

# Providing Throughput and Fairness Guarantees in Virtualized WLANs through Control Theory

Albert Banchs · Pablo Serrano · Paul Patras · Marek Natkaniec

**Abstract** With the increasing demand for mobile Internet access, WLAN virtualization is becoming a promising solution for sharing wireless infrastructure among multiple service providers. Unfortunately, few mechanisms have been devised to tackle this problem and the existing approaches fail in optimizing the limited bandwidth and providing virtual networks with fairness guarantees. In this paper, we propose a novel algorithm based on control theory to configure the virtual WLANs with the goal of ensuring fairness in the resource distribution, while maximizing the total throughput. Our algorithm works by adapting the contention window configuration of each virtual WLAN to the channel activity in order to ensure optimal operation. We conduct a control-theoretic analysis of our system to appropriately design the parameters of the controller and prove system stability, and undertake an extensive simulation study to show that our proposal optimizes performance under different types of traffic. The results show that the mechanism provides a fair resource distribution independent of the number of stations and their level of activity, and is able to react promptly to changes in the network conditions while ensuring stable operation.

---

A. Banchs (✉)  
Institute IMDEA Networks  
Avenida del Mar Mediterraneo, 22  
28918 Leganés (Madrid), Spain  
Phone: +34-91-6246245  
Fax: +34-91-6248749  
E-mail: albert.banchs@imdea.org

P. Serrano · A. Banchs  
Universidad Carlos III de Madrid  
Leganés, Spain

P. Patras  
Hamilton Institute of the National University of Ireland  
Maynooth, Ireland

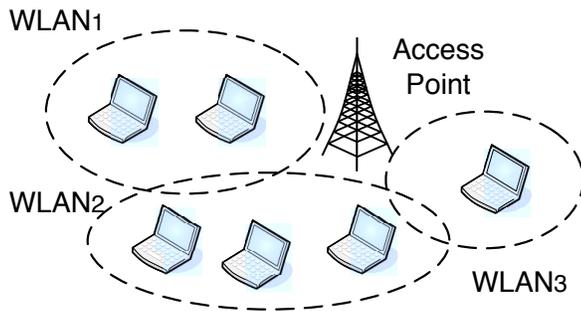
M. Natkaniec  
AGH University of Science and Technology  
Krakow, Poland

**Keywords** Wireless LAN · virtualization · 802.11 · control theory · throughput optimization · fairness

## 1 Introduction

As portable devices are becoming widespread and users increasingly prefer connecting to the Internet through wireless access points (APs), Internet Service Providers (ISPs) are competing to provide wireless broadband services in popular locations such as airports, cafés, hotels, *etc.* As the infrastructure on such premises is usually managed by local businesses, network operators seeking to enable roaming services for their existing customers or to gain additional revenue from temporary users are often required to share the resources of a single AP with other parties. The solutions range from setting up a unique client authentication mechanism on the AP [1], which enables virtual networking for each provider across the AP's gateway connection, to establishing virtual APs (VAP) on the same device, that will manage the operation of independent virtual WLANs, exposing a unique service set identifier (SSID) for the users of each operator. The latter is enabled by the recent hardware/software advances that allow the virtualization of a single physical interface and the creation of multiple logical AP entities [2–4].

Although the existing virtualization techniques solve the problem of sharing a single wireless resource, they do not provide fairness guarantees among VAPs that serve different number of clients, as in the case illustrated in Fig. 1. Specifically, as the IEEE 802.11 MAC protocol [5] grants stations equal opportunities of accessing the channel [6], in such scenarios the throughput performance of the VAPs will be proportional to their number of users, and thus overloaded virtual WLANs will significantly affect the performance of the coexisting networks. Considering the example of Fig. 1 with saturated stations running standard EDCA, WLAN<sub>2</sub> will be



**Fig. 1** Scenario under study: a single Access Point hosting multiple virtual Wireless LANs on the same channel, each with different number of users.

taking 50% of the network throughput, while  $\text{WLAN}_1$  will receive 66% of the remaining bandwidth. Hence, the default configuration of the 802.11 protocol yields significant *inter-VAP unfairness*. If operators decide to share a given Access Point using virtualization, and they evenly share the deployment and maintenance cost of the infrastructure, this default behavior is highly undesirable, as those VAPs with few stations will obtain a small share of the wireless resource for the same cost. Based on this observation, we argue that a fair distribution of wireless resources between VAPs is required. Recent work [7] addresses this problem by proposing an architecture to deploy algorithms that enforce equal airtime among groups of stations. However, the solution requires clients to run a software application that involves non-negligible signaling with a central controller located at the AP and employ traffic shaping to limit the sending rates, which challenges its practical use. Additionally, this and previous proposals [7–10] do not address throughput optimization in virtualized WLANs, which is essential given the scarce nature of the wireless medium.

In this paper, we propose *C-VAP (Control-theoretic optimization of Virtual APs)*, a novel algorithm that maximizes the total throughput shared by virtual APs while providing fairness guarantees. The key technique of C-VAP is to employ control-theoretic tools to adjust the contention window ( $CW$ ) configuration of the stations within each VAP, to drive the WLAN to the optimal point of operation and evenly share the resources among the virtual networks. Specifically, with our approach the AP runs an independent proportional integrator (PI) controller for each VAP, which monitors the channel activity and drives the empty slot probability to the optimal value that maximizes performance, while simultaneously equalizing the probabilities of successful transmissions among the VAPs. To this end, each controller computes the optimal  $CW$  to be used by the clients of the VAP and broadcasts it to the stations by means of beacon frames, a feature specified in the current standard [5].

We conduct a performance analysis of the virtualized WLAN to characterize the optimal point of operation, which provides the foundations for the design of our algorithm. We configure the parameters of the PI controllers and prove system stability by undertaking a control-theoretic analysis of the WLAN. The key advantages of our solution are that (i) it is fully compliant with the 802.11 standard as it requires no modifications at the client side, while solely relying on existing AP functionality, (ii) it provides the same throughput performance for all the VAPs sharing the wireless resources irrespective of their number of users and their traffic patterns, (iii) it guarantees that non-saturated stations see all of their traffic served, and (iv) it maximizes the total throughput of the network.

The performance of the algorithm has been evaluated by means of simulation experiments under different network scenarios. The results show that our proposal significantly outperforms the default 802.11 scheme in terms of throughput, while providing fairness gains of up to 50% as compared to both EDCA and the static configuration of the  $CW$  that maximizes throughput in the whole system. Furthermore, we show that our approach maximizes performance even when not all the stations are backlogged; in this case, non-saturated VAPs see their throughput demand satisfied while the remaining resources are equally shared among the more demanding VAPs.

The rest of the paper is organized as follows. In Sec. 2 we present our system model and derive the optimal point of operation of a virtualized WLAN, Sec. 3 describes the proposed algorithm, a control-theoretic analysis is conducted in Sec. 4 to configure the algorithm’s parameters and in Sec. 5 we evaluate the performance of our proposal through simulation experiments. Sec. 6 summarizes the related work and, finally, Sec. 7 concludes the paper.

## 2 System Model and Optimization

We consider the case of  $N$  different virtual WLANs sharing the resources of a single AP and operating on the same carrier frequency.<sup>1</sup> We assume ideal channel conditions, and that all stations are in carrier-sensing range of each other, regardless of the virtual AP they are associated with. In this way, collisions are the only source of frame losses. Such ideal channel conditions have been widely used in the past (see e.g. [6, 12, 13]) and been proven to yield a good level of accuracy in experimental scenarios [14]. We consider stations are using a single transmission queue (note that following [15] the analysis can be easily extended to account for multiple active EDCA queues per station). Given that our

<sup>1</sup> In case of overlapping BSS scenarios, our mechanism can be independently implemented on each AP, as long as they employ appropriate dynamic channel assignment schemes such as, e.g., [11].

approach computes the optimal point of operation according to the observed network conditions, the exponential backoff scheme is not required, (furthermore, this would increase jitter) and therefore we set  $CW_{min,i} = CW_{max,i} = CW_i$  (we refer the reader to [15] for a detailed discussion and validation of this argument). We denote with  $CW_i$  the configuration of the contention window parameter that the virtual AP  $i$  (VAP $_i$ ) announces to its  $n_i$  associated stations, with  $i \in \{1, \dots, N\}$ . Assuming that all clients operate in saturation conditions, i.e., they always have a frame ready for transmission,<sup>2</sup> the probability that a station transmits at a randomly chosen slot time is given by [12],

$$\tau_i = \frac{2}{1 + CW_i}, \quad (1)$$

and the total throughput obtained by clients associated with VAP $_i$ , denoted as  $R_i$ , can be computed as [13]

$$R_i = \frac{\mathbb{E}[\text{payload VAP}_i/\text{slot}]}{\mathbb{E}[\text{slot length}]} = \frac{S_i L}{P_e T_e + (1 - P_e) T_o}, \quad (2)$$

where  $S_i$  is the probability that a slot contains a successful transmission from VAP $_i$ ,  $L$  is the average frame length,  $P_e$  is the probability that a slot is empty,  $T_e$  is the corresponding slot length in this case and  $T_o$  is the average length of an occupied slot, as derived in [12].<sup>3</sup>  $P_e$  is expressed as

$$P_e = \prod_{k=1}^N (1 - \tau_k)^{n_k}, \quad (3)$$

while  $S_i$  can be computed as

$$S_i = n_i \tau_i (1 - \tau_i)^{n_i - 1} \prod_{k=1, k \neq i}^N (1 - \tau_k)^{n_k} = \frac{n_i \tau_i}{1 - \tau_i} P_e. \quad (4)$$

The above completes our throughput analysis. Based on this model, we next address the optimization of the  $CW_i$  parameters of all VAPs in order to fulfill two key objectives. Namely, our goal is to design an algorithm that ensures the following two requirements:

1. All VAPs obtain the same performance when the network is fully loaded, regardless of their number of stations, i.e.,

$$R_i = R_j \quad \forall i, j$$

2. The overall network performance is maximized, i.e.,

$$\max \sum R_i$$

<sup>2</sup> Later on we relax this assumption and demonstrate that performance is optimized even when some stations are not saturated.

<sup>3</sup> Although for simplicity reasons we assume throughout the paper a fixed frame length, this assumption could be relaxed following our previous work [13].

To derive the condition that guarantees the **first objective** is achieved, we rewrite (2) as

$$R_i = \frac{\frac{n_i \tau_i}{1 - \tau_i} P_e L}{T_o - (T_o - T_e) P_e}. \quad (5)$$

With the above, it can be easily seen that the first objective imposes the following constraint on the transmission probabilities:

$$\frac{n_i \tau_i}{1 - \tau_i} = \frac{n_j \tau_j}{1 - \tau_j}, \quad (6)$$

which, assuming  $\tau_i \ll 1 \forall i$ , can be approximated by<sup>4</sup>

$$n_i \tau_i \approx n_j \tau_j. \quad (7)$$

Based on this result, we next address the **second objective** of our algorithm, namely, maximizing the throughput obtained by any VAP $_i$  (given the first objective, this is equivalent to maximizing the total throughput). Using the same approximation  $\tau \ll 1$  on (5) yields

$$R_i \approx \frac{n_i \tau_i L}{T_o / P_e - (T_o - T_e)} = \frac{n_i \tau_i L}{T_o \prod_k (1 - \tau_k)^{-n_k} - (T_o - T_e)},$$

which can be further approximated as

$$R_i \approx \frac{n_i \tau_i L}{T_o e^{\sum \tau_k n_k} - (T_o - T_e)} = \frac{n_i \tau_i L}{T_o e^{N \tau_i n_i} - (T_o - T_e)}.$$

The optimal  $\tau_i$ , denoted as  $\tau_i^*$ , can be obtained by solving

$$\frac{dR_i}{d\tau_i} = 0,$$

which leads to the following non-linear equation:

$$n_i [T_o e^{N \tau_i n_i} - (T_o - T_e)] - (n_i \tau_i) T_o e^{N \tau_i n_i} N n_i = 0.$$

To solve this equation, we proceed as in [12, 16] and use a Taylor expansion to approximate the exponential, i.e.,  $e^x = 1 + x + x^2/2 + \dots$ , and, given  $\tau_i \ll 1$ , we neglect the  $\tau_i$  terms above second order, which leads to the following expression for  $\tau_i^*$ :

$$\tau_i^* = \frac{1}{N n_i} \sqrt{\frac{2 T_e}{T_o}} \quad (8)$$

Thus, we obtain the optimal  $CW$  configuration by substituting the above in (1),

$$CW_i = \frac{2}{\tau_i^*} - 1. \quad (9)$$

Finally, we compute the probability of an empty slot,  $P_e$ , when all stations are configured as above, which will characterize the **point of operation** of the WLAN under optimal

<sup>4</sup> Note that this assumption is reasonable, as large values of the transmission probability would lead to high collision probability and hence to an inefficient utilization of the WLAN.

configuration. To this end, we substitute (8) in (3), which results in

$$P_e^* = \prod_k \left( 1 - \frac{1}{Nn_k} \sqrt{\frac{2T_e}{T_o}} \right)^{n_k}, \quad (10)$$

This expression can be approximated as

$$P_e^* \approx \prod_k e^{-\frac{1}{N} \sqrt{\frac{2T_e}{T_o}}} = e^{-\sqrt{\frac{2T_e}{T_o}}}. \quad (11)$$

The above shows that, under optimal operation with saturated stations, the probability of an empty slot is a constant independent of the number of VAPs and stations. The key approximation of this paper is to assume that this also holds when there are non-saturated stations in the system. The accuracy of this approximation will be assessed in Sec. 5.

### 3 C-VAP Algorithm

From the analysis of Sec. 2 we know that the optimal point of operation of the system as given by  $P_e^*$  does not depend on the number of VAPs, the number of stations, or their activity. This suggests that  $P_e^*$  can be used as a **reference signal**, to assess how far the network is operating from this optimal point and react accordingly. A key challenge, though, is to appropriately react when the system deviates from  $P_e^*$ : if the reaction is not quick enough, this will result in wastage of channel time; on the other hand if the reaction is too prompt, the system may turn unstable due to the inherent randomness of the EDCA mechanism.

Control theory is a particularly suitable tool to address this challenge, since it provides the necessary apparatus to guarantee the convergence and stability of adaptive algorithms. Therefore, in this paper we propose C-VAP (Control-theoretic optimization of Virtual APs), a mechanism based on the classic control system depicted in Fig. 2, where each VAP runs an independent controller in order to compute the  $CW$  configuration of its clients.<sup>5</sup> Specifically, we employ a proportional integral (PI) controller [17], a well-known device from classic control theory that has been previously applied to a number of networking algorithms in the literature [16, 18, 19]. A key advantage of using a PI controller is that it is simple to design, configure and implement with existing hardware [16].

As shown in the figure, the PI controller of  $VAP_i$  takes the error signal  $e_i$  as input and provides the control signal  $o_i$  as output, which is then used to compute the  $CW_i$  announced by  $VAP_i$ , thereby controlling the aggressiveness of the  $n_i$  stations. The error signal serves to evaluate the state of the system. If the system is operating at the desired point,

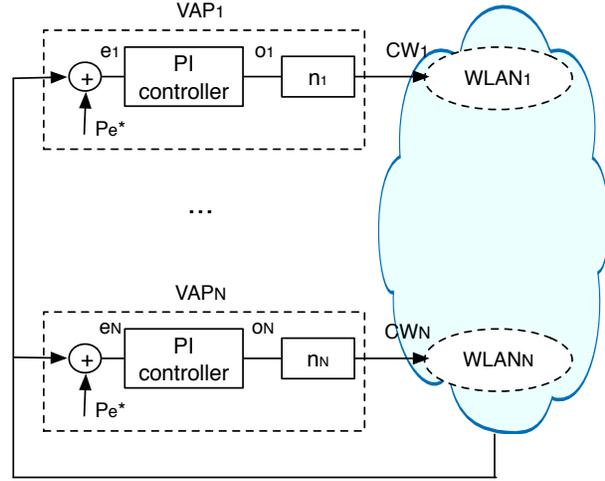


Fig. 2 Use of a different PI controller per Virtual AP.

the error signal of all VAPs will be zero. Otherwise, a non-zero error will drive the system from its current state towards the optimal state. In our approach, the error signal  $e_i$  is designed to fulfill the two objectives identified previously, namely (i) VAPs fairly share the system resources, and (ii) the overall throughput is maximized.

In order to satisfy the above requirements, we take the error signal as the sum of two terms. The first one is given by:

$$e_{opt} = P_e^* - P_e, \quad (12)$$

where  $P_e$  is the estimated probability of an empty slot and  $P_e^*$  is the optimal value resulting from our analysis. This term ensures that if the network operation yields an empty slot probability higher than the desired value (corresponding to a suboptimal utilization of the channel), the error will be negative, thus triggering a decrease of the  $CW_i$  and therefore an increase in the channel activity.

The second term of the error signal is:

$$e_{fair,i} = (N-1)S_i - \sum_{j \neq i} S_j. \quad (13)$$

This term of the error ensures that if  $VAP_i$  is obtaining a share of the total bandwidth larger than the average of the other  $(N-1)$  VAPs due to employing a smaller  $CW$  configuration, the error will be positive, thus reducing the aggressiveness of the stations associated to  $VAP_i$ .

The combination of (12) and (13) leads to the following error signal:

$$e_i = e_{opt} + e_{fair,i} = P_e^* - P_e + (N-1)S_i - \sum_{j \neq i} S_j \quad (14)$$

<sup>5</sup> Although in the figure we represent each VAP as a different block, they all run on the same physical device and therefore they can easily share operation parameters, e.g., sniffed frames.

Theorem 1 included in the Appendix guarantees that there exists a unique point of operation at which both terms of

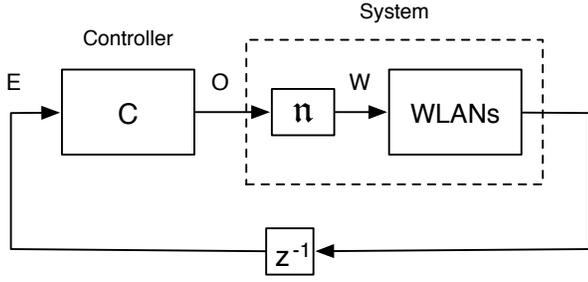


Fig. 3 Control system.

the above error signal are equal to zero, this being our target configuration specified by (8). Note that this result is of particular importance, as it ensures that there exists a single point of operation for the whole system despite the independent PI instances running at each VAP.

#### 4 Control-Theoretic Analysis

To appropriately configure our PI controllers, we conduct a control-theoretic analysis of the closed-loop system depicted in Fig. 2, which can be expressed in the form of Fig. 3, where

$$E = \begin{pmatrix} e_1 \\ e_2 \\ \vdots \\ e_N \end{pmatrix} = \begin{pmatrix} P_e^* - P_e + (N-1)S_1 - \sum_{j \neq 1} S_j \\ P_e^* - P_e + (N-1)S_2 - \sum_{j \neq 2} S_j \\ \vdots \\ P_e^* - P_e + (N-1)S_N - \sum_{j \neq N} S_j \end{pmatrix}, \quad (15)$$

$$O = \begin{pmatrix} o_1 \\ o_2 \\ \vdots \\ o_N \end{pmatrix}, \quad (16)$$

and

$$W = \begin{pmatrix} CW_1 \\ CW_2 \\ \vdots \\ CW_N \end{pmatrix}, \quad (17)$$

Our control system consists of one PI controller responsible for each VAP<sub>*i*</sub>, which takes  $e_i$  as input and gives  $o_i$  as output. Each VAP takes this output signal and multiplies it by the number of associated stations  $n_i$ , and the resulting value is broadcast to the associated stations as the  $CW$  to use during the next interval. Following this behavior, we can express the relationship between  $E$  and  $W$  as follows:

$$W(z) = \mathfrak{N} \cdot O = \mathfrak{N} \cdot C \cdot E(z), \quad (18)$$

where

$$\mathfrak{N} = \begin{pmatrix} n_1 & 0 & \cdots & 0 \\ 0 & n_2 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & n_N \end{pmatrix}, \quad (19)$$

and

$$C = \begin{pmatrix} C_{PI}(z) & 0 & \cdots & 0 \\ 0 & C_{PI}(z) & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & C_{PI}(z) \end{pmatrix}, \quad (20)$$

with  $C_{PI}(z)$  being the  $z$ -transform of a PI controller, i.e.,

$$C_{PI}(z) = K_P + \frac{K_I}{z-1}. \quad (21)$$

In order to analyze this closed loop we need to characterize the cluster of VAPs as a system with a transfer function  $H$  that takes as input the  $o_i$ 's and provides as output the error signals  $e_i$ 's. Since our system acts with beacon frequency, typically 100 ms, we can safely assume that the channel measurements obtained over a beacon interval correspond to stationary conditions. This implies that the error does not depend on the previous values, but only on the output value computed in the previous interval; this is modeled by the term  $z^{-1}$  in the figure, which shows that the error signal at a given instance is computed with the output signal of the previous interval.

Following the above,  $E$  can be computed from  $O$  by multiplying its elements by their respective  $n_i$ 's to obtain the  $W$  vector, using (1) to compute the respective  $\tau_i$ 's, and expressing  $P_e$  and the  $S_i$ 's as a function of the  $\tau_i$ 's, following (4) and (3). This gives a nonlinear relationship between  $E$  and  $O$ . In order to express this relationship as a transfer function, we linearize it when the system suffers small perturbations around its stable point of operation. Note that the stability of the linearized model guarantees that our system is locally stable [18], which is confirmed by the performance evaluation results presented in Section 5.

We express the perturbations around the point of operation as follows:

$$o_i = o_{i,opt} + \delta o_i \quad (22)$$

where  $o_{i,opt}$  is the  $o_i$  value that yields  $\tau_i^*$ .

With the above, the perturbations suffered by  $E$  can be approximated by

$$\delta E = H \cdot \delta O \quad (23)$$

where

$$H = \begin{pmatrix} \frac{\partial e_1}{\partial o_1} & \frac{\partial e_1}{\partial o_2} & \cdots & \frac{\partial e_1}{\partial o_N} \\ \frac{\partial e_2}{\partial o_1} & \frac{\partial e_2}{\partial o_2} & \cdots & \frac{\partial e_2}{\partial o_N} \\ \vdots & \vdots & \ddots & \vdots \\ \frac{\partial e_N}{\partial o_1} & \frac{\partial e_N}{\partial o_2} & \cdots & \frac{\partial e_N}{\partial o_N} \end{pmatrix}. \quad (24)$$

The above partial derivatives can be computed as

$$\frac{\partial e_i}{\partial o_j} = \frac{\partial e_i}{\partial \tau_j} \frac{\partial \tau_j}{\partial CW_j} \frac{\partial CW_j}{\partial o_j}, \quad (25)$$

where we have, according to our system,

$$\frac{\partial CW_j}{\partial o_j} = n_j, \quad (26)$$

while (1), evaluated at the stable point of operation, yields,

$$\frac{\partial \tau_j}{\partial CW_j} = -\frac{1}{2} \tau_j^2. \quad (27)$$

We next compute  $\partial e_i / \partial \tau_j$  for  $j \neq i$  that, after some operations, yields the following

$$\frac{\partial e_i}{\partial \tau_j} = \frac{n_j P_e}{1 - \tau_j} \left( 1 - (N - 1) \frac{n_i \tau_i}{1 - \tau_i} - \frac{1}{1 - \tau_j} + \sum_{k \neq i} \frac{n_k \tau_k}{1 - \tau_k} \right), \quad K_U = \frac{T_o}{P_e^* T_e}. \quad (28)$$

which, evaluated at the stable point of operation (with  $n_i \tau_i \approx n_j \tau_j$ ) and assuming  $\tau_j \ll 1$ , results in the following

$$\frac{\partial e_i}{\partial \tau_j} \approx 0. \quad (29)$$

If we now compute  $\partial e_i / \partial \tau_i$ , we obtain

$$\frac{\partial e_i}{\partial \tau_i} = \frac{n_i P_e}{1 - \tau_i} \left( 1 + \frac{N - 1}{1 - \tau_i} - \frac{(N - 1) n_i \tau_i}{1 - \tau_i} + \sum_{k \neq i} \frac{n_k \tau_k}{1 - \tau_k} \right), \quad (30)$$

which, evaluated at the stable point of operation, results in

$$\frac{\partial e_i}{\partial \tau_i} \approx N n_i P_e. \quad (31)$$

Combining all the above yields

$$H = K_H I \quad (32)$$

where

$$K_H = -\frac{P_e^* T_e}{N T_o} \quad (33)$$

Thus, our system is now fully characterized by the matrices  $C$  and  $H$ . The next step is to configure the  $K_P$  and  $K_I$  parameters of the PI controller. Following Theorem 2 (provided in the Appendix), we have that the  $\{K_P, K_I\}$  setting has to fulfill the following condition for the system to be stable:

$$K_I < K_P < \frac{N T_o}{P_e^* T_e} + \frac{1}{2} K_I \quad (34)$$

In addition to guaranteeing stability, our goal in the configuration of the PI parameters is to find the right trade-off

between speed of reaction to changes and oscillation under stable conditions. To find this trade-off we use the *Ziegler-Nichols* rules [20] as follows: (i) we first compute the  $K_P$  value that leads to instability when  $K_I = 0$ , denoted as  $K_U$ , and configure  $K_P = 0.4 K_U$ ; (ii) we then compute the oscillation period  $T_I$  when the system is unstable, and configure  $K_I = K_P / (0.85 T_I)$ .

To compute  $K_U$  we set  $K_I = 0$  in (34), which gives

$$K_P < \frac{N T_o}{P_e^* T_e}. \quad (35)$$

Since the above is a function of  $N$ , to find a bound independent of the number of VAPs we set  $N = 1$ , as this constitutes the most restrictive case on  $K_P$ , which leads to

$$K_U = \frac{T_o}{P_e^* T_e}. \quad (36)$$

During unstable operation, a given set of input values may change their sign up to every time interval, yielding an oscillation period of two ( $T_I = 2$ ). Thus, we obtain the following configuration for the  $K_P$  and  $K_I$  parameters:

$$K_P = 0.4 \frac{T_o}{P_e^* T_e}, \quad (37)$$

$$K_I = \frac{0.2}{0.85} \frac{T_o}{P_e^* T_e}.$$

It is easy to verify that this configuration meets the condition of (34) and therefore guarantees the stability of the system.

## 5 Performance Evaluation

To evaluate the performance of the proposed algorithm, we conducted an extensive set of simulation experiments. For this purpose, we have extended the simulator used in [13, 16],<sup>6</sup> which is an event-driven network simulator based on the OMNeT++<sup>7</sup> framework that closely follows the details of the MAC protocol of 802.11 EDCA for each contending station. The simulations are performed with the system parameters of the IEEE 802.11a physical layer [21] and the 54 Mbps PHY rate, assuming a channel in which frames are only lost due to collisions and considering stations transmit frames with a payload size of 1000 bytes. We present averages over 10 simulation runs, each lasting 300 seconds. We also compute 95% confidence intervals for the throughput figures, and confirm that in all cases their width is well below 1% of the average.

Unless otherwise specified, we assume that all stations are saturated. We compare the performance of our proposal,

<sup>6</sup> The source code of the simulator used in [13, 16] is available at <http://enjambre.it.uc3m.es/~ppatras/owsim/>.

<sup>7</sup> <http://www.omnetpp.org>

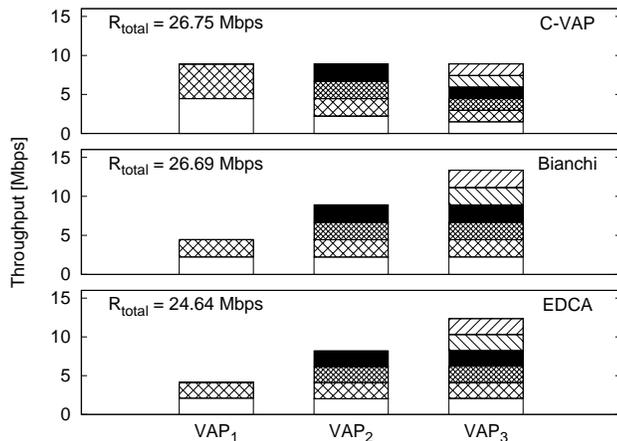


Fig. 4 Throughput distribution among VAPs.

C-VAP, against the following two alternatives: (i) the standard default configuration, denoted as ‘EDCA’ [5], and (ii) the static  $CW$  configuration that maximizes total throughput of the WLAN (regardless of VAPs associations) under saturation conditions [12], labeled as ‘Bianchi’ in the plots.

### 5.1 Throughput & Fairness

Our first aim is to validate that C-VAP is able to maximize the throughput performance in the WLAN while providing all VAPs with a fair share of the resources. To this end, we consider the case of three VAPs with  $n_i = \{2, 4, 6\}$  saturated stations, respectively, and compute the throughput obtained by each station. The results, grouped by VAP, are presented in Fig. 4.

The figure shows that C-VAP succeeds in providing all VAPs with the same throughput (8.9 Mbps approximately) regardless of their number of users (the stacked boxes show the throughput attained by each station). In contrast, the other two alternatives fail to provide fairness among VAPs, and instead favor the VAPs with higher number of associated stations. Precisely, the Jain Fairness Index (JFI) [22] for the per-VAP throughput distribution yields values of 1, 0.85 and 0.86 for C-VAP, Bianchi and EDCA, respectively. Note that C-VAP is able not only to enforce fairness among VAPs, but also to maximize the overall throughput in the system; indeed, the total throughput obtained with C-VAP and Bianchi is approximately 26.7 Mbps, while the default EDCA configuration proves to be too aggressive for the considered number of stations and yields a total throughput of 24.6 Mbps.

We next analyze how the performance of the three approaches varies when the number of stations associated with the VAPs changes. For this purpose, we consider the case of two VAPs, with  $n_1 = 5$  stations and an increasing number of saturated stations associated with VAP<sub>2</sub>. For each considered case, we obtain the total throughput in the system and

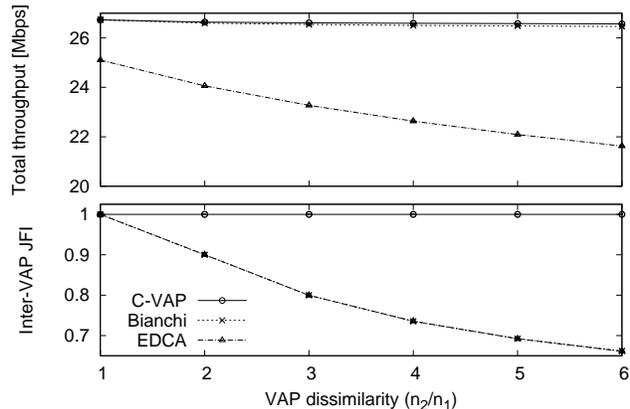


Fig. 5 Throughput performance and inter-VAP fairness.

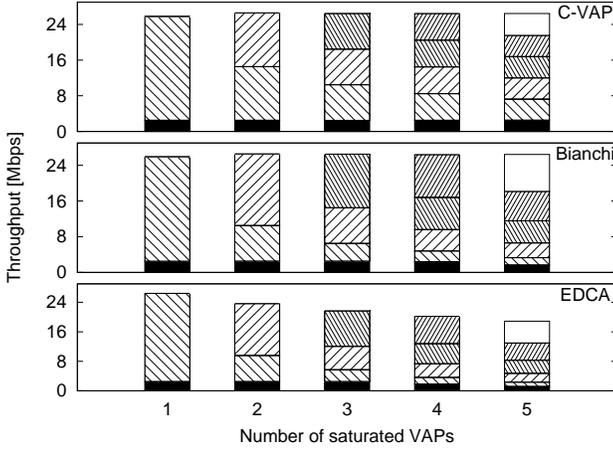
the *Inter-VAP* JFI as in the previous case. The results are depicted in Fig. 5.

The figure confirms the results obtained in the previous scenario. First, it can be seen that as the total number of stations increases, the EDCA configuration results overly aggressive and therefore the overall throughput performance is degraded; in contrast, both C-VAP and Bianchi’s approach are able to maximize the total throughput. On the other hand, only C-VAP is able to provide a fair resource distribution with  $JFI \approx 1$  in all cases, while the other two approaches excessively favor VAP<sub>2</sub> as its number of users increases, which results in JFI values significantly smaller than one. More specifically, although Bianchi’s approach optimally configures the  $CW$  and maximizes the overall throughput, it does not take into account users associations and therefore penalizes the VAP with the least number of stations.

The above results confirm that, in saturation conditions, our mechanism is able to maximize the overall throughput in the system while guaranteeing a fair distribution of the resources among VAPs. In what follows, we study the case of non-saturation scenarios and assess the effectiveness of the configuration of the PI controller under both steady operation and dynamic conditions.

### 5.2 Non-saturation Scenarios

We next analyze the behavior of the proposed algorithm in non-saturated traffic conditions, to confirm that the good properties of C-VAP are maintained even when stations are not constantly backlogged with frames to transmit. Note that under non-saturation conditions our goals are the following: (i) non-saturated stations see all their traffic served, as long as they generate less than the saturation rate; (ii) VAPs with saturated stations fairly share resources regardless of the number of stations; and (iii) the overall network performance is maximized.

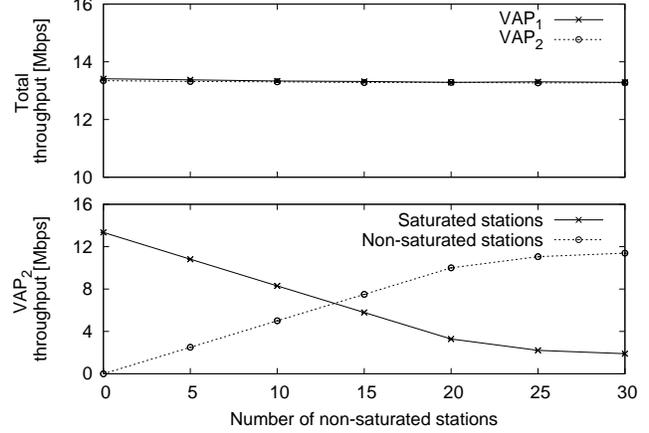


**Fig. 6** Throughput performance with non-saturated stations associated to one VAP.

In our first set of experiments, we consider the case of one VAP with  $n_0 = 5$  stations generating 500 kbps Poisson traffic, and an increasing number of VAPs, each having  $n_i = 5 \cdot i$  saturated stations, i.e., the first VAP with saturated stations associated has  $n_1 = 5$ , the second VAP that we add has  $n_2 = 10$ , and so on. The aggregated throughput per VAP is depicted in Fig. 6 for the three considered mechanisms. We mark with solid black the throughput obtained by the non-saturated VAP, while the other boxes depict the throughput of the saturated VAPs. The results can be summarized as follows:

- C-VAP satisfies all the considered objectives, as the non-saturated VAP always sees all of its traffic served, while the other VAPs fairly share the available bandwidth, which is furthermore maximized.
- The optimal-throughput configuration (Bianchi) only satisfies the non-saturated VAP as long as the number of saturated VAPs is below 5. Otherwise, despite the overall throughput is maximized as with C-VAP, the uneven distribution of resources harms the performance of non-saturated traffic and favors the VAPs with more associated stations.
- Finally, EDCA fails to fulfill all the above objectives, as it does not serve non-saturation traffic appropriately, the throughput is not maximized, and resources are unevenly shared.

The above scenario confirms that C-VAP is able to guarantee a fair sharing of resources when a VAP with non-saturated stations is contending vs. other VAPs with saturated stations. We next analyze the case when there are saturated and non-saturated stations associated with the same VAP. For this purpose, we consider the case of two VAPs, with VAP<sub>1</sub> having  $n_1 = 5$  saturated stations, and VAP<sub>2</sub> having 5 saturated stations and a varying number of non-saturated stations associated (like in the previous case, non-



**Fig. 7** Performance vs. increasing number of non-saturated stations.

saturated stations generate 500 kbps Poisson traffic). We compute the aggregated throughput per VAP, and the throughput distribution within VAP<sub>2</sub>, and depict the results in Fig. 7.

The figure shows that the good properties of the throughput distribution are maintained also in this case. Indeed, in all cases the VAPs fairly share the resources like in the previous cases, each one getting about 13.5 Mbps (top subplot of the figure). Examining the throughput distribution within VAP<sub>2</sub> (bottom subplot of the figure), again we see that saturated stations are able to maximize their performance as long as non-saturated stations see their traffic served. Once the number of non-saturated stations increases above 20, the resources are fairly distributed among the stations within the VAP.

### 5.3 Configuration of the Controller

The main objective in the setting of the  $K_P$  and  $K_I$  parameters proposed in Sec. 4 is to achieve a good tradeoff between stability and speed of reaction to changes in the system.

To validate that our system guarantees a stable behavior, we consider the case of three VAPs, with  $n_i = \{5, 10, 15\}$  saturated stations each, and analyze the evolution over time of the  $CW$  announced by each virtual AP for our  $\{K_P, K_I\}$  setting proposed in (37) and a configuration of these parameters 10 times larger. The results are depicted in Fig. 8. We observe from the figure that with the proposed setting (labeled “ $K_P, K_I$ ”) the systems performs stably with minor deviations of the  $CW$ s around their average values; in contrast, for the other setting (labeled “ $K_P * 10, K_I * 10$ ”) the announced values drastically oscillate and the system shows unstable behavior.

We next investigate the speed with which the system reacts to changes in the working conditions. To this end, we consider the case of two VAPs, namely VAP<sub>1</sub> and VAP<sub>2</sub>. The first one has associated  $n_1 = 5$  saturated stations, while

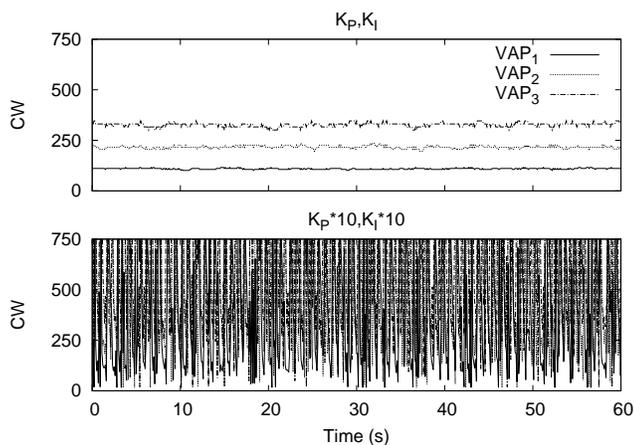


Fig. 8 Stability of the PI controller configuration.

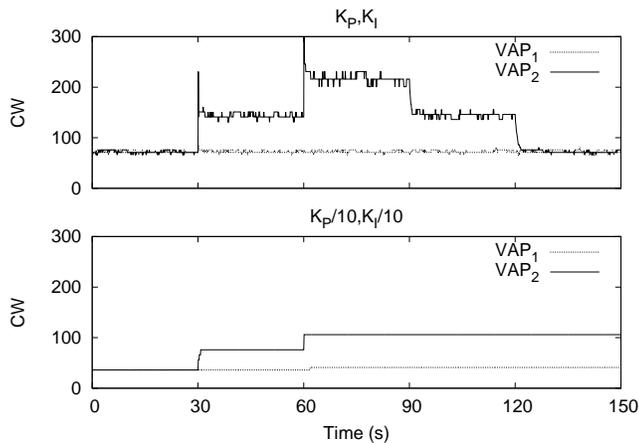


Fig. 9 Speed of reaction to changes.

for VAP<sub>2</sub> the number of associated stations varies over time as follows: in the beginning there are  $n_2 = 5$  stations, at  $t = 30$  s 5 more stations join the network, and subsequently 5 more stations join the VAP at  $t = 60$  s, resulting in a total of  $n_2 = 15$  stations. Then, after 30 s, 5 stations leave VAP<sub>2</sub>, and again 5 more stations leave at  $t = 120$  s, the WLAN returning to the initial state with  $n_1 = 5$  and  $n_2 = 5$ . For this experiment, we examine the evolution over time of the  $CW$  announced by each VAP for our  $\{K_P, K_I\}$  setting, as well as for a configuration of these parameters 10 times smaller. The results are depicted in Fig. 9.

The figure shows that with our setting (“ $K_P, K_I$ ”), the system reacts fast to the changes described above, as the  $CW$  announced by VAP<sub>2</sub> reaches the new value almost immediately. In contrast, for the other setting (“ $K_P/10, K_I/10$ ”), the system cannot keep up with the changes and reacts too slowly.

We conclude that the proposed setting of  $\{K_P, K_I\}$  provides a good tradeoff between stability and speed of reaction to changes, since with a larger setting the system suffers

from instability and with a smaller one it reacts too slowly to changes.

## 6 Related Work

WLAN virtualization has recently become an important issue addressed by the research community. Wireless networks virtualization architectures are proposed in [23–25] and a virtual networking infrastructure using open source techniques is introduced in [3]. Design and implementation of solutions for supporting multiple virtual WiFi interfaces with a single physical device are discussed in [2,4]. TDMA-based approaches to WLAN virtualization are studied in [8] and [9], while the strengths and drawbacks of space and time based virtualization techniques are compared in [26]. AP virtualization for enabling efficient mobility management is described in [27, 28]. Client virtualization is employed for supporting simultaneous connectivity to multiple APs and achieving bandwidth aggregating in [29–32], while [10] exploits virtualization and multi-AP connectivity to improve video streaming performance. In [7] the problem of fair sharing of the uplink airtime across groups of users is considered in a network virtualization scenario.

However, none of the above works address the problem of throughput optimization in virtualized WLANs while providing fairness guarantees to virtual APs, which significantly limits their applicability to practical scenarios where service providers seek to maximize revenue from their wireless subscribers. In contrast to these works, we propose a standard compliant solution that can be easily deployed at the AP and which successfully maximizes the network performance while evenly sharing the resources among the virtual networks, irrespective of their number of users.

## 7 Conclusions

It is becoming increasingly common that operators share a physical device to create different virtual WLANs, for reasons varying from lack of available channels (and therefore to increase efficiency in coordinating with competitors), to infrastructure being owned by local businesses. In such circumstances, it is critical to guarantee fair sharing of resources between virtual WLANs while maximizing throughput and, therefore, revenue. While previous approaches have provided the means to enable virtualization or to optimally configure a single-owner WLAN, the problem of an optimal yet fair configuration has not been addressed. Furthermore, without a proper configuration, the default access scheme favors those operators with more clients, thus degrading the performance of the users attached to lightly loaded networks.

In this paper we proposed C-VAP, a novel mechanism that maximizes performance in virtualized WLANs scenar-

ios while ensuring fairness among competing providers. In contrast to previous work that introduces non-trivial changes to both the AP and the stations, our approach runs exclusively at the AP and relies only on standard functionality. Furthermore, by building on foundations from control theory, C-VAP is able to adapt to changes in the WLANs while guaranteeing system stability. Extensive simulations confirm the good properties of our mechanism, and results show that (i) our scheme outperforms the standard configuration in terms of throughput, (ii) it maintains fairness among virtual WLANs regardless of the network conditions, either in terms of number of stations or traffic patters (in contrast to the standard or the throughput-optimal configurations), and (iii) it promptly reacts to changes in network conditions while ensuring stable operation. Following our implementation experiences [14], we plan as part of our future work to assess the performance of C-VAP in a real-life testbed.

## Acknowledgements

This work has been supported by the European Community's Seventh Framework Programme (FP7-ICT-2009-5) under grant agreement n. 257263 (FLAVIA project).

## Appendix

**Theorem 1** *Given the definition of  $e_i$  in (14), there exists an unique solution to the system defined by  $e_i = 0 \forall i$  that satisfies  $e_{opt} = 0$  and  $e_{fair,i} = 0 \forall i$ .*

*Proof* By subtracting  $e_j$  from  $e_i$  we obtain

$$e_i - e_j = (N-1)S_i - S_j - (N-1)S_j + S_i = N(S_i - S_j), \quad (38)$$

and therefore, given that  $e_i = 0 \forall i, j$ , the above results in  $S_i = S_j \forall i, j$ , and therefore we have that  $e_{fair,i} = 0 \forall i$ . Furthermore, this results in the following relation (as already expressed in (6)),

$$\frac{n_i \tau_i}{1 - \tau_i} = \frac{n_j \tau_j}{1 - \tau_j}, \quad (39)$$

which specifies, for a given  $(n_i, n_j)$  pair, a one-to-one relationship between  $\tau_j$  and  $\tau_i \forall i, j$ , and therefore we can take e.g.  $\tau_1$  as reference. In this way, if we express  $e_{opt} = 0$  as

$$\prod (1 - \tau_k)^{n_k} = P_e^*, \quad (40)$$

we have that the rhs of the above equation is a constant between 0 and 1, while the lhs is a decreasing function of  $\tau_1$  from 1 to 0. Therefore there exists a unique solution that solves the above equation, thus ensuring also that  $e_{opt} = 0$ .

**Theorem 2** *The  $K_P$  and  $K_I$  relationship specified by (34) guarantees stability.*

*Proof* According to [33], we need to check that the following transfer function is stable

$$(I - z^{-1}CH)^{-1}C. \quad (41)$$

Computing the above yields

$$(I - z^{-1}CK_H I)^{-1}C = \frac{K_P + \frac{K_I}{z-1}}{1 - z^{-1} \left( K_P + \frac{K_I}{z-1} \right) K_H} I, \quad (42)$$

which can be expressed as

$$(I - z^{-1}CK_H I)^{-1}C = \frac{P(z)}{z^2 + za_1 + a_2} I, \quad (43)$$

where  $P(z)$  is a polynomial and

$$a_1 = -1(1 + K_P K_H) \quad (44)$$

$$a_2 = K_H(K_P - K_I) \quad (45)$$

According to [33], a sufficient condition for stability is that the zeros of the pole polynomial fall within the unit circle. This can be ensured by choosing the coefficients  $\{a_1, a_2\}$  that belong to the stability triangle [17]:

$$a_2 < 1, \quad (46)$$

$$a_1 < a_2 + 1, \quad (47)$$

$$a_1 > -1 - a_2. \quad (48)$$

Equation (46) is satisfied as long as  $K_P > K_I$ , while (48) is satisfied if  $K_I > 0$ . By operating in (47) we obtain the relationship  $K_P < -K_H^{-1} + K_I/2$ , which combined with the previous relations results in the conditions expressed by (34).

## References

1. T. Janevski, A. Tudzarov, P. Stojanovski, and D. Temkov, "Applicative Solution for Easy Introduction of WLAN as Value-Added Service in Mobile Networks," in *IEEE Vehicular Technology Conference (VTC) Spring*, April 2007, pp. 1096–1100.
2. T. Hamaguchi, T. Komata, T. Nagai, and H. Shigeno, "A Framework of Better Deployment for WLAN Access Point Using Virtualization Technique," in *IEEE International Conference on Advanced Information Networking and Applications Workshops (WAINA)*, April 2010, pp. 968–973.
3. G. Aljabari and E. Eren, "Virtualization of wireless LAN infrastructures," in *IEEE International Conference on Intelligent Data Acquisition and Advanced Computing Systems (IDAACS)*, vol. 2, September 2011, pp. 837–841.
4. L. Xia, S. Kumar, X. Yang, P. Gopalakrishnan, Y. Liu, S. Schoenberg, and X. Guo, "Virtual WiFi: bring virtualization from wired to wireless," in *ACM SIGPLAN/SIGOPS international conference on Virtual execution environments*, ser. VEE '11, Newport Beach, California, USA, 2011, pp. 181–192.

5. IEEE 802.11 WG, *Information Technology - Telecommunications and Information Exchange between Systems. Local and Metropolitan Area Networks. Specific Requirements. Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) specifications*. IEEE Std 802.11, 2007.
6. G. Berger-Sabbatel, A. Duda, O. Gaudoin, M. Heusse, and F. Rousseau, "Fairness and its impact on delay in 802.11 networks," in *IEEE Global Telecommunications Conference (GLOBECOM)*, vol. 5, November 2004, pp. 2967–2973.
7. G. Bhanage, D. Vete, I. Seskar, and D. Raychaudhuri, "SplitAP: Leveraging Wireless Network Virtualization for Flexible Sharing of WLANs," in *GLOBECOM 2010, 2010 IEEE Global Telecommunications Conference*, December 2010, pp. 1–6.
8. G. Smith, A. Chaturvedi, A. Mishra, and S. Banerjee, "Wireless virtualization on commodity 802.11 hardware," in *ACM international workshop on Wireless network testbeds, experimental evaluation and characterization*, ser. WinTECH '07, Montreal, Quebec, Canada, 2007, pp. 75–82.
9. S. Perez, J. Cabero, and E. Miguel, "Virtualization of the wireless medium: A simulation-based study," in *IEEE Vehicular Technology Conference (VTC) Spring*, April 2009, pp. 1–5.
10. S.-W. Ahn and C. Yoo, "Network interface virtualization in wireless communication for multi-streaming service," in *IEEE International Symposium on Consumer Electronics (ISCE)*, June 2011, pp. 67–70.
11. R. Akl and A. Arepally, "Dynamic Channel Assignment in IEEE 802.11 Networks," in *IEEE International Conference on Portable Information Devices (PORTABLE)*, May 2007, pp. 1–5.
12. G. Bianchi, "Performance Analysis of the IEEE 802.11 Distributed Coordination Function," *IEEE Journal on Selected Areas in Communications*, vol. 18, no. 3, pp. 535–547, March 2000.
13. P. Serrano, A. Banchs, P. Patras, and A. Azcorra, "Optimal Configuration of 802.11e EDCA for Real-Time and Data Traffic," *Vehicular Technology, IEEE Transactions on*, vol. 59, no. 5, pp. 2511–2528, June 2010.
14. P. Serrano, P. Patras, A. Mannocci, V. Mancuso, and A. Banchs, "Control Theoretic Optimization of 802.11 WLANs: Implementation and Experimental Evaluation," 2012. [Online]. Available: <http://arxiv.org/abs/1203.2970v1>
15. A. Banchs and L. Voller, "Throughput analysis and optimal configuration of 802.11e EDCA," *Computer Networks*, vol. 50, no. 11, pp. 1749–1768, 2006.
16. P. Patras, A. Banchs, P. Serrano, and A. Azcorra, "A Control-Theoretic Approach to Distributed Optimal Configuration of 802.11 WLANs," *IEEE Transactions on Mobile Computing*, vol. 10, pp. 897–910, June 2011.
17. K. Aström and B. Wittenmark, *Computer-controlled systems, theory and design*, 2nd ed. Prentice Hall International Editions, 1990.
18. C. Hollot, V. Misra, D. Towsley, and W.-B. Gong, "A Control Theoretic Analysis of RED," in *Proceedings of IEEE INFOCOM 2001*, Anchorage, Alaska, April 2001.
19. L. Grieco, G. Boggia, S. Mascolo, and P. Camarda, "A control theoretic approach for supporting quality of service in IEEE 802.11e WLANs with HCF," in *IEEE Conference on Decision and Control*, vol. 2, 2003, pp. 1586–1591.
20. G. F. Franklin, J. D. Powell, and M. L. Workman, *Digital Control of Dynamic Systems*, 2nd ed. Addison-Wesley, 1990.
21. 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.
22. R. Jain, Chiu, D.M., and W. Hawe, "A Quantitative Measure of Fairness and Discrimination for Resource Allocation in Shared Systems," *DEC Research Report TR-301*, 1984.
23. Y. He, J. Fang, J. Zhang, H. Shen, K. Tan, and Y. Zhang, "MPAP: virtualization architecture for heterogenous wireless APs," *SIGCOMM Comput. Commun. Rev.*, vol. 41, pp. 475–476, August 2010.
24. K.-K. Yap, M. Kobayashi, R. Sherwood, T.-Y. Huang, M. Chan, N. Handigol, and N. McKeown, "OpenRoads: empowering research in mobile networks," *SIGCOMM Comput. Commun. Rev.*, vol. 40, pp. 125–126, January 2010.
25. R. Matos, S. Sargento, K. Hummel, A. Hess, K. Tutschku, and H. de Meer, "Context-based wireless mesh networks: a case for network virtualization," *Telecommunication Systems*, pp. 1–14, March 2011.
26. R. Mahindra, G. Bhanage, G. Hadjichristofi, I. Seskar, D. Raychaudhuri, and Y. Zhang, "Space versus time separation for wireless virtualization on an indoor grid," in *Next Generation Internet Networks (NGI)*, April 2008, pp. 215–222.
27. Y. Grunenberger and F. Rousseau, "Virtual Access Points for Transparent Mobility in Wireless LANs," in *IEEE Wireless Communications and Networking Conference (WCNC)*, April 2010, pp. 1–6.
28. M. Berezin, F. Rousseau, and A. Duda, "Multichannel Virtual Access Points for Seamless Handoffs in IEEE 802.11 Wireless Networks," in *IEEE Vehicular Technology Conference (VTC Spring)*, May 2011, pp. 1–5.
29. R. Chandra and P. Bahl, "MultiNet: connecting to multiple IEEE 802.11 networks using a single wireless card," in *INFOCOM 2004*, vol. 2, March 2004.
30. S. Kandula, K. C.-J. Lin, T. Badirhanli, and D. Katabi, "Fat-VAP: aggregating AP backhaul capacity to maximize throughput," in *USENIX Symposium on Networked Systems Design and Implementation (NSDI)*, San Francisco, California, 2008.
31. D. Giustiniano, E. Goma, A. Lopez, and P. Rodriguez, "WiSwitcher: an efficient client for managing multiple APs," in *SIGCOMM workshop on Programmable routers for extensible services of tomorrow (PRESTO)*, Barcelona, Spain, 2009, pp. 43–48.
32. D. Giustiniano, E. Goma, A. Lopez Toledo, I. Dangerfield, J. Morillo, and P. Rodriguez, "Fair WLAN backhaul aggregation," in *ACM International conference on Mobile computing and networking (MobiCom)*, Chicago, Illinois, USA, 2010, pp. 269–280.
33. T. Glad and L. Ljung, *Control Theory: Multivariable and Nonlinear Methods*. Taylor & Francis, 2000.