



ELSEVIER

Contents lists available at SciVerse ScienceDirect

Computer Networks

journal homepage: www.elsevier.com/locate/comnet

Control theoretic optimization of 802.11 WLANs: Implementation and experimental evaluation

Pablo Serrano^{a,*}, Paul Patras^b, Andrea Mannocci^{a,c}, Vincenzo Mancuso^{a,c}, Albert Banchs^{a,c}

^a University Carlos III de Madrid, Avda. Universidad, 30, 28911 Leganés (Madrid), Spain

^b Hamilton Institute, National University of Ireland Maynooth, Maynooth, Co. Kildare, Ireland

^c Institute IMDEA Networks, Avenida del Mar Mediterraneo, 22, 28918 Leganés (Madrid), Spain

ARTICLE INFO

Article history:

Received 14 November 2011

Received in revised form 5 June 2012

Accepted 1 September 2012

Available online xxxx

Keywords:

WLAN

802.11

Control theory

Experimentation

ABSTRACT

In 802.11 WLANs, adapting the contention parameters to network conditions results in substantial performance improvements. Even though the ability to change these parameters has been available in standard devices for years, so far no adaptive mechanism using this functionality has been validated in a realistic deployment. In this paper we report our experiences with implementing and evaluating two adaptive algorithms based on control theory, one centralized and one distributed, in a large-scale testbed consisting of 18 commercial off-the-shelf devices. We conduct extensive measurements, considering different network conditions in terms of number of active nodes, link qualities, and data traffic. We show that both algorithms significantly outperform the standard configuration in terms of total throughput. We also identify the limitations inherent in distributed schemes, and demonstrate that the centralized approach substantially improves performance under a large variety of scenarios, which confirms its suitability for real deployments.

© 2012 Elsevier B.V. All rights reserved.

1. Introduction

The IEEE 802.11 standard for Wireless LANs [1] has become one of the most commonly used technologies to provide broadband connectivity to the Internet. The default channel access mechanism employed in IEEE 802.11 networks is based on a CSMA/CA scheme, regulated by a set of parameters that determines the aggressiveness of the stations when trying to access the channel. In particular, the contention window (CW) parameter controls the probability that a station defers or transmits a frame once the medium has become idle, and therefore has a key impact on the WLAN performance.

Commercial devices implement a fixed CW configuration, which is known to yield suboptimal performance [2]. Indeed, for a fixed CW, if too many stations contend

the collision rate will be very high, while if few stations are backlogged the channel will be underutilized most of the time. This behavior has been analyzed by several works in the literature, e.g. [2–4], which have shown that adapting the CW to the number of backlogged stations significantly improves performance.

Following the above result, an overwhelming number of solutions have proposed to adapt the 802.11 MAC behavior to the observed network conditions with the goal of maximizing the WLAN performance [5–15]. However, as we detail in the related work section, these previous works suffer from both theoretical and practical limitations.

In this paper, we present our experiences with the implementation of two adaptive algorithms, namely the *Centralized Adaptive Control* (CAC) [16] and the *Distributed Adaptive Control* (DAC) [17], both based on a Proportional Integrator (PI) controller that dynamically tunes the CW configuration to optimize performance. In contrast to previous proposals [5–13,18,14,15], both algorithms are supported by solid theoretical foundations from control theory and can be easily implemented with unmodified

* Corresponding author. Address: Universidad Carlos III de Madrid, Avenida de la Universidad 30, E-28911 Leganés, Madrid, Spain. Tel.: +34 91 624 6236; fax: +34 91 624 8749.

E-mail address: pablo@it.uc3m.es (P. Serrano).

existing devices. Further, as compared to the few existing implementations of adaptive MAC mechanisms that require complete interface redesign and enhanced computational capabilities [15] or rely on proprietary firmware code and are tight to a specific platform [18], our prototypes can be deployed with off-the-shelf devices, demand minimal system resources and do not alter the networking stack, as they are executed as userland processes.

CAC and DAC have been carefully designed following detailed performance analyses of 802.11 and employing control-theoretic techniques, to ensure both optimal performance and system stability. While previous simulation results illustrate the benefits of these algorithms [16,17], in this paper we demonstrate the feasibility of running these mechanisms on real devices, and extensively assess their behavior in a real deployment under a wide range of network conditions.

First, we provide a detailed description of the implementation of our adaptive mechanisms with commodity hardware and open-source drivers. The algorithms run as user space applications and rely on standardized system calls to estimate the contention level in the WLAN and adjust the CW configuration of 802.11 stations. We also provide insights into the differences between the theoretical design and the practical implementation of the algorithms, which arose with the inherent limitations of the real devices. Second, by conducting exhaustive experiments in a large-scale testbed consisting of 18 devices, which is able to emulate realistic deployments such as office environment, and seminar room, we evaluate the performance of our proposals under non-ideal channel effects and different traffic conditions. Additionally, we compare the performance of our algorithms against the default IEEE 802.11 configuration, and identify those scenarios where a network deployment can benefit from using such adaptive mechanisms.

Our results confirm that both approaches outperform the standard's default scheme, improving the performance by up to 50%. Our experiments also reveal that the distributed algorithm suffers from a number of problems with heterogeneous radio links, which are inherent in its distributed nature and the limitations of the wireless interfaces. In contrast, the centralized scheme exhibits remarkable performance under a wide variety of network conditions. The conclusions drawn from our analysis prove the feasibility of using adaptive MAC mechanisms in realistic scenarios and provide valuable insights for their design.

The remainder of the paper is organized as follows. Section 2 summarizes the IEEE 802.11 EDCA protocol and the underlying principles of CAC and DAC. Section 3 details the implementation of the functionality comprised by the proposed schemes. Section 4 describes our testbed and Section 5 validates the implementation of the algorithms. Section 6 presents a thorough experimental study of the algorithms in a wide set of network conditions. Finally, Section 7 summarizes the related work and Section 8 concludes the paper.

2. Background

This section summarizes the behavior of IEEE 802.11 EDCA [1] and the two adaptive protocols implemented in this paper.

2.1. IEEE 802.11 EDCA

The IEEE 802.11 Enhanced Distributed Channel Access (EDCA) mechanism [1] is a CSMA/CA-based protocol that operates as follows. If a station with a new frame to transmit senses the channel idle for a period of time equal to the arbitration interframe space parameter (*AIFS*), the station transmits. Otherwise, if the channel is busy (either immediately or during the *AIFS* period), the station continues to monitor the channel until it is sensed idle for an *AIFS* interval, and then executes a backoff process.

Upon starting the backoff process, stations compute a random integer uniformly distributed in the range $[0, CW - 1]$, and initialize their backoff time counter with this value. The *CW* value is called the contention window, and depends on the number of failed transmission attempts. For the first transmission attempt the minimum contention window (CW_{min}) is used. In case of a collision, its value doubles, up to a maximum value CW_{max} . The backoff time counter is decremented once every time slot if the channel is sensed idle, frozen when a transmission is detected on the channel, and reactivated when the channel is sensed idle again for an *AIFS* time. When the backoff time counter reaches zero, the station transmits its frame in the next time slot.

When two or more stations start transmitting simultaneously, a collision occurs. Acknowledgment (ACK) frames are used to notify a transmitting station of successfully received frames. In the case of a failed transmission, the station doubles its *CW* and reenters the backoff process. Once a frame has been successfully transmitted or the retry limit has been exceeded, the *CW* value is set again to CW_{min} . To prevent duplicates, the standard uses a retry flag *R* to mark those frames that are being retransmitted, i.e., the flag is set to 0 on the first transmission attempt, and set to 1 on every retransmission (see Fig. 1). As we discuss later, our algorithms exploit this functionality to infer the network conditions and adapt the *CW* of the stations accordingly.

To support service differentiation, EDCA implements different access categories (ACs) at every station, each having a different backoff configuration. The parameters of each AC are announced by the Access Point using the Beacon Frames. In the rest of the paper we do not consider service differentiation and assume that all stations only execute the Best Effort AC.

2.2. Optimal point of operation of the WLAN

Both CAC and DAC share the goal of adjusting the *CW* to drive the WLAN to the optimal point of operation that maximizes the total throughput given the observed network conditions. Let p denote the probability that a transmission attempt collides. Following the work of [2], that derives the optimal transmission probability τ_{opt} , the corresponding optimal conditional collision probability can be computed as

$$p_{col} = 1 - (1 - \tau_{opt})^{n-1} = 1 - \left(1 - \frac{1}{n} \sqrt{\frac{2T_e}{T_c}}\right)^{n-1}, \quad (1)$$




Fig. 1. Retry flag marking upon collisions.




Fig. 2. CAC algorithm.

which, following [16,17], can be approximated for large n by

$$p_{opt} \approx 1 - e^{-\sqrt{\frac{2T_e}{T_c}}}, \quad (2)$$

where T_e is the duration of an idle slot (a PHY layer constant) and T_c is the average duration of a collision. Further, previous simulations demonstrate that this approximation does not impact the performance even when n is small [16,17]. Therefore, p_{opt} does not depend on the number of stations, but only on the average duration of a collision T_c , i.e.,

$$T_c = T_{PLCP} + \frac{L}{C} + EIFS, \quad (3)$$

where T_{PLCP} is the duration of the Physical Layer Convergence Protocol (PLCP) preamble and header, C is the modulation rate, $EIFS$ is a PHY layer constant and L denotes the frame length.¹

2.3. Centralized Adaptive Control algorithm

The *Centralized Adaptive Control* (CAC) algorithm [16], illustrated in Fig. 2, is based on a PI controller located at the Access Point (AP). This controller computes the configuration of the CW_{min} parameter as an integer ranging between the default minimum and maximum values defined by the standard specification, while CW_{max} is set

¹ The duration of a collision is given by the longest packet involved in a collision. For simplicity, here we consider a fixed frame length, but the analysis can be easily extended for the case of different frame sizes following [4].

as $CW_{max} = 2^m CW_{min}$, following the standard binary exponential backoff procedure.² By adjusting the CW configuration to be used by the stations, the AP controls the collision probability in the WLAN, with the goal of driving the network to the optimal point of operation that maximizes throughput.

Following the above, the controller performs two tasks every beacon interval (approx. 100 ms)³: (i) it estimates the current point of operation of the WLAN as given by the observed collision probability p_{obs} , and (ii) based on this estimation and p_{opt} , it computes the CW configuration to be used during the next beacon interval and sends it to the stations in a beacon frame.

The computation of p_{obs} is based on the observation of the retry flag of successful frames. Let us denote by R_1 (R_0) the number of observed frames with the retry bit set (unset) during a beacon interval. Assuming that no frames exceed the retry limit given by the `MAX_RETRY` parameter,⁴ and that transmission attempts collide with a constant and independent probability,⁵ the observed proba-

² In our experiments we use the PHY layer parameters of IEEE 802.11a, hence $m = 6$ is chosen.

³ This is the default beacon interval value recommended by the 802.11 specification and standard compliant devices must be able to update their configuration with this frequency. Existing devices support setting this parameter to different values, however larger values may slow the reaction to network changes.

⁴ Note that this assumption is accurate as in an optimally configured WLAN the collision probability is very low.

⁵ This assumption has been widely used and shown to be accurate, see e.g. [2].

bility of a collision in the WLAN can be estimated with (see [16]):

$$p_{obs} = \frac{R_1}{R_0 + R_1}, \quad (4)$$

which can also be seen as the probability that the first transmission attempt from a station collides. Given the short length of a frame transmission relative to the beacon interval, on average there will be a large number of samples to accurately compute p_{obs} . Still, in order to prevent inaccuracies, if the number of available samples of R_0 and R_1 is smaller than 20, the update is deferred until the next beacon, which results in a good accuracy for our purposes as experiments demonstrate.

To calculate the new CW_{min} , CAC employs a PI controller that takes as input an error signal e , computed as the difference between the observed collision probability p_{obs} and the target value p_{opt} :

$$e = p_{obs} - p_{opt}. \quad (5)$$

In this way, when the observed collision probability is above the target value, the error signal will be positive and trigger an increase of the CW_{min} , and consequently a decrease of the collision rate in the next beacon interval. Similarly, when the collision probability is below the target value, CW_{min} is decreased in order to increase the activity on the channel. The operation of CAC is summarized in Algorithm 1.

Algorithm 1. Centralized Adaptive Control algorithm.

```

1: while true do
2:   repeat
3:     if new frame sniffed then
4:       retrieve retry flag
5:       if retry flag is set then
6:         Increment  $R_1$ 
7:       else
8:         Increment  $R_0$ 
9:       end if
10:    end if
11:   until new beacon interval
12:   compute  $p_{obs}[t]$  using (4)
13:    $e[t] = p_{obs}[t] - p_{opt}$ 
14:    $CW_{min}[t] = CW_{min}[t - 1] + K_p \cdot e[t]$ 
15:      $+ (K_I - K_p) \cdot e[t - 1]$ 
16:   send beacon with new  $CW$  configuration
17: end while

```

The $\{K_p, K_I\}$ parameters of the PI controller are obtained using the Ziegler–Nichols rules, to achieve a proper trade-off between stability and speed of reaction to changes. Specifically, these are computed as $K_p = 0.4K_u$, respectively $K_I = K_p / (0.85T_i)$, where K_u is the K_p value that turns the system unstable when $K_I = 0$ and T_i is the oscillation period under these conditions [19]. Then, K_p and K_I are configured as follows:

$$K_p = \frac{0.8}{p_{opt}^2 \left(1 + p_{opt} \sum_{k=0}^{m-1} (2p_{opt})^k\right)}; \quad (6)$$

$$K_I = \frac{0.4}{0.85 \cdot p_{opt}^2 \left(1 + p_{opt} \sum_{k=0}^{m-1} (2p_{opt})^k\right)}.$$

The detailed computation of the $\{K_p, K_I\}$ parameters and the proof that the system behaves stably with this configuration are given in [16].

2.4. Distributed Adaptive Control algorithm

The *Distributed Adaptive Control* (DAC) algorithm [17] employs an independent PI controller at each station to compute its CW configuration, to drive the overall collision probability to the target value p_{opt} . As illustrated in Fig. 3, each controller computes the CW_{min} value employed by its Network Interface Card (NIC), based on the locally observed network conditions. Similarly to CAC, CW_{max} is set as $CW_{max} = 2^m CW_{min}$.

While with centralized approaches all stations use the same configuration provided by a single entity, and therefore fairly share the channel, with distributed approaches this is not necessarily the case. Indeed, given that the collision probability in a network of n stations is $p = 1 - \prod_{i=1}^n (1 - \tau_i)$, in general there are multiple (τ_1, \dots, τ_n) solutions that satisfy $p = p_{opt}$, with $\tau_i \neq \tau_j$. These solutions maximize the total network throughput but result in dissimilar allocations among stations. Therefore, to guarantee a fair throughput distribution, the error signal utilized in DAC consists of two terms: one to drive the WLAN to the desired point of operation, and another one to achieve fairness between stations. More specifically, the error signal at station i is given by

$$e_i = e_{collision,i} + e_{fairness,i}. \quad (7)$$

The first term of (7) ensures that the collision probability in the network is driven to the target value:

$$e_{collision,i} = p_{obs,i} - p_{opt}, \quad (8)$$

where $p_{obs,i}$ denotes the collision probability as measured by station i . When the collision probability observed by station i is larger than the target value, the above term yields a positive error that increases the CW of station i , thereby reducing the collision probability.

The second term of (7) is computed as

$$e_{fairness,i} = p_{obs,i} - p_{own,i}, \quad (9)$$

where $p_{own,i}$ is the collision probability experienced by station i . The purpose of this second component of e_i is to drive the CW of all stations to the same value. Indeed, the higher the CW_{min} , the lower the number of collisions caused, and thereby, the lower the observed collision probability $p_{obs,i}$ is. Therefore, a station will increase its CW_{min} if it experiences less collisions than the others.

With the above, the error signal at the input of the controller can be expressed as:

$$e_i = 2 \cdot p_{obs,i} - p_{own,i} - p_{opt}. \quad (10)$$

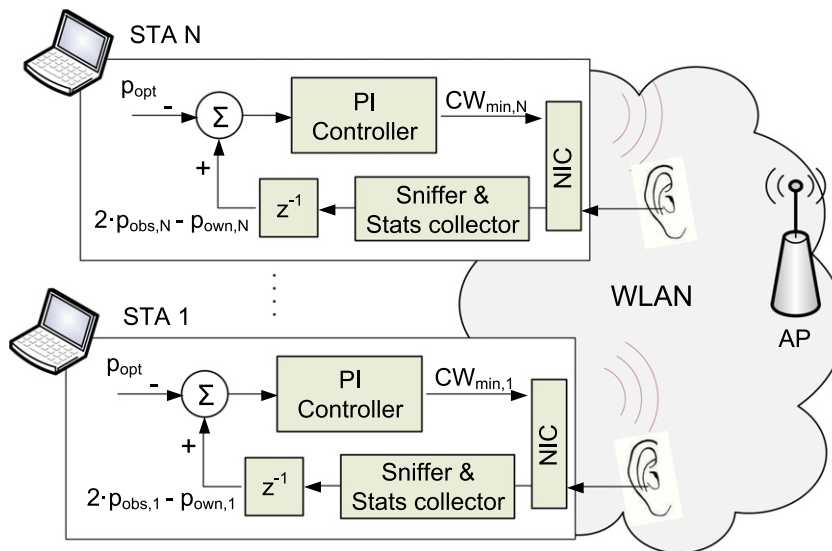


Fig. 3. DAC algorithm.

Algorithm 2. Distributed Adaptive Control algorithm.

```

1: while true do
2:   repeat
3:     if new frame sniffed then
4:       retrieve retry flag and
5:       increment  $R_0$  or  $R_1$  accordingly
6:     end if
7:   until beacon received
8:   compute  $p_{obs,i}$  using (4)
9:   fetch  $T$  and  $F$  from driver stats
10:  compute  $p_{own,i}$  using (11)
11:   $e[t] = 2 \cdot p_{obs,i}[t] - p_{own,i}[t] - p_{opt}$ 
12:   $CW_{min}[t] = CW_{min}[t-1] + K_p \cdot e[t]$ 
13:   $+ (K_I - K_p) \cdot e[t-1]$ 
14:  update the local CW configuration
15: end while

```

To compute the error signal, each station needs to measure $p_{obs,i}$ and $p_{own,i}$. The former is computed as p_{obs} in CAC. For the computation of $p_{own,i}$, we rely on the following statistics which are readily available from wireless cards: the number of successful transmission attempts T and the number of failed attempts F . With these statistics, $p_{own,i}$ is computed as:

$$p_{own,i} = \frac{F}{F+T}. \quad (11)$$

Each station will estimate $p_{obs,i}$ and $p_{own,i}$ and compute the error signal e_i , which is provided to the PI controller for the computation of the new $CW_{min,i}$. Like in CAC, we choose to trigger an update of the $CW_{min,i}$ every beacon interval, as this is compatible with existing 802.11 hardware, which is able to update the EDCA configuration at the beacon frequency.

Although the analysis of DAC, based on multivariable control theory, significantly differs from the analysis of CAC, based on standard control theory, the $\{K_p, K_I\}$ parameters that each station uses are the same ones of (6), as proved in [17]. The DAC operation is summarized in Algorithm 2.

Finally, it is important to notice that both approaches do not require to estimate the number of active stations in the WLAN and do not introduce any signaling among nodes, thus scale well with the network size. However, DAC fits more naturally the ad hoc mode of operation where there is no entity maintaining node synchronization and propagating the MAC parameters through beacons. On the other hand, we will show that by having a global view of the WLAN conditions, CAC can ensure better fairness among stations. Therefore, there exists tradeoff between throughput fairness and operation paradigm when choosing one of the two approaches in a practical deployment.

3. Implementation details

A major advantage of CAC and DAC is that they are based on functionalities already available in IEEE 802.11 devices, and therefore can be implemented with commercial off-the-shelf (COTS) hardware. We address the WLAN operation with nodes employing a single access category, namely best-effort, as this is the usual case in existing WLANs.⁶ In what follows we describe the hardware used in our deployment and the implementation of the functionality required by CAC and DAC.

⁶ Applications rarely request service differentiation using e.g. the "Type of Service" field (ToS) within the IP header. Therefore the MAC layer is unable to classify the received packets, which are assigned to a single queue. Still, it should be noted that our scheme could be extended to the case of multiple access categories, as we did in [20] for the case of video transmission in the presence of data traffic.

3.1. Implementation overview

We have implemented our algorithms using Soekris net4826-48 devices [21]. These are low-power, low-costs computers equipped with 233 MHz AMD Geode SC1100 CPUs, 2 Mini-PCI sockets, 128 Mbyte SDRAM and 256 Mbyte compact flash circuits for data storage. To accommodate the installation of current Linux distributions, we have extended the storage capacity of the boards with 2-GB USB drives. As wireless interfaces, we used Atheros AR5414-based 802.11a/b/g devices.

As software platform we installed Gentoo Linux OS (kernel 2.6.24) and the popular MadWifi open-source WLAN driver (version v0.9.4) [22], which we modified as follows: (i) we enabled the dynamic setting of the EDCA parameters for the best effort AC, since the driver disables by default changing the configuration of the parameters for this category (which is not inline with the standard specification); (ii) we overwrote the drivers' EDCA values for the best-effort traffic with the standard recommended ones [1], and (iii) for the case of DAC we modified the driver to enable the stations to employ the locally computed EDCA configuration using standardized system calls (as described in Section 3.4). The source code of the modified drivers and our implemented prototypes is available online [23].

Fig. 4 illustrates the main modules of our implementation of CAC and DAC. The algorithms do not require introducing modifications to the hardware/firmware nor have tight timing constraints, and therefore they can run as user-space applications that communicate with the driver by means of `IOCTL` calls. We also take advantage of the ability of the MadWifi driver to support multiple virtual devices using different operation modes (master/managed/monitor) with a single physical interface. In the following we detail the implementation of the different modules.

3.2. Estimation of p_{obs}

Both algorithms require to estimate the collision probability observed in the WLAN. For the case of CAC this is performed only at the AP and results in p_{obs} , while for the case of DAC this is performed independently at each station i and results in $p_{obs,i}$. The estimators are computed with (4), which relies on observing the retry flag of the overheard frames. We next explain how these values are obtained from a practical perspective.

To overhear frames, we utilize a virtual device operating in the so called `monitor` mode with promiscuous configuration. With this configuration, the device passes all traffic to user-space applications, including frames not addressed to the station. We also configure the device to pass the received frames with full IEEE 802.11 link layer headers, such that the Frame Control field of the frames (where the retry flag resides) can be examined.

With this set-up, the algorithms open a `raw` socket to the driver, which enables the reception of Layer 2 frames. Through this socket, the algorithms listen for transmitted frames and process their headers in an independent thread (the "Frame Sniffer" module of Fig. 4). For every observed frame, one of the counters used in the estimation of the collision probability is incremented: R_0 if the retry flag was unset, R_1 if the retry flag was set. Every beacon interval the computation of p_{obs} or $p_{obs,i}$ using (4) is triggered, and then the counters are reset to zero.

3.3. Estimation of p_{own}

In addition to the observed collision probability $p_{obs,i}$, the DAC algorithm requires to estimate the experienced collision probability $p_{own,i}$. We perform this computation in the "Statistics Collector" module of Fig. 4 using information recorded by the wireless driver. More specifically, at

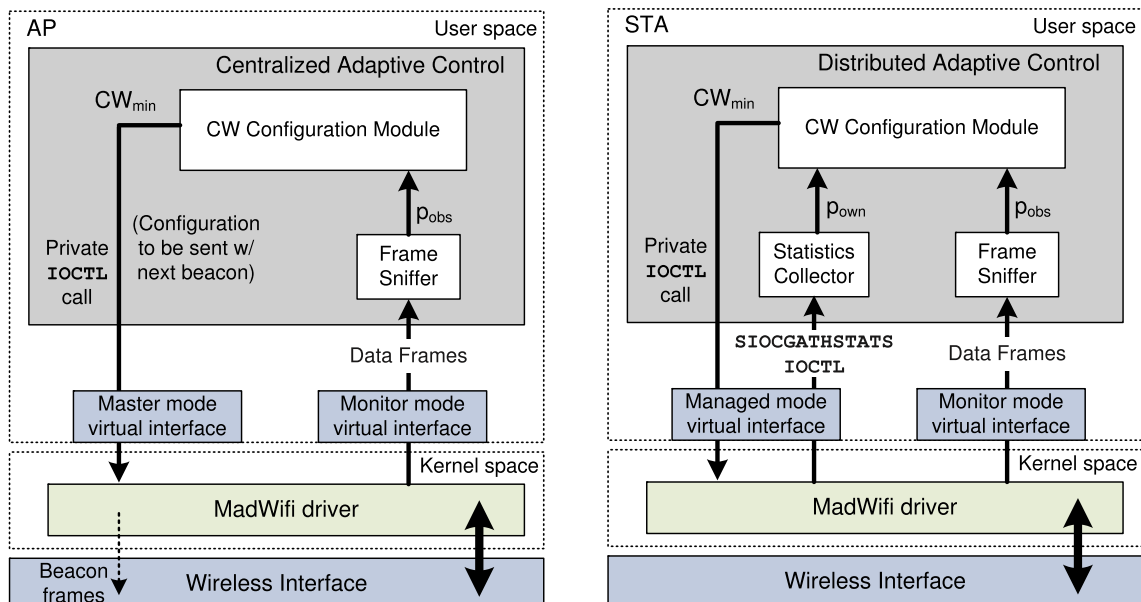


Fig. 4. CAC (left) and DAC (right) implementations.

the end of a beacon interval we open a communication channel with the driver instance, configured in `managed` mode, and perform a `SIOCGATHSTATS` `IOCTL` request. Upon this request, the driver populates an `ath_stats` data structure, which contains detailed information about the transmitted and received frames since the Linux kernel has loaded the driver module. Out of the statistics retrieved, the records that are of particular interest for our implementation are:

- `ast_tx_packets`: number of unique frames sent to the transmission interface.
- `ast_tx_noack`: number of transmitted frames that do not require ACK.
- `ast_tx_longretry`: number of transmission retries of frames larger than the RTS threshold. As we do not use the RTS/CTS mechanisms, this is the total number of retransmissions.
- `ast_tx_xretries`: number of frames not transmitted due to exceeding the retry limit, which is set by the `MAX_RETRY` parameter.

To compute $p_{own,i}$ we need to count the number of successful transmissions and the number of failed attempts. To compute the former, we subtract from the number of unique frames those that are not acknowledged (e.g., management frames) and those that were not delivered,

$$\begin{aligned} \text{Successes} &= \text{ast_tx_packets} - \text{ast_tx_xretries} \\ &\quad - \text{ast_tx_noack}. \end{aligned} \quad (12)$$

Similarly, to compute the number of failed attempts, out of the total number of retransmissions we do not count those retransmissions caused by frames that were eventually discarded because the `MAX_RETRY` limit was reached, therefore,

$$\begin{aligned} \text{Failures} &= \text{ast_tx_longretry} \\ &\quad - \text{ast_tx_xretries} \cdot \text{MAX_RETRY}. \end{aligned} \quad (13)$$

With the above, the terms F and T of (11) used to estimate $p_{own,i}$ are computed as

$$\begin{aligned} F[t] &= \text{Failures}[t] - \text{Failures}[t-1], \\ T[t] &= \text{Successes}[t] - \text{Successes}[t-1], \end{aligned} \quad (14)$$

where t denotes the time of the current beacon interval and $t-1$ the previous one. Note that in the above we take the difference between the total number of failures and successes, respectively, counted by the driver at the end of the current and previous intervals, as these counters accumulate values since the driver module has been loaded in the kernel and cannot be reset by the user.

3.4. Contention window update

With the estimated collision probabilities, CAC and DAC compute the error signal at the end of a beacon interval according to (5) and (7), respectively. Depending on this value, the PI controller triggers an update of the CW_{min} to be used in the next beacon interval t , according to the following expression:

$$CW_{min}[t] = CW_{min}[t-1] + K_P \cdot e[t] + (K_I - K_P) \cdot e[t-1]. \quad (15)$$

To ensure a safeguard against too large and too small CW_{min} values we impose lower and upper bounds for the CW_{min} . We set these bounds to the default CW_{min}^{DCF} and CW_{max}^{DCF} values specified by the standard, which are 16 and 1024, respectively, for IEEE 802.11a [24].

The algorithms assume that the CW_{min} can take any integer value in the [16, 1024] range. However, with our devices only integer powers of 2 are supported (i.e., $CW_{min} \in \{16, 32, \dots, 1024\}$). Therefore, the value actually used is obtained as:

$$CW[t] = \text{pow}(2, \text{rint}(\log_2(CW_{min}[t]))). \quad (16)$$

where `rint(x)` is a function that returns the integer value nearest to x .

To commit the computed CW configuration, first we retrieve the list of private `IOCTLs` supported by the device to search for the call that sets the CW_{min} . Once this call has been identified, we prepare an `iwreq` data structure with the following information: the interface name, the base-2 exponent of the CW computed with Eq. (16), the access category index as defined by the standard (0 for Best Effort) and an additional parameter that identifies if the value is intended to be used locally or propagated. For the case of DAC this value is set to 0, as the CW is only intended to the local card, while for the case of CAC is set to 1, thereby requesting the driver to broadcast the new CW within the EDCA Parameter Set element of the next scheduled beacon frame.

4. Testbed description

In this section we describe our testbed and analyze the link qualities between each node and the AP, showing that our set-up is able to mimic a realistic deployment with significant differences in terms of SNR.

Our testbed is located in the Torres Quevedo building at University Carlos III de Madrid. It consists of 18 devices deployed under the raised floor, a placement that provides physical protection as well as radio shielding to some extent (see [25]).

Fig. 5 illustrates the location of the nodes. We placed one node towards the center of the testbed, which acts as the Access Point throughout our experiments and we denote as AP, while stations are distributed at different distances from the AP and numbered from 1 to 17 in decreasing order of their links' SNR values towards the AP, which we measure as explained next. In this way, our deployment emulates several popular scenarios, e.g. cafés, offices, conference rooms, both in terms of number of nodes and mobility. Indeed, in such scenarios nodes are mostly static, experiencing dissimilar link qualities towards the AP with small variations, as our measurements confirm. Although we focus on a single AP topology, in the considered scenarios careful channel allocation is usually performed to avoid inter-AP interference, and therefore our algorithms can be independently executed within each WLAN.




Fig. 5. Deployed testbed.

All nodes are equipped with 5 dBi omnidirectional antennas and are configured to operate on channel 64 (5.32 GHz) of IEEE 802.11a standard [24], where no other WLANs were detected. All nodes use the 16-QAM modulation and coding scheme, which provides 24 Mbps channel bit rate, as calibration measurements showed that this was the highest rate achievable by the node with the worst link to the AP (node 17). Additionally, we disabled the RTS/CTS, rate adaptation, turbo, fast frame, bursting and unscheduled automatic power save delivery functionality, as well as the antenna diversity scheme for transmission/reception.

Unless otherwise specified, all nodes use the same transmission power level of 17 dBm. Given the node placement of Fig. 5, this setting results in dissimilar link qualities between each station and the AP (e.g., node 1 is very close to the AP, while e.g. node 16 is far away, which will affect the perceived SNR). To confirm this link heterogeneity, we designed the following experiment. For a given node, we ran a 10-s ping test between the station and the AP, recording the SNR values of the received frames as obtained by the `wireshark` packet analyzer [26] from the `radiotap` header.⁷ This test was performed on a

⁷ With the `radiotap` option, the driver provides additional information about received frames to user-space applications, including the signal-to-noise ratio.

node-by-node basis, and repeated for 18 h. The average and standard deviation of the SNR for each link are shown in Fig. 6.

5. Validation of the algorithms

Our first set of experiments aims at confirming that the good operation properties of CAC and DAC, obtained ana-

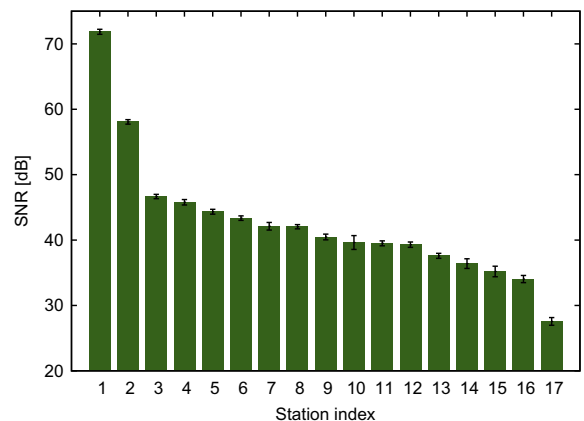


Fig. 6. SNR of the links between each node and the Access Point.

lytically and via simulations in [16,17], are also achieved in a real testbed. Specifically, we want to confirm that the use of the algorithms results in stable behavior despite the described hardware/software limitations and the impairments introduced by the channel conditions, and also assess their resource consumption in terms of CPU and memory usage.

5.1. Point of operation

We consider a scenario with $N = 10$ stations randomly selected, namely 1, 2, 5, 6, 8, 9, 10, 12, 13, and 14, which corresponds to a network with heterogeneous link qualities. Stations are constantly backlogged with 1500-Byte UDP frames, which they send to the AP utilizing `iperf` [27] for a duration of 5 min. For the case of the centralized algorithm (CAC) we log its key variables, namely, the CW announced in beacon frames and the observed collision probability p_{obs} . Both are obtained every 100 ms and depicted in Fig. 7.

As the 30-s snapshot in the figure shows, CAC drives the WLAN to the desired point of operation. Indeed, the announced CW oscillates between the two power of 2 values closest to the optimal CW_{min} (i.e., 64 and 128), while p_{obs} fluctuates stably around the desired p_{opt} given by (2). We conclude that, despite the hardware limitations imposed on the values of CW and the channel impairments, CAC is able to drive the WLAN to the desired point of operation.

Next we validate the operation of the distributed algorithm (DAC). We consider the same scenario as before, logging the key parameters of the algorithm at each station, namely $CW_{min,i}$, $p_{own,i}$ and $p_{obs,i}$. In Fig. 8 we depict in the upper subplot the evolution of the CW_{min} used by four representative nodes. More specifically, we select out of the 10 considered nodes the one with the best link (node 1) and the one with the worst link (node 14), as well as two stations with similar link qualities, (nodes 10 and 13). Additionally, in the lower subplot we show the collision probabilities estimated by node 10 ($p_{obs,10}$ and $p_{own,10}$).

From the two subplots we see that DAC also drives the average collision probability in the WLAN to the desired value. However, there is a key difference as compared to the previous case: while with CAC all stations use the same




Fig. 7. Announced CW_{min} and observed collision probability with CAC.




Fig. 8. CW_{min} used by four selected nodes (1, 10, 13 and 14) and the estimated p_{obs} and p_{own} for node 10 with DAC.

CW_{min} value, with DAC they operate at different average CW_{min} . Indeed, the four stations considered in the experiment use average CW_{min} values of 300, 92, 92 and 64, respectively. As we will explain in Section 6.2, this behavior is caused by the relative differences in link qualities, combined with the inability of the wireless interface to identify the reasons for a packet loss.

5.2. Resource consumption

In addition to analyzing the performance of CAC and DAC, it is also important to assess their resource consumption. For this purpose, we analyzed the CPU and memory usage of the algorithms utilizing the `top` Linux application, which provides a dynamic real-time view of a running system. With this tool, we recorded the used shares of the CPU time and available physical memory with a frequency of 1 sample per second and computed the average usage. CAC, which runs exclusively at the AP, demands on average 14% of the CPU time and only 0.8% of the physical memory. For the case of DAC, which runs at every station, the average CPU time consumption is 17%, while the physical memory consumption is 0.9%. Given the low speed of the nodes' CPU (233 MHz) and their reduced physical memory (128 MB), these results show that both CAC and DAC are suitable for commercial deployments as they do not pose stringent resource requirements.⁸

6. Performance evaluation

We next assess the performance of the algorithms under a large number of different scenarios and evaluate their improvements over the default EDCA configuration, which we use as a benchmark. Each considered experiment runs for 2 min and is repeated 10 times to obtain average values of the measured metrics with good statistical significance.

⁸ In order to have a benchmark to compare against, we measured the CPU and memory usage of the popular packet sniffer `tcpdump`, running on the same device type. Measurements revealed that `tcpdump` consumes 28% of the CPU time, while its memory usage reaches 4.3%, these values being significantly larger than the ones for CAC and DAC.

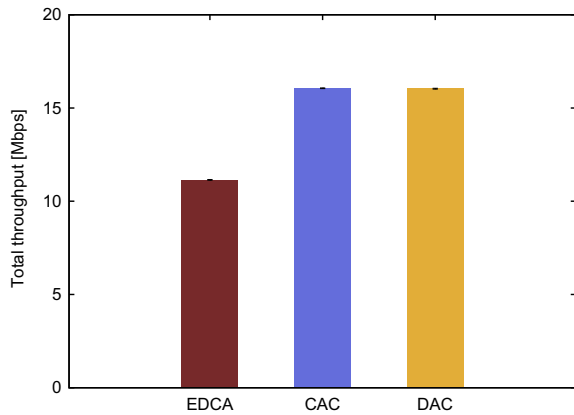


Fig. 9. Total throughput with UDP traffic.

We evaluated the behavior of the proposed algorithms with both unidirectional (UDP) and bidirectional (TCP) traffic, under diverse conditions in terms of link qualities and number of active nodes. Unless otherwise stated, we assume nodes have similar traffic demands and are constantly backlogged, thus the network is mostly saturated.

6.1. UDP throughput

We first measure the achievable throughput between the nodes and the AP when all the stations are transmitting UDP traffic at the same time. Fig. 9 plots the average and standard deviation of the total throughput obtained with each mechanism. We observe that the EDCA default configuration achieves around 11 Mbps, while the use of DAC and CAC results in a performance gain of approximately 45%. Therefore, we confirm that both approaches, by properly adapting the CW configuration to the number of contending stations, achieve a much higher efficiency.

To further examine the performance of the algorithms we plot the per-station throughput in Fig. 10. According to the figure, the use of the EDCA recommended values not only provides the lowest overall throughput figures, but also fails to provide a fair sharing of the available bandwidth. Indeed, it can be seen that, e.g., the node with the best link quality to the AP (node 3) achieves more than three times the throughput obtained by the station with the poorest link (node 15).

While DAC provides a larger total throughput than EDCA, it does not improve the level of fairness. Actually, it results in a somehow *opposite* performance as the one obtained with EDCA: stations that obtained a relatively large bandwidth with EDCA (e.g., nodes 3 and 6) now obtain a relatively small bandwidth with DAC. Furthermore, due to the poor short-term fairness properties of 802.11 MAC [28], the results show a notable variability. The use of CAC, on the other hand, provides the best performance both in terms of total throughput and fairness, as it provides all stations with very similar throughput values.

To quantify the throughput fairness achieved by the considered mechanisms we compute the Jain's fairness in-

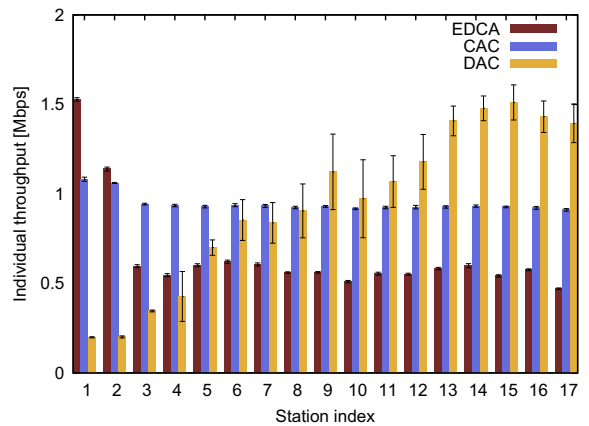


Fig. 10. Throughput per station with UDP traffic.

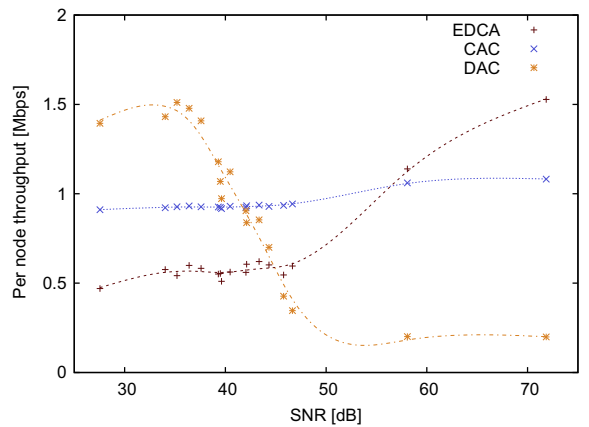


Fig. 11. Per node throughput obtained vs. SNR.

dex (JFI) [29].⁹ The resulting JFI values are 0.865, 0.997 and 0.817 for the case of EDCA, CAC and DAC, respectively. These figures confirm the good fairness properties of CAC, and shows that DAC and EDCA suffer from a higher level of unfairness, a result that we analyze next.

6.2. Impact of SNR on throughput

We have seen that link quality affects throughput distribution, in particular for EDCA and DAC. To analyze this impact, we plot in Fig. 11 the average UDP throughput per station vs. the SNR of the link between the station and the AP. Note that, in this analysis we do not modify stations' SNR by tuning their transmission power, but instead we investigate how the SNR dissimilarity that arises due to their relative position to the AP affects the throughput performance of CAC and DAC as compared to EDCA. For ease of visualization we also plot *natural smoothing splines* over the data points.

⁹ Since in most scenarios the network operates close to saturation and all nodes have the same traffic requirements, we argue that JFI is the most suitable metric for assessing the fairness properties of CAC and DAC. Specifically, our goal is to quantify the dissimilarity in the resource distribution, independent of the number of users.

From the figure we observe that: (i) for EDCA there is a noticeable and positive correlation between SNR and throughput; (ii) for CAC, performance is not much affected by SNR dissimilarities, as significantly better link qualities result in very small throughput improvements; (iii) for DAC there is a large and negative correlation between SNR and throughput, with small differences in terms of SNR causing large differences in terms of throughput.

For the case of EDCA, the positive correlation is caused by the *capture effect* [30]. With this effect, in case of a collision the receiver can decode the packet with the higher SNR. As a result, stations with better link quality obtain higher throughput. In contrast, the use of CAC reduces the number of collisions in the WLAN, and therefore the impact of the capture effect is significantly reduced.

For the case of DAC, the negative correlation is also driven by the capture effect as follows. Nodes with high capture probability will experience smaller collision rates than the others, and therefore will have $p_{own,i}$ smaller than $p_{obs,i}$. This will cause a positive error signal according to the $e_{fairness,i}$ term in (9), which will result in large CW_{min} values. Conversely, nodes with low capture probability will experience larger $p_{own,i}$ values and smaller $p_{obs,i}$ ones, and therefore will have smaller CW_{min} configurations. In this way, capturing nodes will transmit less often and therefore will obtain low throughput figures, while the other nodes will transmit more often and experience a higher throughput. Additional experiments with different transmission power settings, not reported due to space constraints, confirmed that a careful *equalization* of the link qualities is able to restore fairness to some extent.

6.3. Hidden nodes scenario

Our adaptive algorithms have been designed for scenarios where all stations are in radio range of each other and coordinate their transmissions by means of carrier sensing. However, in real deployments hidden nodes may be present, and therefore we want to investigate their behavior under such circumstances.

To this aim, we ran extensive measurements, selecting different topologies and different transmission power settings, to determine the most pathological scenario. This is obtained when node 1 acts as AP, and nodes 10 and 13 act as stations, using a transmission power level of 5 dBm. With this setting, each EDCA station transmitting in isolation (i.e., with the other station silent) obtains about 16.3 Mbps throughput, while if both stations transmit simultaneously the throughput of each one drops to 1.6 Mbps. Thereby we managed to reproduce a hidden node scenario.

We then repeated the experiment with CAC and DAC, and obtained the results depicted in Fig. 12. We observe that the use of DAC does not improve performance over EDCA. In contrast, CAC provides a dramatic throughput increase, i.e., more than three times the throughput attained with the other mechanisms. We conclude that CAC detects the large collision rate and commands hidden nodes to be less aggressive by announcing a higher CW_{min} , which lessens (but does not eliminate) the hidden node problem. On the other hand, a station running DAC is not able to over-




Fig. 12. Performance with hidden nodes.

hear MAC (re-)transmissions from hidden nodes, and hence cannot correctly estimate the collision probability in the network.

6.4. Impact network size

We next evaluate the performance of the algorithms as a function of the number of stations. To this aim, we measure the total throughput and JFI for an increasing number of contending nodes, adding new stations in ascending order of their link quality. We plot the obtained results in Fig. 13.

We observe that for both DAC and CAC the total throughput performance is practically flat, regardless of the number of stations. This result confirms that both approaches are able to adapt the CW to the number of stations present in the WLAN.

For the case of EDCA, performance degrades with the number of stations, which is the expected result from the use of a fixed set of (relatively small) contention parameters. However, for $N > 15$ the total throughput performance slightly grows again, a behavior caused by the capture effect as the last nodes to be added in our experiments are the ones experiencing better link qualities




Fig. 13. Total throughput and fairness for different number of stations.




Fig. 14. Total throughput of FTP-like traffic.

(nodes 2 and 1). This is confirmed by the fairness values, as for $N > 15$ there is a drop in the JFI for the case of EDCA. JFI values also confirm that DAC is more sensitive to heterogeneous link conditions, as its performance noticeably degrades with N . In contrast, with CAC the fairness index is practically constant for all N values.

6.5. TCP throughput

We next evaluate performance in scenarios in which stations use TCP. We start by evaluating the throughput and fairness performance when all stations are constantly backlogged sending TCP traffic to the AP, replicating *bulky* FTP transfers. Note that this scenario is substantially different from the ones considered in the previous subsections, as TCP congestion control¹⁰ introduces a “closed loop” that can lead to extreme unfairness conditions and even starvation [32]. Precisely, we expect nodes that use a more aggressive CW configuration to leave less channel access opportunities for the more conservative ones, thus causing the latter to reach congestion much faster and reduce their TCP windows more frequently. This behavior will exacerbate the difference in the throughput values they obtain and impact the fairness performance more severely.

We plot in Fig. 14 the total throughput values for the three mechanisms. According to the results, both CAC and DAC significantly outperform EDCA, improving throughput by 50% and 40%, respectively.

The per-station throughput distribution is depicted in Fig. 15. With EDCA, the node with the poorest link quality (node 17) suffers from a large performance degradation, this being worse than in the UDP case (see Fig. 9). The use of DAC with TCP traffic also exacerbates the unevenness in the traffic distribution, with node 15 clearly outstanding among the other nodes. DAC results also present a large deviation, caused by relatively frequent TCP time-outs from nodes with weak radio link (e.g., node 17). Conversely, CAC yields a remarkably fair and stable throughput distribution.

¹⁰ The Linux distribution used in our deployment executes the TCP CUBIC variant [31].

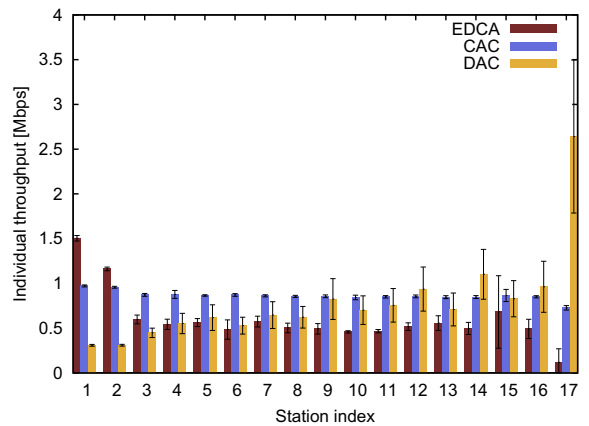


Fig. 15. Throughput per station with FTP-like traffic.

Like in the UDP case, we compute the JFI values for the resulting throughput distributions. In this case, the values for EDCA, CAC and DAC are 0.787, 0.996 and 0.692, respectively. We conclude that, as expected, the performance of EDCA and DAC worsens with TCP, while CAC preserves its good properties in this scenario.

6.6. TCP transfer delay

We finally consider a scenario involving finite-size TCP connections. More specifically, all stations alternate periods of activity—during which a transmission of 10 MB occurs—with silent periods exponentially distributed with mean λ^{-1} [33]. We consider three different values for λ , corresponding to three different levels of activity, namely high, moderate and low. For each case we ran 1-h experiments, logging all transfer durations and computing the per-station average delay. We use a *box-and-whisker* diagram to illustrate the distribution of the average delay among nodes: we provide the median, first and third quartiles of the average delay, as well as its maximum and minimum values.

Results are depicted in Fig. 16. With $\lambda^{-1} = 30$ s, which corresponds to high activity, we see that CAC provides the smallest and most uniform distribution of transfer delay among nodes, with practically no difference between the best and worst performing node. In case of EDCA, the delay shows a larger median and higher variability. However, the small distance between the first and third quartiles shows that most of the stations experience similar performance. Finally, for the case of DAC, despite the median is similar to the one of CAC, results show a much larger dispersion. This further confirms that, although DAC achieves the main objective of optimizing the total throughput, nodes adjust their CW based on different, locally observed conditions, which results in some stations being more aggressive than others, and correspondingly yields large variations in the delay figures.

When the traffic activity is moderate ($\lambda^{-1} = 60$ s), the absolute values decrease, but the relative results are similar, i.e., CAC provides again the smallest and most uniform delays among nodes. Finally, when the activity of the




Fig. 16. TCP delay performance.

nodes is low ($\lambda^{-1} = 90$ s), medians are very similar but still CAC provides the most fair distribution of the transfer delays. From these experiments, we conclude that CAC also provides the best performance under dynamic traffic scenarios.

7. Related work

The scientific literature offers many examples of MAC optimization approaches. Many of them are based on a centralized entity, responsible for monitoring system performance and adapting the system parameters to current

conditions. Other works focus on distributed approaches to adapt MAC parameters. Very little experimental work is available, and it is based on complex algorithms, non-standard functionality and small-sized networks. In the following we review the most significant contributions in each of these areas and describe the novelty of our work.

7.1. Centralized approaches

A significant number of approaches exists in the literature [5,7,14,16] that use a single node to compute the set of MAC parameters to be used in the WLAN. With the exception of our CAC algorithm [16], the main drawbacks of these approaches are that they are either based on heuristics, thereby lacking analytical support for providing performance guarantees [5,7], or they do not consider the dynamics of the WLAN under realistic scenarios [14].

7.2. Distributed approaches

Several works [8,10,13,34,35] have proposed mechanisms that independently adjust the backoff operation of each stations in the WLAN. The main disadvantages of these approaches are that they change the rules of the IEEE 802.11 standard and therefore require introducing significant hardware or firm-ware modifications.

7.3. Implementation experiences

Very few schemes to optimize WLAN performance have been developed in practice [14,15,18]. While the idea behind Idle Sense [13] is fairly simple, its implementation [18] entails a significant level of complexity, introducing tight timing constraints that require programming at the firmware level. Nonetheless, the microcode that achieves the desired functionality is proprietary and thus subject to portability constraints. Similar limitations hold for the approach of [15], which introduces changes to the MAC protocol that require redesigning the whole NIC implementation. Further, this involves complex DSP and FPGA programming and demands non-negligible computational resources. Finally, the work of [14] does not propose or evaluate any adaptive algorithm to adapt the CW but just evaluates the performance of static configurations. Additionally, our testbed is substantially larger than the one used in these previous works.

8. Conclusions

We have prototyped with standard 802.11 devices two adaptive mechanisms that tune the contention window based on the observed network conditions. In contrast to other proposals that require complex modifications, these mechanisms rely on functionalities already supported by COTS hardware/firmware, and do not introduce any extensions to the standard 802.11 MAC. We have extensively evaluated the performance of the mechanisms in an 18-nodes testbed, considering a large variety of network conditions. With our experimental study we have identified the key limitations of the distributed scheme, inherent in

realistic scenarios, and we have confirmed that the centralized mechanism significantly improves network throughput, transfer delay and fairness among stations in a broad range of circumstances, including the pathological case of hidden nodes. A major conclusion from our work is that, by simply adding a few lines of code at the AP to exploit the functionality readily available, we can achieve performance improvements of up to 50%. We believe that the results presented herein pave the way for a widespread deployment of the centralized mechanism.

Acknowledgement

This work has been supported by the European Community's Seventh Framework Programme (FP7-ICT-2009-5) under Grant Agreement No. 257263 (FLAVIA project).

References

- [1] IEEE 802.11, Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications, Revision of IEEE Std 802.11-1999, 2007.
- [2] G. Bianchi, Performance analysis of the IEEE 802.11 distributed coordination function, *IEEE Journal on Selected Areas in Communications* 18 (3) (2000) 535–547.
- [3] A. Banchs, L. Voller, Throughput analysis optimal configuration of 802.11e EDCA, *Computer Networks* 50 (11) (2006).
- [4] P. Serrano, A. Banchs, P. Patras, A. Azcorra, Optimal configuration of 802.11e EDCA for real-time data traffic, *IEEE Transactions on Vehicular Technology* 59 (5) (2010) 2511–2528.
- [5] J. Freitag, N.L.S. da Fonseca, J.F. de Rezende, Tuning of 802.11e network parameters, *IEEE Communications Letters* 10 (8) (2006) 611–613.
- [6] L. Scalia, I. Tinnirello, J.W. Tantra, Chuan Heng Foh, Dynamic MAC parameters configuration for performance optimization in 802.11e networks, in: *Proc. GLOBECOM*, San Francisco, CA, USA, Dec 2006, pp. 1–6.
- [7] A. Nafaa, A. Ksentini, A. Ahmed Mehaoua, B. Ishibashi, Y. Iraqi, R. Boutaba, Sliding contention window (SCW): towards backoff range-based service differentiation over IEEE 802.11 wireless LAN networks, *IEEE Network* 19 (4) (2005) 45–51.
- [8] Y. Yang, J.J. Wang, R. Kravets, Distributed optimal contention window control for elastic traffic in single-cell wireless LANs, *IEEE Transactions on Networking* 15 (6) (2007) 1373–1386.
- [9] Q. Xia, M. Hamdi, Contention window adjustment for IEEE 802.11 WLANs: a control-theoretic approach, in: *Proc. International Conference on Computer Communications (ICC)*, Istanbul, Turkey, June 2006.
- [10] Q. Ni, I. Aad, C. Barakat, T. Turletti, Modeling and analysis of slow CW decrease for IEEE 802.11 WLAN, in: *Proc. IEEE Personal, Indoor, and Mobile Radio Communications Conference (PIMRC)*, Beijing, 2003.
- [11] Lijun Chen, Steven H. Low, John C. Doyle, Joint congestion control and media access control design for wireless ad hoc networks, in: *Proc. IEEE INFOCOM*, Miami, Florida, March 2005.
- [12] Lijun Chen, Steven H. Low, John C. Doyle, Random access game and medium access control design, *IEEE Transactions on Networking* 18 (4) (2010) 1063–6692.
- [13] Martin Heusse, Franck Rousseau, Romaric Guillier, Andrzej Duda, Idle Sense: an optimal access method for high throughput and fairness in rate diverse wireless LANs, in: *Proc. ACM SIGCOMM*, Philadelphia, PA, USA, 2005, pp. 121–132.
- [14] Vasilios A. Siris, George Stamatakis, Optimal CWmin selection for achieving proportional fairness in multi-rate 802.11e WLANs: test-bed implementation and evaluation, in: *Proceedings of the 1st International Workshop on Wireless Network Testbeds, Experimental Evaluation & Characterization*, Los Angeles, CA, USA, 2006, pp. 41–48.
- [15] Ralph Bernasconi, Silvia Giordano, Alessandro Puiatti, Raffaele Bruno, Enrico Gregori, Design and implementation of an enhanced 802.11 MAC architecture for single-hop wireless networks, *EURASIP Journal on Wireless Communications and Networking* 2007 (1) (2007) 33.
- [16] P. Patras, A. Banchs, P. Serrano, A control theoretic approach for throughput optimization in IEEE 802.11e EDCA WLANs, *Mobile Networks and Applications (MONET)* 14 (6) (2009) 697–708.
- [17] P. Patras, A. Banchs, P. Serrano, A. Azcorra, A control theoretic approach to distributed optimal configuration of 802.11 WLANs, *IEEE Transactions on Mobile Computing* 10 (6) (2011) 897–910.
- [18] Yan Grunenberger, Martin Heusse, Franck Rousseau, Andrzej Duda, Experience with an implementation of the idle sense wireless access method, in: *Proceedings of the ACM CoNEXT conference*, New York, New York, 2007, pp. 1–12.
- [19] G.F. Franklin, J.D. Powell, M.L. Workman, *Digital Control of Dynamic Systems*, second ed., Addison-Wesley, 1990.
- [20] Paul Patras, Albert Banchs, Pablo Serrano, A Control Theoretic Scheme for Efficient Video Transmission over IEEE 802.11e EDCA WLANs, *ACM Transactions on Multimedia Computing, Communications, and Applications (TOMCCAP)* 8 (3) (2012).
- [21] Soekris Engineering, Inc. <<http://www.soekris.com/>>.
- [22] The MadWifi Project. <<http://madwifi-project.org/>>.
- [23] Personal Page, <<http://www.hamilton.ie/ppatras/#code>>.
- [24] IEEE 802.11, Supplement to Wireless LAN Medium Access Control and Physical Layer Specifications: High-speed Physical Layer in the 5 GHz Band, *IEEE Std 802.11a*, 1999.
- [25] Pablo Serrano, Carlos J. Bernardos, Antonio de la Oliva, Albert Banchs, Ignacio Soto, Michael Zink, FloorNet: Deployment and Evaluation of a Multihop Wireless 802.11 Testbed, *EURASIP Journal on Wireless Communications and Networking*, 2010.
- [26] Wireshark Network Protocol Analyzer. <<http://wireshark.org>>.
- [27] Iperf Traffic Generator. <<http://sourceforge.net/projects/iperf/>>.
- [28] Gilles Berger-Sabbatel, Andrzej Duda, Martin Heusse, Franck Rousseau, "Short-term fairness of 802.11 networks with several hosts, in: *Proceedings of the Sixth IFIP TC6/WG6.8 Conference on Mobile and Wireless Communication Networks (MWCN 2004)*, vol. 162, Paris, France, 2004, pp. 263–274(25–27).
- [29] R. Jain, Chiu, D.M., W. Hawe, A Quantitative Measure of Fairness and Discrimination for Resource Allocation in Shared Systems, *DEC Research Report TR-301*, 1984.
- [30] Z. Hadzi-Velkov, B. Spasenovski, Capture effect in IEEE 802.11 basic service area under influence of Rayleigh fading and near/far effect, in: *Proc. IEEE Personal, Indoor, and Mobile Radio Communications Conference (PIMRC)*, 2002, vol. 1, pp. 172–176.
- [31] Sangtae Ha, Injong Rhee, Lisong Xu, CUBIC: a new TCP-friendly high-speed TCP variant, *ACM SIGOPS Operating Systems Review* 42 (2008) 64–74.
- [32] O. Gurewitz, V. Mancuso, Jingpu Shi, E.W. Knightly, Measurement and modeling of the origins of starvation of congestion-controlled flows in wireless mesh networks, *IEEE Transactions on Networking* 17 (6) (2009).
- [33] P. Barford, M.E. Crovella, Generating representative web workloads for network and server performance evaluation, in: *Proc. ACM SIGMETRICS*, 1998, pp. 151–169.
- [34] L. Bononi, M. Conti, E. Gregori, Runtime optimization of IEEE 802.11 wireless LANs performance, *IEEE Transactions on Parallel and Distributed Systems* 15 (1) (2004) 66–80.
- [35] F. Cali, M. Conti, E. Gregori, IEEE 802.11 Protocol: design performance evaluation of an adaptive backoff mechanism, *IEEE Journal on Selected Areas in Communications* 18 (9) (2000) 1774–1786.



Pablo Serrano received his Telecommunication Engineering degree and his PhD from the University Carlos III of Madrid (UC3M) in 2002 and 2006, respectively. He has been with the Telematics Department of UC3M since 2002, where he currently holds the position of Assistant Professor. In 2007 he was a Visiting Researcher at the Computer Network Research Group at Univ. of Massachusetts Amherst partially supported by the Spanish Ministry of Education under a José Castillejo grant. His current work focuses on performance evaluation of wireless networks. He has over 30 scientific papers in peer-reviewed international journal and conferences. He also serves as TPC member of several international conferences, including IEEE Globecom and IEEE INFOCOM.



Paul Patras received the M.Sc. and Ph.D. degrees in telematics engineering from University Carlos III of Madrid in 2008 and 2009, respectively. He is currently a Post-Doctoral research fellow at the Hamilton Institute. In 2010 he was a visiting intern at the Rice Networks Group of Rice University, Houston, USA. He serves as a reviewer for several international journals and conferences, including Elsevier Computer Networks, IEEE/ACM Transactions on Networking, IEEE Transactions on Communications, IEEE INFO-

COM, IEEE Globecom and IEEE ICC. His current research interests include performance optimization of IEEE 802.11 wireless LANs, adaptive MAC mechanisms, dynamic spectrum access, prototype implementation and testbeds.



Andrea Mannocci received a M.Sc. degree in Computer Science Engineering from University of Pisa in 2010. His Masters thesis, entitled “Stepwise Evolution of Java Applications using Dynamic Updates”, was developed during an Erasmus exchange period at The Maersk Mc-Kinney Moller Institute, Syddansk Universitet (University of Southern Denmark, Odense, Denmark) under the supervision of Prof. Bo Nørregaard Jørgensen and Dr. Allan Raundahl Gregersen. In 2011 he also received a M.Sc. degree in Telematics Engineering from

University Carlos III of Madrid. He is a research assistant at Institute IMDEA Networks. His current research activities mainly focuses on the IEEE 802.11 and 802.16 protocols.



Vincenzo Mancuso received his master degree in Electronics from the University of Palermo, Italy, in 2001, and a Ph.D. in Electronics, Computer Science and Telecommunications from the same University in 2005. In 2004–2005 he was with University of Roma “Tor Vergata”, working on SATNEX, an IST Network of Excellence. He was then back to University of Palermo as postdoc, from June 2005 to May 2009. Also, during that period he has been visiting scholar at the ECE department of Rice University, Houston, Texas

(September 2006 to January 2008). In June 2009 he joined the MAESTRO team at INRIA Sophia Antipolis, France. Since September 2010, Vincenzo is with Institute IMDEA Networks in Madrid, Spain.



Albert Banchs received his Telecommunications Engineering degree from the Polytechnical University of Catalonia in 1997, and the Ph.D. degree from the same university in 2002. His Ph.D. received the national award for best thesis on broadband networks. He was a visitor researcher at ICSI, Berkeley, in 1997, worked for Telefonica I+D, in 1998, and for NEC Europe Ltd., Germany, from 1998 to 2003. Since 2003, he is with the University Carlos III of Madrid and since 2009 he is also Deputy Director of IMDEA Networks. His

authors over 50 publications in peer-reviewed journals and conferences and five patents. He is associated editor for IEEE Communications Letters and has been guest editor for IEEE Wireless Communications and for Computer Networks. He has served on the TPC of a number of conferences and workshops including IEEE INFOCOM, IEEE ICC and IEEE Globecom, and was TPC chair for European Wireless 2010.