

CSMA/CA in Time and Frequency Domains

Julien Herzen*, Albert Banchs†, Vsevolod Shneer‡, Patrick Thiran*

*EPFL, Switzerland † Heriot-Watt University, United Kingdom

‡Institute IMDEA Networks and University Carlos III of Madrid, Spain

* first.name.lastname@epfl.ch † banchs@it.uc3m.es ‡ v.shneer@hw.ac.uk

Abstract—It has recently been shown that *flexible channelization*, whereby wireless stations adapt their spectrum bands on a per-frame basis, is feasible in practice. In this paper, we propose TF-CSMA/CA, an algorithm for flexible channelization that schedules packets in time and frequency domains. TF-CSMA/CA is a simple extension of the CSMA/CA protocol used by IEEE 802.11. Contrary to existing channelization schemes, it is entirely distributed and it reacts only to packet collisions, successful transmissions and carrier sensing.

With TF-CSMA/CA, when a station is involved in a collision, it performs backoff in both time and frequency domains. Backing off also in the frequency domain allows the transmitters to be much more efficient and aggressive in the time domain, which significantly reduces the severe overheads present with recent 802.11 PHY layers. The main challenge, however, is that the stations need some level of self-organization in order to find spectrum bands of variable widths that minimize interference, while still efficiently using the available spectrum.

Using analysis and simulations, we show that such an extension of CSMA/CA to the frequency domain drastically improves both throughput and fairness. Notably, it enables the stations to find interference-free spectrum bands of appropriate size using no communication – relying only on collisions and successes as implicit signals.

I. INTRODUCTION

Future wireless networks will be able to use *flexible channelization*, whereby the spectrum consumed by each station can be adapted on a per-frame basis. Recently, significant progress has been made in system design, which has shown that flexible channelization is feasible in practice [1], [2], [3], [4], [5]. It is known that this paradigm has the potential to drastically increase the efficiency, fairness and load-balancing properties of wireless networks [6], [2], [7], [3]. In particular, it provides the following advantages. First, adding frequency-domain decisions to the contention resolution process can mitigate the severe time-domain overheads of 802.11, which are exacerbated by recent PHY layers. Second, adapting the amount of consumed spectrum becomes crucial to avoid interference in recent 802.11 amendments such as 802.11ac, which can use large channel bandwidths (up to 160 MHz) and currently requires very careful spectrum planning [8]. Third, modulating spectrum on a per-frame basis departs from the usual static channel assignment perspective, and enables spectrum-allocation schemes to finely adapt to instantaneous traffic loads.

Despite important promises in terms of performance improvements, finding efficient schedules in time and frequency

The work presented in this paper was supported by the Swiss National Science Foundation under grant number 200021-146423.

domains is difficult. It requires that the stations reach some level of coordination, because for each frame they need to choose “time-spectrum blocks” that (i) do not overlap (to avoid interference) and (ii) consume as much of the available spectrum as possible (to maximize performance). For this reason, to the best of our knowledge, all schemes for flexible channelization proposed so far rely on different forms of explicit signaling, synchronization, spectrum scanning or central control, in order to coordinate neighboring stations and efficiently organize transmissions (see e.g., [6], [2], [7] – this is also the case for more traditional spectrum assignment schemes operating at slower timescales, e.g., [9], [10]). Employing such extra signaling introduces extra overhead and complexity, and typically adapts poorly to variable traffic.

We follow a different approach for scheduling packets in time and frequency, which is completely decentralized and requires no synchronization, explicit signaling, control traffic, nor spectrum scans. We propose TF-CSMA/CA, an extension of the time-domain CSMA/CA backoff mechanism of 802.11 to the frequency domain. In addition to the well-known contention window and backoff counter used in the time domain, TF-CSMA/CA also adjusts dynamically the channel bandwidth and center frequency used for each frame, which determine the spectral-domain behavior. When a station is involved in a collision, it hops to another spectrum band and (with a certain probability) decreases both its time-domain aggressiveness and its (average) spectrum consumption. In contrast, when a station experiences a successful transmission, it remains in its current spectrum band with a large probability, and it increases its (average) spectrum consumption with a small probability.

TF-CSMA/CA respects the design and engineering principles of 802.11: it is a purely random-access mechanism that adapts its time-spectrum aggressiveness based only on transmission outcomes (collisions or successes) and carrier sensing. Although the proposed additional decision rules are relatively simple to describe, we will see that they produce non-trivial self-organizing behaviors, whereby stations avoid interference while efficiently using the available spectrum in both time and frequency domains.

Compared to time-domain random access, TF-CSMA/CA provides several important advantages. First, it drastically reduces the inefficiencies caused by the recent PHY layers of 802.11n and 802.11ac. These amendments deliver up to multi-gigabit raw transmission rates, by using techniques such as MU-MIMO, aggressive modulations, and larger channel

bandwidths (up to 40 MHz for 802.11n and up to 160 MHz for 802.11ac). Although these techniques drastically reduce the time required to transmit a frame, they also increase correspondingly the time-domain overheads due to backoff, acknowledgments, PHY preambles, and other MAC overheads. To mitigate this, 802.11n and 802.11ac amendments have the ability to use frame-aggregation mechanisms, in order to increase the transmission durations. The sizes of the aggregated frames can reach up to 65 kB for 802.11n and up to 4.5 MB for 802.11ac [8]. Although heavy aggregation increases efficiency, it does not help applications producing chatty traffic, or real-time traffic such as video, VoIP or gaming, which cannot afford to wait for large buffers to fill up. In contrast, TF-CSMA/CA drastically reduces these inefficiencies, by (i) reducing the channel width in case of interference (thus reducing the fraction of time consumed by overheads, as reducing the bandwidth increases the transmission duration while maintaining the same overheads) and (ii) being much more aggressive in the time domain (it is able to use minimum contention windows as low as $CW_{min} = 2$ while maintaining excellent fairness, compared to $CW_{min} = 16$ with current 802.11).

In addition to improving efficiency, TF-CSMA/CA also serves to dynamically find non-overlapping channels under interfering networks. Indeed, the use of larger channel widths makes it increasingly difficult to assign non-overlapping channels to neighboring networks (in the US, there is currently only one contiguous 160 MHz band available in the 5.17-5.33 GHz range). 802.11ac can use different channel widths of 20, 40, 80 and 160 MHz and can decide to use channel bonding on a per-frame basis. However, this decision amounts only to deciding whether to employ or not the non-primary channel, and it offers only limited additional flexibility because the primary channel remains fixed. In fact, 802.11ac requires very careful spectrum planning in order to manage interference when large channel widths are employed [8]. TF-CSMA/CA finds interference-free schedules and spectrum allocations directly at the MAC layer, as determined by instantaneous traffic loads.

In summary, the main contribution of this paper is the design and analysis of a mechanism for scheduling packets in time and frequency domains without requiring any form of control traffic. Notably, we will see that even without synchronization, the stations can self-organize to find variable-width spectrum bands that avoid interference while efficiently using the available spectrum.

We organize the remainder of the paper as follows. In the next section, we give some background on the time-domain CSMA/CA mechanism used by 802.11, and explore some tradeoffs involved with packet scheduling. We present TF-CSMA/CA in Section III. In Section IV, we analyze TF-CSMA/CA and show that it converges to interference-free spectrum allocations. Then, in Section V, we use packet-level simulation to thoroughly evaluate the performance of TF-CSMA/CA, both in terms of throughput and short-term fairness, in a wide variety of settings. Finally, we present related work in Section VI and give concluding remarks in Section VII.

II. BACKGROUND AND MOTIVATION

A. The IEEE 802.11 Distributed Coordination Function

To arbitrate transmissions and avoid collisions, 802.11 specifies a distributed coordination function (DCF) based on CSMA/CA. When a station receives a new packet for transmission from the upper layer, it selects a *backoff counter* BC uniformly at random in $\{0, \dots, CW - 1\}$, where CW denotes the contention window and is initially set to a minimum value CW_{min} . The backoff mechanism employs a discrete time scale; for each time slot during which the medium is sensed to be idle (i.e., below the carrier-sensing threshold), the station decreases its backoff counter BC by 1. When the medium is sensed busy, the station freezes its backoff counter until the medium is sensed idle again for a duration equal to DIFS (DCF Interframe Space). The station transmits when the backoff counter reaches 0. If the destination station successfully receives the frame, it waits for a duration equal to SIFS (Short Interframe Space) and replies with an ACK. If there is a collision (detected by a missing ACK), this is interpreted as contention and the transmitting station reduces its aggressiveness by doubling CW (up to a CW_{max} value). It then repeats the process.

The time slot duration has to last long enough to perform reliable carrier-sensing (i.e., measure the energy level), switch the RF front-end from receiving to transmitting, and account for possible propagation delays. It appears that these durations are mostly incompressible; the 802.11a/g/n/ac amendments have been using time slot durations given by $t_{slot} = 9 \mu s$ for more than a decade. Similarly, SIFS needs to account for the time required to process the incoming frame and to switch the mode of the RF front-end to transmit the ACK. 802.11a/n/ac use SIFS durations given by $t_{SIFS} = 16 \mu s$. These time constraints also propagate to DIFS, which is set to SIFS + 2 time slots and is equal to $t_{DIFS} = 34 \mu s$ for 802.11a/n/ac. Finally, each frame starts with the transmission of a PHY preamble, which is required to detect and to decode frame transmissions, as well as to set the spectrum and modulation parameters. In total, 802.11ac uses PHY preambles lasting for durations of $t_{PHY} = 44 \mu s$ [8].

Let us define the (*normalized*) *throughput* (or *efficiency*) of a medium access control protocol as the product of (i) the fraction of time and (ii) the fraction of spectrum that are used for successful transmission of payload traffic. Since 802.11 uses 100% of its channel, its efficiency is only determined by its time-domain operation. To analyze the efficiency of 802.11 as a function of the PHY rate, we can adopt a simple analytical model like the one proposed by Tan et al. [2]. When there is only one transmitting station (and thus no collision), the average value of BC , which we denote by \overline{BC} , is given by $\overline{BC} = (CW_{min} - 1)/2$. We can thus easily compute the efficiency as

$$\text{eff}_{802.11} = \frac{t_{data}}{t_{DIFS} + \overline{BC}t_{slot} + t_{PHY} + t_{data} + t_{SIFS} + t_{ACK}},$$

where t_{data} denotes the time required to transmit the payload and t_{ACK} is the total time required to send the ACK. In Figure 1




Fig. 1. Efficiency of the 802.11 DCF as a function of the PHY rate, for several frame sizes (left) and different numbers of stations N (right). The results on the left subplot are computed analytically for $N = 1$. The results on the right subplot are obtained by simulation with frames of 1 kB.

(left), we show this throughput for different packet sizes as a function of the physical data rate. Although faster transmission rates reduce the total time required for transmitting a frame, they exacerbate the time-domain overheads explained above: when sending 1 kB frames with a PHY rate of 600 Mbps (the maximum rate achievable with 802.11n, but well below the rates achievable with 802.11ac), the efficiency is below 10%. This is also true when the number N of contending stations is larger, as shown in Figure 1 (right) using simulation results (we give more details on our simulator in Section V).

B. Improving Efficiency

We now present two techniques for improving efficiency, which are used by TF-CSMA/CA.

1) *Reducing Backoff Durations:* Current 802.11 amendments use $CW_{min} = 16$. One obvious solution for improving efficiency is to reduce the overhead due to the backoff process, by employing smaller contention windows (i.e., smaller CW_{min} values). Of course, there are good reasons for employing a reasonably large CW_{min} . If the stations transmit too aggressively, they can increase the collision probability (harming the overall efficiency) and even cause starvation.

Too small values for CW_{min} can cause poor short-term fairness (i.e., fairness evaluated on short time horizons), as some stations might starve for long durations before successfully sending a packet. In order to quantify this short-term (un)fairness and starvation effect, we define the *inter-transmission time* ITX as the time duration between two successive successful transmissions of a given station. A scheduling algorithm is perfectly short-term fair (and prevents starvation) if ITX is constant and equal for all the transmissions of all stations. We can therefore use the standard deviation of ITX (over all inter-transmissions of all stations) as a measure of unfairness, which we call σ_{ITX} . The larger σ_{ITX} is, the less the protocol is short-term fair, and the more likely it is for the stations to experience starvation¹. In Figure 2, we show σ_{ITX} , as well as the normalized throughput, for several values of CW_{min} with 802.11 and $N = 5$ stations. With the default CW_{min} , 802.11 gets a throughput of about 6.7%. The throughput can be increased to more than 10% by

¹Our measure of short-term unfairness follows what has been proposed in previous studies on MAC layer short-term fairness, which also use inter-transmission times [11]. Note that short-term fairness implies long-term fairness, whereas the converse is not true.




Fig. 2. Variation of the inter-transmission durations and throughput as functions of CW_{min} , for 802.11 with $N = 5$ stations sending frames of 1 kB, using a PHY rate of 600 Mbps. Small CW_{min} values increase the throughput but correspond to situations where a station monopolizes the medium for long durations. The hatched region ($CW_{min} < 4$) corresponds to complete starvation, where some stations could not experience a single successful transmission over the whole simulation time.

decreasing CW_{min} . The cases with small CW_{min} , however, correspond to situations where a station monopolizes the medium for long durations (indicated by large σ_{ITX} values)². In the extreme cases where $CW_{min} < 4$, some stations could not experience any successful transmission at all during the whole simulation time (which is set to one second in this case).

2) *Using Narrow Channels for Multiple Stations:* Even with dangerously small CW_{min} and backoff durations, 802.11 still obtains relatively low efficiencies (about 10% for the example shown in Figure 2). A solution to further improve efficiency is to reduce channel bandwidths; narrow channels require longer durations to send a given number of payload symbols, and thus amortize the time-domain overheads. This idea has been previously proposed by Chintalapudi et al. [3] and others. Note, however, that for a single station, simply dividing a wide-band channel into several narrow-band channels to send several longer frames effectively requires buffering more payload bits and is thus equivalent to performing aggregation on the original wide-band channel. However, when multiple stations compete for access, it is possible to increase efficiency by having each station transmit in parallel on different narrow bands (without requiring more payload to be buffered).

In the remainder of the paper, we show that it is possible to implement the two above-mentioned solutions (reduction of backoff durations and narrow channels for multiple stations), by extending the contention resolution process of 802.11 to the frequency domain. Backing off in the frequency domain enables TF-CSMA/CA to use very small CW_{min} values and reach efficiencies much higher than 802.11 (or any other time-domain scheduling mechanism), while maintaining excellent fairness and removing the starvation problem existing for 802.11 with small CW_{min} values. Overall, when $N = 1$, the efficiency gain comes only from a reduction in backoff duration. When $N > 1$, the gain comes from a combination of reduced backoff durations and reduced overheads over narrower bandwidths. Notably, we will see that when $N > 1$, the stations naturally converge to operating points where they use an average amount of spectrum proportional to $1/N$ – without knowing the number of stations N .

²We use $CW_{max} = 1024$ in these experiments. Reducing CW_{max} together with CW_{min} (i.e., by using a fixed number of backoff stages) avoids starvation but increases the collision rates and produces worse throughputs than the default configuration.

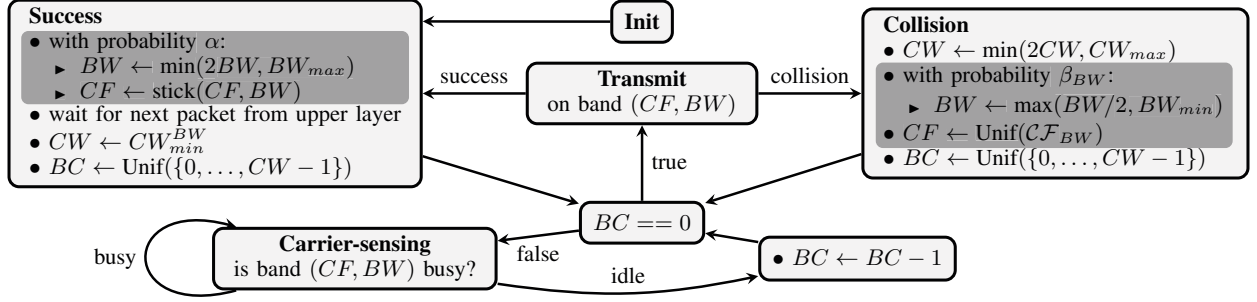


Fig. 3. Finite state machine of TF-CSMA/CA, as running in a station. In addition to the contention window CW and backoff counter BC maintained by 802.11, TF-CSMA/CA also maintains the current center frequency (CF) and channel bandwidth (BW). The changes with respect to 802.11 are highlighted in dark gray. The function $\text{stick}(CF, BW)$ returns the center-frequency in \mathcal{CF}_{BW} that is the closest to CF (breaking ties uniformly at random).

III. SCHEDULING IN TIME AND FREQUENCY DOMAINS

We now present TF-CSMA/CA. We start by introducing some necessary notations, and then present the algorithm itself.

A. System Model and Notations

We assume that the stations use a flexible baseband design such as the one proposed in [5], which lets the receivers detect the center frequency and bandwidth used by incoming transmissions (e.g., using PHY-layer preambles) and process frames accordingly. We focus on the case where the stations use contiguous chunks of spectrum (i.e., without fragmentation), which is simpler in terms of system design. Hence, with TF-CSMA/CA, in addition to its contention window CW and backoff counter BC , each station also maintains its current center frequency CF and bandwidth BW . These parameters are the spectrum parameters used at any point in time for packet transmissions and carrier sensing. To describe spectrum constraints, we denote by \mathcal{CF}_{BW} the set of center frequencies that can be used with a given bandwidth BW (for example, in the 5.17-5.33 GHz band, we can have $\mathcal{CF}_{160\text{ MHz}} = \{5.25\text{ GHz}\}$ and $\mathcal{CF}_{80\text{ MHz}} = \{5.21\text{ GHz}, 5.29\text{ GHz}\}$, etc.). We write BW_{min} and BW_{max} for the minimum and maximum available bandwidths, respectively (e.g., in 802.11ac settings, we can have $BW_{min} = 20\text{ MHz}$ and $BW_{max} = 160\text{ MHz}$). For simplicity of exposition, we assume throughout the paper that bandwidths are powers of 2, so that switching to the next larger (resp. next smaller) bandwidth is obtained by multiplying (resp. dividing) the current bandwidth by 2 (in a similar way as CW for 802.11). Finally, TF-CSMA/CA employs a value of CW_{min} that depends on the current bandwidth BW , and which we denote CW_{min}^{BW} .

B. Description of TF-CSMA/CA

TF-CSMA/CA is based on the following two observations:

- *Reaction to collisions in the frequency domain:* In the presence of contention, the stations should preferably separate their transmissions in the frequency domain. This is because orthogonal transmissions in the frequency domain enable simultaneous transmission of packets, and narrow bands reduce the time-domain overheads mentioned in Section II. Therefore, upon experiencing a

collision, a station should seek another spectrum band by changing its center frequency. In addition, frequent collisions should be interpreted as a signal that the station is using too much spectrum and should thus reduce its channel bandwidth to be able to find a free spectrum band.

- *Reaction to successes in the frequency domain:* Repeated successes indicate that a station operates alone in its spectrum band. The station should thus remain in this band or, with a small probability, try to increase its bandwidth in order to check if it is possible to use more spectrum.

We show the operation of TF-CSMA/CA at a single station as a finite-state machine in Figure 3. The stations start in any arbitrary combination of center frequency and bandwidth. The time-domain backoff mechanism is strictly equivalent to that of 802.11. Upon receiving a data packet from the upper layer, the station draws BC uniformly at random in $\{0, \dots, CW_{min}^{BW} - 1\}$. It then performs carrier-sensing on the current band that is specified by the tuple (CF, BW) . For each slot during which the band is sensed idle, the station decreases BC by 1 (the time slots have the same duration as for 802.11). When BC reaches 0, the station attempts a transmission. If the destination station successfully receives the frame, it sends an ACK on the same band (CF, BW) (after a SIFS duration). If the transmission collides (as detected by a missing ACK), the station doubles CW . If the transmission succeeds, the station sets CW to CW_{min}^{BW} .

The differences compared to 802.11 are shown in dark gray on Figure 3 and consist of the following additional actions. If a collision occurs, the station re-selects a new center frequency CF uniformly at random. In addition, it divides BW by 2 with a probability β_{BW} that depends on the current bandwidth. In contrast, in the event of a successful transmission, the station doubles BW with a probability α . Finally, if BW changes because of a successful transmission, the station also re-selects a new CF that is as close as possible to its current CF . This action is represented by the “stick” function in Figure 3: the function $\text{stick}(CF, BW)$ simply returns the center frequency in \mathcal{CF}_{BW} that is the closest to CF (breaking ties uniformly at random).

Note that the parameters BW and CF play roles in the frequency domain that are similar to CW and BC in the time domain. BW determines aggressiveness in the frequency domain, similarly to CW in the time domain. Likewise, CF and BC determine the localizations of the resource chunks consumed in the frequency and time domains, respectively.

C. Time-Domain Behavior and Configuration of CW_{min}

As we will see in Sections IV and V, the stations running TF-CSMA/CA converge to using non-overlapping spectrum bands that are well spread over the entire available spectrum. Although TF-CSMA/CA uses the same time-domain mechanism as 802.11, the fact that it can self-organize in the spectral domain makes it possible to configure the time-domain backoff mechanism in a more efficient way.

When the stations use large bandwidths, TF-CSMA/CA attempts to separate their transmissions in the frequency domain, by reducing their bandwidth and letting them transmit on orthogonal subbands. As a result, contention can be resolved entirely in the frequency domain and the stations operating with large bandwidths can be much more aggressive in the time domain (i.e., employ very short backoff durations) without risking to starve other stations. In contrast, when the stations already use narrow bandwidths (for example, if there are many stations using orthogonal bands with the minimum bandwidth BW_{min}), some stations may have to share some spectrum bands. Therefore, in this case, the stations should also use the time domain to separate their transmissions (i.e., employ reasonably long backoff durations – note, however, that the time spent in backoff represents a smaller overhead when using small bandwidths).

Overall, the importance of the time domain in the contention-resolution process should thus depend on the bandwidth. In particular, CW_{min}^{BW} should be a decreasing sequence of BW . In this paper, we propose to use CW_{min} values given by $CW_{min}^{BW} = \lceil \frac{16}{BW/BW_{min}} \rceil$. This sequence is such that $CW_{min}^{BW_{min}} = 16$, which corresponds to the current default CW_{min} employed by 802.11. In 802.11ac settings, the corresponding sequence is $CW_{min}^{20\text{ MHz}} = 16$, $CW_{min}^{40\text{ MHz}} = 8$, $CW_{min}^{80\text{ MHz}} = 4$ and $CW_{min}^{160\text{ MHz}} = 2$.

D. Mechanism for Adapting Contention Bandwidth

TF-CSMA/CA, as described above, uses spectrum efficiently, but it can create problematic situations in terms of short-term fairness. When several stations transmit simultaneously on orthogonal narrow bands, it is possible that a given wide band, which contains some of these narrow bands, rarely becomes entirely free. Thus, if a station is contending on this wide band, it might have to freeze its backoff counter for long durations. To avoid this undesirable situation, TF-CSMA/CA uses the following additional mechanism (not shown in Figure 3), which incurs no performance penalty but improves short-term fairness.

Bandwidth Adaptation after Carrier Sensing: Each station halves its bandwidth BW with a small probability $\epsilon \ll 1$

after having sensed the medium busy due to a transmission by another station.

Although this mechanism is simple and requires no additional state, it ensures that each station waits on average no more than $1/\epsilon$ transmissions from other stations before reducing the bandwidth on which it contends. It is useful when there are many stations, as it ensures that each station adapts the amount of spectrum on which it contends, without actually experiencing a collision (or waiting for one).

IV. ANALYSIS AND CONFIGURATION

In this section, we first introduce a Markov-chain model to study the spectral self-organization of TF-CSMA/CA. The main purpose of this analysis is to show that a simple frequency-domain scheduling scheme based on random access such as TF-CSMA/CA can exhibit self-organization. We conclude from the analysis that if the parameter α is small enough, the stations spend the vast majority of their time in states without interference. Then, in Section IV-B, we use the results of the analysis, as well as arguments related to the transient regime of TF-CSMA/CA, to configure the parameters of the algorithm, namely β_{BW} , α and ϵ . In particular, the arguments related to the transient regime consider the tradeoff between exploration (i.e., converging quickly and thus adapting to variable traffic) and exploitation (i.e., remaining in good states as much as possible to optimize performance in steady state).

A. Steady-State Model of Spectrum Consumption

Let $C := BW_{max}/BW_{min}$ be the number of smallest orthogonal subbands. For simplicity of exposition, we restrict our analysis to the case where $N = C$. For these values, there exists exactly one state without interference³. We first detail our Markov-chain model and provide an example where $N = 2$ and $C = 2$ and then we extend our results to general N . We consider the case where the N stations belong to a single contention domain, and we assume that the channel quality is sufficiently high so that packet losses are due only to collisions. Without loss of generality, we set $BW_{min} = 1$ and $BW_{max} = C$. In addition, we make a modeling assumption similar to the decoupling assumption introduced by Bianchi in the time domain [12], and we assume that every station attempts a transmission with a fixed probability p at any given time slot. Let $n_i, 1 \leq i \leq C$, denote the number of nodes using a band that overlaps with the i -th subband of width 1. We build a discrete-time Markov chain whose states represent all the possible patterns according to which the N stations can occupy the spectrum. Precisely, each state belongs to the set $\mathcal{S} := \{n_i : 1 \leq i \leq C, 0 \leq n_i \leq N\}$ ⁴. With TF-CSMA/CA, the stations change their spectral configuration after a transmission attempt with probability α (in case of success) or β_{BW}

³The case $N < C$ corresponds to an easier problem, in terms of finding interference-free assignments, and it can be treated similarly. Note that there does not exist a state without interference when $N > C$. Yet, we will see in Section V that TF-CSMA/CA performs well for all N .

⁴ \mathcal{S} describes the set of all possible states, also if stations could fragment their spectrum. If the stations do not fragment their spectrum (as is the case for TF-CSMA/CA), the possible spectral patterns belong to a subset of \mathcal{S} .




Fig. 4. Markov chain for the case of two stations and two bandwidths. The states are denoted A , B , C and D and the spectrum settings for the two stations within each state are shown by segments. We do not show self-transition probabilities.

(in case of collision). Therefore, the transitions of the Markov chain from one state to the next occur upon a transmission attempt by any one of the stations (following the assumption of geometric backoff durations).

1) *Example with Two Stations and Two Subbands*: It is helpful to first consider the case with two stations and two subbands, as the states can be easily enumerated. In this case there are two bandwidths: one bandwidth corresponding to using all the band (i.e., $BW = BW_{max}$) and the other corresponding to half of the band (i.e., $BW = BW_{max}/2$). The Markov chain is represented in Figure 4. There are four possible states, denoted A , B , C and D : They correspond to the different combinations of spectrum occupation (the spectrum configurations of the two stations are represented by segments in Figure 4).

As there are only two bandwidths, we only need one β_{BW} , as stations can only decrease their bandwidth when $BW = BW_{max}$; hence, we define $\beta := \beta_{BW_{max}}$. The transition probabilities are easy to obtain from the reaction of TF-CSMA/CA to successes and collisions. For example, the transition probabilities from A to B and from A to D are $\frac{1}{2}p\beta^2$, because the other station (the one that does not trigger the state transition) has to transmit (which happens with probability p) and the two stations have to independently choose to reduce their bandwidth (with probability β^2).

In the case under study, the most desirable state is D because there is no frequency-domain interference and the whole spectrum is used in this state. The following theorem states that, if α is small enough, TF-CSMA/CA spends an arbitrarily large fraction of time in state D .

Theorem 1. *Let π_i be the stationary distribution of state $i \in \{A, B, C, D\}$. We have*

$$\pi_D \xrightarrow{\alpha \downarrow 0} 1.$$

Proof. Using the balance equation for D , we get

$$\pi_D = \pi_D(1 - \alpha - p\alpha^2) + \pi_A \frac{1}{2}p\beta^2 + \pi_B \frac{1}{2}p + \pi_C \frac{1}{4}p\beta.$$




Fig. 5. Proportion of time spent in states A , B , C and D for the Markov chain of Figure 4 and a packet-level simulation of TF-CSMA/CA, with the settings $p = 0.05$, $\alpha = 10^{-3}$ and $\beta = 1$.

Let us define $\beta' := \min\{\frac{1}{4}p\beta, \frac{1}{2}p\beta^2\}$. We have

$$(\alpha + p\alpha^2)\pi_D \geq \sum_{i \in \{A, B, C\}} \pi_i \beta',$$

and thus

$$\frac{\pi_D}{\sum_{i \in \{A, B, C\}} \pi_i} \geq \frac{\beta'}{\alpha + p\alpha^2},$$

which concludes the proof. \square

The result of Theorem 1 also holds if $\alpha = 0$, in which case D becomes an absorbing state. However, in this case the chain is no longer ergodic and, for general configurations of N and C , it might remain “stuck” in absorbing states that avoid interference but under-utilize the spectrum. For this reason, TF-CSMA/CA employs a small but non-zero value of α (we elaborate further on this point in Sections IV-B and V). In Figure 5, we show the fraction of time spent in the states A , B , C and D by the Markov chain of Figure 4, as well as by packet-level simulations of TF-CSMA/CA (see Section V for more details on our simulation settings). Although our Markov model makes simplifying assumptions, it correctly captures the tendency of TF-CSMA/CA to spend the vast majority of the time in the best possible state in this scenario.

2) *N Stations and N Subbands*: We now extend Theorem 1 to the general case of N subbands (with $C = N$).

Theorem 2. *Let $s^* \in \mathcal{S}$ be the (unique) interference-free state. We have*

$$\pi_{s^*} \xrightarrow{\alpha \downarrow 0} 1.$$

Proof. Let us denote the bandwidth used by a station u in state $s \in \mathcal{S}$ by BW_u^s . We define

$$\mathcal{S}_1 := \left\{ s : \max_{u \in \{1, \dots, N\}} \{BW_u^s\} \leq 2 \right\} \setminus \{s^*\},$$

which is the set of states that are one transition away from s^* .

For any two states s and s' , let $P_{s \rightarrow s'}$ denote the transition probability from s to s' . Now, when the network is in state s^* and a station transmits, there could be a random number, say k , of other stations transmitting at the same time, and k follows a binomial distribution of parameters $N - 1$ and p . Then, the network remains in state s^* if and only if none of the $k + 1$ transmitting stations decides to double its bandwidth.

Therefore, the probability of staying in state s^* is

$$\begin{aligned}
P_{s^* \rightarrow s^*} &= \sum_{k=0}^{N-1} \binom{N-1}{k} p^k (1-p)^{N-1-k} (1-\alpha)^{k+1} \\
&= (1-\alpha) \sum_{k=0}^{N-1} \binom{N-1}{k} (p(1-\alpha))^k (1-p)^{N-1-k} \\
&= (1-\alpha) (p(1-\alpha) + 1-p)^{N-1} \\
&= (1-\alpha) (1-p\alpha)^{N-1} \\
&\geq (1-\alpha)^N \\
&\geq 1-N\alpha.
\end{aligned}$$

We can thus use the balance equation for s^* to obtain

$$\pi_{s^*} \geq \pi_{s^*} (1-N\alpha) + \sum_{s \in \mathcal{S}_1} \pi_s P_{s \rightarrow s^*}.$$

Let $\beta_{min} := \min_{BW} \{\beta_{BW}\}$. It is easy to see that

$$P_{s \rightarrow s^*} \geq C^{-N} p^{N-1} (\beta_{min})^N$$

for any state s in \mathcal{S}_1 . We thus have

$$\pi_{s^*} \geq \pi_{s^*} (1-N\alpha) + \sum_{s \in \mathcal{S}_1} \pi_s C^{-N} p^{N-1} (\beta_{min})^N,$$

from which we obtain

$$\frac{\pi_{s^*}}{\sum_{s \in \mathcal{S}_1} \pi_s} \geq \frac{p^{N-1} (\beta_{min})^N}{C^N N \alpha}$$

and thus, for any state $s \in \mathcal{S}_1$,

$$\pi_s \leq A(N\alpha) \pi_{s^*}, \quad (1)$$

with $A := C^N / (p^{N-1} (\beta_{min})^N)$.

We now need to iterate this reasoning over the states that are not in \mathcal{S}_1 and need more than one transition to reach s^* . To this end, we extend the definition of \mathcal{S}_1 and define

$$\mathcal{S}_k := \left\{ s : \max_{u \in \{1, \dots, N\}} \{BW_u^s\} = 2^k \right\},$$

for $k \geq 2$. Now, for any $k \geq 2$ and any state $s_k \in \mathcal{S}_k$, let \mathcal{N}_{s_k} denote the set of stations that use bandwidth 2^k in s_k . Note that $|\mathcal{N}_{s_k}| > 0$ by construction of \mathcal{S}_k , and so there exists a state $s_{k-1} \in \mathcal{S}_{k-1}$ that is obtained by halving the bandwidth of the stations in s_k that use bandwidth 2^k (and having them use any valid center-frequency). It is again easy to see that

$$P_{s_k \rightarrow s_{k-1}} \geq C^{-N} p^{N-1} (\beta_{min})^N$$

and so from the balance equation of s_{k-1} ,

$$\pi_{s_{k-1}} \geq \pi_{s_k} C^{-N} p^{N-1} \beta_{min}^N.$$

We can now iterate this argument k times and combine it with inequality (1) in order to obtain (noting that $k \leq \lceil \log_2(N) \rceil$)

$$\pi_s \leq A^{\lceil \log_2(N) \rceil} (N\alpha) \pi_{s^*}$$

for any possible state $s \in \mathcal{S}$, which concludes the proof. \square

This shows that, by setting α sufficiently small, we can ensure that TF-CSMA/CA spends an arbitrarily large fraction of the time in the most desirable state. Based on this and other considerations, we next discuss the setting of α , as well as the other parameters of the algorithm.

B. Parameters Configuration

Let us now give some high-level comments on the setting of the parameters of TF-CSMA/CA, namely α , β_{BW} and ϵ .

Let us start with β_{BW} . A collision indicates that a station uses a band that overlaps with another station. In this case, the station should change its center frequency and find a new (hopefully non-overlapping) band and, if it is using a bandwidth that is too large to find a free spectrum band, it needs to reduce it. The average number of collisions needed to reduce BW is given by $1/\beta_{BW}$: this determines the time that a station has to find an interference-free configuration. Therefore, on the one hand, β_{BW} should be sufficiently small so that the stations are given enough time to find an interference-free configuration, if it exists, before reducing their bandwidths. On the other hand, it should not be smaller than needed, as otherwise the stations might lose time looking for an interference-free configuration that does not exist.

Hence, in order to find an appropriate setting for β_{BW} , we need to compute the time needed to find an interference-free configuration for a given bandwidth, in situations where the stations should not reduce their bandwidth. This problem is similar to the one addressed in [13], which analyses the time it takes a balls-into-bins algorithm to find a configuration in which all bins have the same number of balls (in our particular case, one ball). In the algorithm of [13], each ball samples randomly each bin until it finds an empty one. This is similar to our algorithm when we have N stations that are using subbands of bandwidth equal to BW_{max}/N . In our case, when a station is in a non-empty subband, it detects this through a collision and randomly chooses another subband until it finds a free one. According to the analysis of [13], the time it takes to find such a configuration is $\mathcal{O}(N)$.

Based on the above reasoning, we set $\beta_{BW} = c \cdot BW$, for some constant c . This is because a station using bandwidth BW is likely to contend with $\mathcal{O}(1/BW)$ stations, and thus the time needed to find an interference-free configuration will be given by $\mathcal{O}(1/BW)$; hence, in this case we set $\beta_{BW} = \mathcal{O}(BW)$. For the choice of c , we set it such that when a station is using BW_{max} , we have $\beta_{BW_{max}} = 1$ (which is clearly the best configuration for this scenario), which leads to setting $\beta_{BW} = BW/BW_{max}$.

For setting α , based on the analysis of the previous subsection, we note that it should be set to a small value, so that the stations experiencing successful transmissions tend to remain on the same band. Whereas, setting α to a non-zero value enables the stations to reclaim possibly unused spectrum. Based on our evaluations of Section V, we set $\alpha = 10^{-3}$, as we observe that it performs well in all settings.

Finally, we set ϵ based on the following reasoning. If we set $\epsilon = 1/x$, then a station has to wait on average up to x transmissions before halving the bandwidth on which it contends, which means that it might not be able to transmit during this time. Based on this, we set $\epsilon = 10^{-2}$, so that each station waits on average for no more than 100 transmissions before halving its bandwidth.




Fig. 6. Comparison of TF-CSMA/CA with 802.11, optimal TDMA and 802.11 operating with optimal spectrum assignment.




Fig. 7. Short-term unfairness σ_{ITX} and collision probability, for 802.11 and TF-CSMA/CA (denoted “TF”).




Fig. 8. Ratio between the efficiency of TF-CSMA/CA and that of 802.11 for $N = 5$. Larger gains occur for smaller transmission delays.

V. PERFORMANCE EVALUATION

A. Simulation Settings

We developed an event-driven packet-level simulator in Python. Our simulator is very similar to several other simulators that have been used in the past to model various MAC layers (see e.g., [12], [14]). We simulate TF-CSMA/CA with the same per-frame timing overheads described in Section II (time slot, SIFS, DIFS, PHY headers and ACK durations). We assume that each station achieves a physical rate proportional to its channel bandwidth, which corresponds to what is observed empirically [15]. Unless otherwise stated, we use 802.11ac settings and set $BW_{max} = 160$ MHz and $BW_{min} = 20$ MHz (and so the set of bandwidths available is $\{20, 40, 80, 160\}$ MHz). When simulating the 802.11 DCF (i.e., without our extensions to the frequency domain), we use the whole 160 MHz channel and the default configuration $CW_{min} = 16$ and $CW_{max} = 1024$ (i.e., 7 backoff stages). TF-CSMA/CA is also simulated with 7 backoff stages in the time domain. In line with 802.11ac, for each bandwidth BW , we use the set of center frequencies \mathcal{C}_{FBW} such that all available bands of width BW do not overlap [8]. The default parameters, which we use unless otherwise specified, are summarized in Table I. We consider scenarios where all stations are interfering (complete interference graph) and each station always has a (unique) frame to send, except in Section V-D, where we consider scenarios with non-complete interference graphs and non-saturated traffic. Finally, in order to isolate the effects of the random-access mechanism, we assume that there is no error due to channel noise (hence all packet losses happen due to collisions). Unless otherwise stated, we use a physical rate of 600 Mbps and 1 kB frames. Each configuration is evaluated using 10 independent simulation runs lasting at least one second of simulated time (which is much larger than convergence times and typically corresponds to several thousands of transmission attempts).

BW	CW_{min}^{BW}	β_{BW}	PHY rate
160 MHz	2	1	600 Mbps
80 MHz	4	1/2	300 Mbps
40 MHz	8	1/4	150 Mbps
20 MHz	16	-	75 Mbps

TABLE I
DEFAULT PARAMETERS USED FOR SIMULATIONS

B. Efficiency and Fairness

In Figure 6, we show the throughput obtained by TF-CSMA/CA as a function of the number of stations N , and we compare it against several other scheduling mechanisms: (i) “802.11 default” denotes 802.11 operating with its default configuration on the single wide-band channel; (ii) “optimal TDMA” shows the performance obtained with a perfect TDMA scheme that uses N distinct time slots for a network with N stations. This corresponds, for instance, to the steady-state of the scheme proposed by Fang et al. [16] and it is also an upper bound on the performance achievable by any enhancement of 802.11 that does not employ channelization (e.g., [12], [17]); (iii) “optimal spectrum” shows the performance obtained by 802.11 when all stations share the spectrum optimally (i.e., spreading their spectrum as evenly as possible). Obtaining this “optimal spectrum” configuration requires perfect information and is an upper bound of what can be achieved using centralized knowledge for the spectrum assignment.

Clearly, even perfect scheduling in the time domain using TDMA is less efficient than a mixture of time and frequency scheduling. Even though TF-CSMA/CA is completely decentralized, its backoff and frequency-repartition mechanism provides significant performance gains compared to time-domain scheduling, and achieves performance close to what can be obtained using a perfect centralized spectrum assignment. For $N = 1$, throughput is increased by roughly $1.5\times$ due to a reduction in backoff durations – the performance in this case is similar to TDMA (that sends packets back-to-back). For $N > 1$, splitting transmission onto smaller bandwidths provides important gains (up to about $6\times$ in this setting).

Importantly, these gains are not obtained at the price of short-term fairness. In Figure 7, we show σ_{ITX} and the collision probability obtained by TF-CSMA/CA and 802.11. TF-CSMA/CA achieves significantly better short-term fairness and smaller collision probabilities, which is a direct result of the parallelization of transmissions onto orthogonal subbands.

Of course, the gains provided in the frequency domain depend on the PHY rate and frame sizes. In Figure 8, we show the performance increase provided by TF-CSMA/CA for various PHY rates and frame sizes. As expected, the gains are the largest for high PHY rates and small packet sizes (i.e., for small overall transmit delays). Note that for the cases




Fig. 9. Interference and spectrum usage over time for $N = 5$. TF-CSMA/CA balances well the two conflicting goals of minimizing interference while maximizing spectrum usage. Furthermore, convergence to steady state happens within 50 ms.

where large frames are transmitted with low PHY rates, TF-CSMA/CA can be slightly less efficient than 802.11. This is because, in these regimes, the efficiency of 802.11 is already high and the small inefficiency introduced by TF-CSMA/CA in the frequency domain is not compensated by drastic gains in the time domain. However, such configurations are unlikely to happen in practice, as the older Wi-Fi standards providing low PHY rates usually do not use frame sizes larger than 1.5 kB.

C. Interference and Self-Organization

TF-CSMA/CA trades off a very high time-domain inefficiency for some frequency-domain inefficiency. By adapting their spectrum bands, the stations pursue two potentially conflicting goals. On the one hand, they aim to avoid using bands that are also used by other stations. On the other hand, they also try to use as much spectrum as possible in order to maximize their transmission rates.

To quantify these two goals, we define the *interference* as the fraction of the total spectrum that is used by *more than one* station at any given time. Similarly, we define the *spectrum usage* as the fraction of total spectrum used by *at least one* station. In Figure 9, we show the interference and spectrum usage over 300 ms of traces (averaged over 100 independent simulation runs using windows of 1 ms) for $N = 5$ stations. All stations start with the same center frequency and bandwidth BW_{max} . Although 5 stations are competing for access, TF-CSMA/CA converges to interference-free spectrum allocations. The nodes spend little transient time using the same spectrum and rapidly self-organize to use the spectrum efficiently; about 70% of the spectrum is used on average. Furthermore, because TF-CSMA/CA acts at the very fast time-scale of packets (re-)transmissions, convergence to steady-state is fast, within about 50 ms – even though the network started in a highly inefficient state in terms of spectrum assignment.

In order to illustrate how resource allocation is performed in time and frequency domains, we show in Figure 10 the average BW and CW parameters that are selected by the stations, as a function of N . Ideally, if all stations were to perfectly share the spectrum, each station should converge towards using a bandwidth BW given by BW_{max}/N (shown by the dotted line on Figure 10). It turns out that TF-CSMA/CA selects values for BW that are on average very close to optimal. This




Fig. 10. Average values of BW and CW selected by TF-CSMA/CA. The optimal spectrum consumption is given by $160/N$ MHz and shown by the dotted line (partly indistinguishable from the BW curve).

is remarkable, considering that the stations do not know N^5 . Note that for $N \geq 8$, TF-CSMA/CA uses mostly $BW = 20$ MHz, as this is the minimum available bandwidth.

D. Dynamic Traffic and Random Topologies

The probability α of doubling BW after a successful transmission responds to a tradeoff between exploration and exploitation. A small α ensures that the stations spend most of their time in states that minimize interference (see Section V). However, a non-zero α is needed to regain available spectrum (for instance, in case the stations that were using that spectrum have left) and to ensure high overall spectrum utilization. We thus expect that a large α should favor situations with high traffic variability, whereas a small α should improve performance in steady state.

To quantify this effect, we introduce random traffic patterns as follows. The packets are generated by an exogenous on/off process at each station. The “on” durations are exponentially distributed with mean $1/\lambda$ and the “off” durations are exponentially distributed with mean $1/\mu$. In addition, the frame sizes are also exponentially distributed with mean 1 kB. In the following experiment, we set $1/\mu = 100$ ms, and we vary $1/\lambda$ between 1 ms and 4 s. In Figure 11, we show the throughput obtained in these settings, as a function of the resulting *average traffic intensity* μ/λ . As expected, a larger α improves performance when μ/λ is small (bursty traffic). This is because the stations experience little contention and gain from re-using the spectrum more aggressively. However, even when using a relatively small $\alpha = 10^{-3}$ while the traffic is highly dynamic and bursty, TF-CSMA/CA performs at least as well as 802.11 (which uses the whole spectrum band).

This observation still holds when traffic intensity varies not only in time, but also in space. We make an experiment where the N stations are spread uniformly at random on a $100\text{m} \times 100\text{m}$ square and use an interference radius of $R = 30\text{m}$ (that is, two stations separated by a distance less than R cannot transmit successfully at the same time on overlapping bands). In Figure 12, we measure the traffic sent when the number of stations varies between $N = 1$ and $N = 512$, for different average traffic intensities (using the

⁵Note that estimating N is not trivial; for instance, a number of works to find the optimal CW_{min} in 802.11 have designed algorithms to estimate N in some way [12], [17].




Fig. 11. Throughput as a function of average traffic load with $N = 5$ for 802.11 and TF-CSMA/CA (denoted “TF”). All stations have a fixed average “off” duration set to $1/\mu = 100$ ms and their average “on” duration $1/\lambda$ is varied between 1 ms (nearly silent) and 4 s (nearly backlogged).




Fig. 12. Traffic sent over 5 seconds of simulated time, as a function of the average spatial density of the network topology.

default value $\alpha = 10^{-3}$). When the stations do not suffer from contention, TF-CSMA/CA and 802.11 offer comparable performances. The gain offered by TF-CSMA/CA increases with traffic intensities and spatial densities.

E. Comparison with FICA

We close this section by providing a comparison with the frequency-domain backoff scheme proposed by FICA [2]. With FICA, the spectrum band is divided into several subchannels, and each station can use one or several subchannel(s) (not necessarily contiguous). FICA introduces a form of RTS/CTS signaling, and the transmissions occur in rounds; all transmitting stations have to simultaneously send an RTS signal and the receiving station elects winner(s) for each subband and announces them using a CTS signal. Note that, compared to FICA, TF-CSMA/CA is considerably simpler, as it does not require any signaling or synchronized transmissions.

The authors of FICA recommend splitting payload traffic into 1.6 kB frames to send over each subchannel for a case with a PHY rate of 580 Mbps and 14 subchannels [2]. In our case we use only 8 subchannels, hence we scale this threshold correspondingly and configure FICA to send 2.8 kB frames on each subchannel. Note that this means that FICA might need to access up to 22.4 kB (8×2.8 kB) of payload traffic in the upper layer’s buffer, when a station decides to transmit on all subchannels simultaneously. We therefore consider two scenarios that correspond to two different saturation levels: (i) the upper layer’s buffer always contains 22.4 kB of payload




Fig. 13. Comparison of TF-CSMA/CA (denoted “TF”) and FICA, for two different amounts of payload traffic available in the upper layer’s buffer.

traffic (in which case, FICA can send up to 8 frames of 2.8 kB simultaneously and TF-CSMA/CA can send a unique frame of 22.4 kB); and (ii) the buffer always contains 2.8 kB of payload traffic (in which case FICA sends a 2.8 kB frame on a unique subchannel and TF-CSMA/CA also sends a 2.8 kB frame on its unique band). For FICA, we use 16 subcarriers in each subchannel for contention resolution and, in the 22.4 kB case, we use the proposed AIMD algorithm for choosing the number of subchannels used for transmission.

We show the results in Figure 13. For both saturation levels, TF-CSMA/CA outperforms FICA, even though TF-CSMA/CA does not rely on control traffic or synchronization primitive in order to organize transmissions. This is mainly due to the fact that FICA introduces extra per-frame overheads for the RTS/CTS signaling in order to explicitly organize transmissions. Such coordination is not needed by TF-CSMA/CA, because it provides self-organization in a purely random-access fashion. Note that, for large N , FICA performs better when only 2.8 kB is available in the transmit buffer. This is because, in these regimes, it is nearly always beneficial to use a single subchannel. Note also that, for large frames, TF-CSMA/CA performs significantly better for $N = 1$ compared to $N = 2$. This is because a unique station always uses the full spectrum band (which is efficient in this case), whereas a scenario with more stations can be less efficient due to the randomness of self-organization. In contrast, for larger N values, TF-CSMA/CA stations almost always contend using the smallest bandwidth, which is less challenging in terms of self-organization.

VI. RELATED WORK

Several recent works have shown the practical feasibility of flexible (or fine-grained) channelization [1], [2], [3], [4], [5]. Among these, [1] and [5] propose schemes to schedule packets in time and on variable amounts of spectrum, but both algorithms rely on a central controller to take the scheduling decisions. FICA [2] proposes a backoff mechanism acting in the frequency domain but, as explained in Section V, it relies on explicit signaling and synchronized transmissions. Relying on synchronized transmissions can introduce additional inefficiencies (e.g., if the traffic is such that the payloads do not have the same durations). [3] presents a novel radio design

that enables the 802.11 DCF function to run independently on several narrow channels. However, the mechanism to decide the subchannel(s) on which each link should contend needs to measure the residual airtime and number of contenders in all subchannels. This approach does not let the stations choose their spectrum on a per-packet basis and is closer to spectrum-assignment schemes acting at slower timescales (e.g. [9], [10]). In addition, it increases efficiency without requiring buffering only if there are enough stations contending: When there are only one or a few stations, splitting the wideband in several narrow bands to send several longer frames in parallel requires buffering more payload traffic. In contrast, in these regimes, TF-CSMA/CA increases efficiency by letting the stations be more aggressive in the time domain.

Some works consider the problem of scheduling packets in the context of flexible channelization. [6] studies the optimization problem of efficiently packing time-spectrum blocks, but the proposed algorithms require centralized control. [18] considers running independent 802.11 DCF on several subchannels (similar to [3]), but does not address the problem of deciding how much spectrum should be used by each station. Recently, [7] proposed an algorithm for scheduling packets in time and frequency domains, but here too the proposed mechanism relies on additional control signals and synchronization (as transmissions occur in synchronized rounds). [19] proposes a generalization of CSMA/CA to contend on several subchannels with variable intensities in the time domain. However, the stations always use a fixed channel for transmission and do not modulate their access intensity in the spectral domain.

Different techniques have been proposed to reduce the time-domain inefficiencies of Wi-Fi [20], [21], [22]. For example, [20] shows that it is possible to improve backoff overheads by resolving contention using signaling on OFDM subcarriers in the frequency domain. However, in these cases, the stations always use a fixed channel and, although contention *resolution* can be done in the frequency domain, the stations do not perform *backoff* in the frequency domain. Furthermore, we have seen in Section V that even completely removing the backoff overhead (using perfect TDMA) provides little gain compared to separating transmissions in the frequency domain.

In contrast to the above works, TF-CSMA/CA modulates its aggressiveness and decides on the schedules in time and frequency domains in a purely random-access fashion, using only collisions, successes and carrier sensing as implicit signals.

VII. CONCLUSIONS

We have proposed TF-CSMA/CA, a scheduling algorithm that adjusts both time and frequency access intensities in a random-access fashion. In contrast to existing schemes acting in time and frequency domains, TF-CSMA/CA is completely decentralized and reacts only to collisions, successes and carrier sensing. Overall, relying only on transmission outcomes provides a simple and effective way to assign channels to stations directly at the MAC layer, in a way that departs from the usual “reservation-based” view of spectrum usage, but that is instead determined by instantaneous traffic loads,

just like CSMA/CA in the time domain. We have shown that (i) it self-organizes in the spectral domain, efficiently packing the spectrum and avoiding interference; (ii) although it is completely decentralized, it outperforms perfect time-domain scheduling, and (iii) it provides performance close to what is achievable when a centralized controller directly assigns spectrum to 802.11 nodes in a perfect (but monolithic) fashion.

There are several interesting directions that remain to be explored. In particular, we would like to study the transient regime of TF-CSMA/CA, and characterize how close it is from being optimal (compared to centralized schemes operating with full information).

REFERENCES

- [1] S. Rayanchu, V. Shrivastava, S. Banerjee, and R. Chandra, “Fluid: improving throughputs in enterprise wireless lans through flexible channelization,” in *ACM MobiCom*, 2011.
- [2] K. Tan, J. Fang, Y. Zhang, S. Chen, L. Shi, J. Zhang, and Y. Zhang, “Fine-grained channel access in wireless lan,” in *ACM SIGCOMM ’10*.
- [3] K. Chintalapudi, B. Radunovic, V. Balan, M. Buettner, S. Yerramalli, V. Navda, and R. Ramjee, “Wifi-nc: Wifi over narrow channels,” in *USENIX NSDI ’12*.
- [4] E. Chai, J. Lee, S.-J. Lee, R. Etkin, and K. G. Shin, “Building efficient spectrum-agile devices for dummies,” in *ACM MobiCom*, 2012.
- [5] S. Yun, D. Kim, and L. Qiu, “Fine-grained spectrum adaptation in wifi networks,” in *ACM MobiCom*, 2013, pp. 327–338.
- [6] T. Moscibroda, R. Chandra, Y. Wu, S. Sengupta, P. Bahl, and Y. Yuan, “Load-aware spectrum distribution in wireless lans,” in *IEEE ICNP’08*.
- [7] S. K. Fayaz, F. Zarinni, and S. Das, “Ez-channel: A distributed mac protocol for efficient channelization in wireless networks,” *Ad Hoc Networks*, 2015.
- [8] M. Gast, *802.11Ac: A Survival Guide*, 1st ed. O’Reilly Media, Inc., 2013.
- [9] B. Kauffmann, F. Baccelli, A. Chaintreau, V. Mhatre, K. Papagiannaki, and C. Diot, “Measurement-based self organization of interfering 802.11 wireless access networks,” in *IEEE INFOCOM*, 2007, pp. 1451–1459.
- [10] J. Herzen, R. Merz, and P. Thiran, “Distributed spectrum assignment for home wlangs,” in *IEEE INFOCOM*, 2013, pp. 1573–1581.
- [11] G. Berger-Sabbatel, A. Duda, O. Gaudoin, M. Heusse, and F. Rousseau, “Fairness and its impact on delay in 802.11 networks,” in *IEEE GLOBECOM*, vol. 5, 2004, pp. 2967–2973.
- [12] G. Bianchi, “Performance analysis of the ieee 802.11 distributed coordination function,” *IEEE JSAC*, 2000.
- [13] P. Berenbrink, K. Khodamoradi, T. Sauerwald, and A. Stauffer, “Balls-into-bins with nearly optimal load distribution,” in *ACM SPAA’13*.
- [14] M. Durvy, O. Dousse, and P. Thiran, “Self-organization properties of csma/ca systems and their consequences on fairness,” *IEEE Transactions on Information Theory*, vol. 55, no. 3, pp. 931–943, 2009.
- [15] R. Chandra, R. Mahajan, T. Moscibroda, R. Raghavendra, and P. Bahl, “A case for adapting channel width in wireless networks,” in *ACM SIGCOMM*, 2008, pp. 135–146.
- [16] M. Fang, D. Malone, K. R. Duffy, and D. J. Leith, “Decentralised learning macs for collision-free access in wlangs,” *Wireless networks*, vol. 19, no. 1, pp. 83–98, 2013.
- [17] P. Patras, A. Banchs, P. Serrano, and A. Azcorra, “A control-theoretic approach to distributed optimal configuration of 802.11 wlangs,” *IEEE Transactions on Mobile Computing*, vol. 10, no. 6, pp. 897–910, 2011.
- [18] R. Maheshwari, J. Cao, A. P. Subramanian *et al.*, “Adaptive channelization for high data rate wireless networks,” *Tech. Report*, 2009.
- [19] H. Kwon, H. Seo, S. Kim, and B. G. Lee, “Generalized csma/ca for ofdma systems: protocol design, throughput analysis, and implementation issues,” *IEEE Transactions on Wireless Communications*, 2009.
- [20] S. Sen, R. Roy Choudhury, and S. Nelakuditi, “No time to countdown: Migrating backoff to the frequency domain,” in *ACM MobiCom 2011*.
- [21] E. Magistretti, K. K. Chintalapudi, B. Radunovic, and R. Ramjee, “Wifi-nano: reclaiming wifi efficiency through 800 ns slots,” in *ACM MobiCom 2011*.
- [22] X. Feng, J. Zhang, Q. Zhang, and B. Li, “Use your frequency wisely: Explore frequency domain for channel contention and ack,” in *IEEE INFOCOM*, 2012, pp. 549–557.