist with high-power nodes in white space networks," in *Proc. ACM CoNEXT*, Nice, France, Dec. 2012.

[4] M. Fang, D. Malone, K. R. Duffy, and D. J. Leith, "Decentralised learning MACs for collision-free access in WLANs," *Wireless Networks*, vol. 19, pp. 83–98, Jan. 2013.

[5] J. Barcelo, A. Toledo, C. Cano, and M. Oliver, "Fairness and Convergence of CSMA with Enhanced Collision Avoidance (ECA)," in *Proc. IEEE ICC*, May 2010.

[6] P. Ren, Y. Wang, Q. Du, and J. Xu, "A survey on dynamic spectrum access protocols for distributed cognitive wireless networks," *EURASIP Journal on Wireless Comms. & Networking*, vol. 60, Feb. 2012.

[7] P. Salvador, V. Mancuso, P. Serrano, F. Gringoli, and A. Banchs, "VoIPiggy: Analysis and implementation of a mechanism to boost capacity in IEEE 802.11 WLANs carrying VoIP traffic," *IEEE Trans. on Mobile Computing*, vol. 13, no. 7, pp. 1640–1652, Jul. 2014.

[8] P. Patras, H. Feghhi, D. Malone, and D. Leith, "Policing 802.11 MAC Misbehaviours," *IEEE Trans. Mobile Computing*, Sept. 2015.