

Feedback based Distributed Admission Control in 802.11 WLANs

Preetam Patil, Vipul Mathur, Varsha Apte and Kannan Moudgalya
Indian Institute of Technology–Bombay
Powai, Mumbai, Maharashtra 400076, India
Email: {pvpatil,vipulmathur,varsha,kannan}@iitb.ac.in

Abstract—A distributed connection admission control (CAC) scheme where stations independently admit or reject flows based on channel utilization threshold is easy to implement for WLANs. However, owing to variable protocol capacity, the advisable threshold also is variable, limiting the efficacy of fixed threshold based CAC. If the admission threshold is tuned to reflect the current protocol capacity, better throughput can be obtained while avoiding WLAN overload. In this paper, we propose a feedback based scheme to tune the utilization threshold for admission. We select the performance metric used for feedback, derive the system model from empirical data, and use control theory to design a feedback controller. Further, we present an analytical model and heuristics using which the controller can be adapted to work under diverse operating scenarios. Simulation results of proportional and proportional-integral controllers in OPNET suggest that the feedback based CAC is able to avoid WLAN overload and achieve high throughput despite large changes in protocol capacity.

I. INTRODUCTION

Connection admission control (CAC) is an important mechanism for avoiding overload and ensuring acceptable quality of service (QoS) in IEEE 802.11 WLANs based on distributed coordination function (DCF). The CAC decision is based on three inputs: (a) the network capacity, (b) admitted load, and (c) the resource demands of the requested flow.

The WLAN capacity is the protocol capacity of the 802.11 MAC that depends on (a) channel data rates used by individual stations, (b) packet sizes of entire WLAN traffic, and (c) the number of stations. As these parameters can vary during network operation, capacity estimation for CAC needs to be performed online.

Distributed CAC is appealing for WLANs because it allows stations to admit flows without explicit message-passing. However, determining the total admitted load is a challenge in a distributed CAC. Instead, channel utilization serves as a useful normalized measure of total load on WLAN since all stations can measure it easily. Thus a CAC scheme based on measured channel utilization being lower than a threshold is easy to implement. Hence many *fixed threshold* based approaches have been proposed [1]–[3]. However, the use of fixed utilization threshold may either lead to overload or underutilization of WLAN. A capacity-dependent admission threshold would provide more consistent performance. The threshold adjustment can be done in a feed-forward manner or using some form of feedback from the WLAN.

An example of feed-forward approach is the use of an analytical model [4] to compute the advisable load from model inputs (average packet size, data rate and number of stations), requiring detailed knowledge of the system being controlled. Alternatively, feedback from the WLAN about performance measures such as delay, queue length or collision rate can be used to tune the threshold. Feedback control can work with an approximate input-output model of the system, as corrective actions in successive steps overcome the inaccuracies of the model itself. However, feedback based systems face the problem of instability or sluggishness and require careful design.

Park et al. [5] have proposed a feedback-based scheme for 802.11 DCF that interacts with TCP's flow control to control TCP flows' sending rate. The feedback based scheduler proposed by Boggia et al. [6] for controlled access mode of 802.11e assigns transmission opportunity (TXOP) to traffic aggregates. Both these schemes monitor interface queue length to assign service rates to flows. Both schemes assume a tight control on flows' packet generation rate, reducing the variability of a flows performance, even on a short time-scale. On the other hand, we wish to develop WLAN CAC that works in presence of load and capacity variations, without requiring application flow control or controlled channel access.

In this paper, we apply control theory to tune the channel utilization threshold in a distributed CAC scheme for WLANs. A station queue length feedback is used to manipulate the channel utilization threshold. The goal of our feedback based CAC is to prevent WLAN overload while trying to maximize WLAN throughput. Using an empirically derived system model, we design proportional as well as proportional-integral controllers for a given WLAN scenario. We also develop an adaptive controller that predicts stable operating points for different WLAN scenarios using heuristics and an analytical model. We test our method by developing a prototype of WLAN CAC and feedback controller in OPNET simulator [7]. Simulations indicate that the feedback based CAC provides significant improvement over fixed threshold CAC.

The rest of the paper is as follows. Section II describes the threshold based CAC. In Section III we formulate the feedback control problem and derive the system model from empirical data. Design of feedback controllers and their performance analysis is presented in Section IV. Section V describes the adaptive controller using operating point prediction. We conclude with Section VI.

II. UTILIZATION THRESHOLD BASED CAC

The CAC mechanism and the feedback controller run on every WLAN station without any central coordination. For empirical data generation and performance evaluation of the feedback controller, we developed a prototype of threshold based CAC in OPNET simulator [7]. The *Admission Control Unit (ACU)* receives channel and interface measurements from the WLAN interface and admits or rejects flows as requested by the *Flow/Packet generator*. The *Flow/Packet generator* starts generating packets as per the specified arrival process only after the ACU admits the flow. Further, if the ACU determines that the current utilization θ_a is over the threshold θ , it terminates flows one by one¹. A terminated flow exits from the system immediately. The feedback controller to tune the threshold is also a part of the ACU. The ACU filters the measured utilization $\tilde{\theta}_a$ through an exponential weighted average filter to reduce the effect of stochastic noise on filtered estimate $\hat{\theta}_a$.

$$\hat{\theta}_a(k+1) = (1-\alpha)\tilde{\theta}_a(k+1) + \alpha\hat{\theta}_a(k), \quad 0 \leq \alpha \leq 1 \quad (1)$$

Further, to protect against frequent actions of flow admission followed by termination, we maintain a $\pm 5\%$ band around the current threshold. The approximate utilization demand θ_f of the requested flow is computed from throughput demand S_f , ignoring overheads. Algorithm 1 outlines three parallel threads of CAC operation.

Algorithm 1 Threshold based CAC

<pre> loop {Flow admission} receive flow request with S_f $\theta_f \leftarrow S_f / (\text{channel data rate})$ if $(\hat{\theta}_a + \theta_f) < 0.95\theta$ then admit flow else reject flow end if end loop </pre>	<pre> loop {Measurement update} receive measurement $\tilde{\theta}_a$ $\tilde{\theta}_a \leftarrow (1-\alpha)\tilde{\theta}_a + \alpha\tilde{\theta}_a$ end loop loop {Flow termination} if $\hat{\theta}_a > 1.05\theta$ then terminate a flow end if end loop </pre>
--	--

A. CAC operation and time scales

The different timescales for measurements and CAC actions, in decreasing order of granularity, are:

- WLAN measurements are fed to the ACU every 0.5s.
- The ACU terminates one flow every 1.7s if $\hat{\theta}_a > \theta$
- The Flow/Packet generator tries to start² a flow every 2s.

The choice of the intervals is governed by two conflicting requirements. While stochastic noise (due to random access WLAN and packet arrival process) makes it necessary to obtain measurements over longer intervals or to use more measurement samples, CAC needs to respond to capacity changes quickly. We select the above time intervals and an α value to achieve a balance between the two requirements.

¹The most recently admitted flow is terminated in our implementation; a different flow termination scheme could be used if required.

²Even if flow requests arrive at any interval, the CAC will need to ensure a minimum separation between consecutively admitted flows.

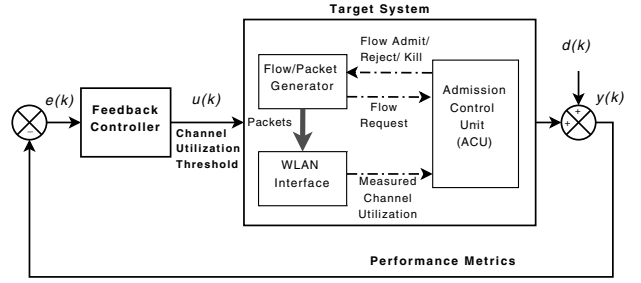


Fig. 1: WLAN with CAC as target system

III. FEEDBACK BASED THRESHOLD TUNING

The goal of feedback control is to keep the system's output (the controlled parameter) at the desired reference. We restrict the discussion here to the domain of digital control [8], [9].

A. Control-theoretic System Formulation

With reference to Fig. 1, the variables marked $r(k)$, $e(k)$, $u(k)$, $d(k)$ and $y(k)$ are reference input, error, control input, disturbance and measured output respectively. The control input $u(k)$ is the channel utilization threshold $\theta(k)$. The measured output $y(k)$ is the WLAN interface queue (IFQ) length $L(k)$ in bytes, as against response time that is cumbersome to measure. Measured $L(k)$ is filtered with exponential weighted average, similar to θ_a in Equation (1).

The evolution of the output variable is described by a first-order system model in the form of a difference equation

$$y(k+1) = ay(k) + bu(k) \quad (2)$$

The disturbance $d(k)$ is the effect on measured output L of factors considered external to the system model. These factors are channel data rate change, WLAN frame errors or any other changes in WLAN protocol capacity. The disturbance need not be measurable for feedback control.

Using linear control for WLANs presents two problems:

- 1) *Knowledge of reference value*: There is no direct performance requirement for L as a measured output. Setting the reference for L to zero minimizes the response time, but at the cost of WLAN underutilization.
- 2) *System nonlinearity*: Two sources of nonlinearity, especially near the saturation point, are present in WLANs: queue length as an exponential function of station load and the dependence of service time on the total WLAN load.

This non-linear system can be approximated by a linear model with reasonable accuracy if it is linearized around the *operating point (OP)* that indicates the desired steady state values of $L(k)$ and $\theta(k)$, denoted as \bar{L} and $\bar{\theta}$ respectively. Therefore, we use

$$y(k) = L(k) - \bar{L} \quad \text{and} \quad u(k) = \theta(k) - \bar{\theta}.$$

The choice of the OP is crucial for non-linear systems. We select the OP based on empirical data for a specific scenario. If the system curve shifts due to disturbance, the OP and the system model needs to be re-evaluated.




Fig. 2: L vs. θ from exploratory experiments

B. System Identification

We construct system models (Equation 2) using empirical data from WLAN simulations with threshold based CAC for different scenarios. A *scenario* defines one set of WLAN parameters such as data rate, number of stations and specifications of flows generated at all stations. The WLAN scenarios considered in this paper correspond to an infrastructure WLAN where the access point (AP) has downlink flows to all other stations and each non-AP station has one uplink flow. All flows have Poisson packet arrivals which helps us test our method in presence of high variability although the method works with any arrival process. The time interval T_s for controller action is 7s, based on the same considerations discussed in Section II-A. The simulation parameters used for identification and FBC performance analysis are listed in Table I.

We first conducted a set of *exploratory experiments* to identify the operating region. A plot of steady-state values of L against a range of θ s helps us visually place the OP sufficiently away from the knee-point of the curve. For example, referring to Fig. 2, the OPs for 100B and 500B packet size are placed at threshold values of 0.67 and 0.7 respectively.

System model parameters are obtained using least squares fitting on the empirical data which is obtained by varying the θ in a sinusoid pattern in the operating region. Table II presents results of empirical identification for different scenarios.

IV. CONTROLLER DESIGN AND PERFORMANCE ANALYSIS

We use pole-placement method [9] to design different controllers. By iterating over design choices, we choose the controller that provides an acceptable performance of CAC, specifically the one that avoids packet drops due to buffer overflow at the IFQ. We evaluate the performance of FBC by simulating scenarios identical to those used for empirical identification, but with FBCs switched on.

A. Proportional Controller

Proportional control (P-control) [9] provides a fast response to error in output y . The proportional control law is

$$u(k) = Ke(k), \quad e(k) = -y(k) \quad (3)$$

Since stochastic noise in queue length causes unwanted variations in threshold value θ , we filter the θ values computed by the controller with exponential weighted average ($\alpha_\theta = 0.6$) before feeding them to the ACU. Further, the θ values are clipped to a suitable ($\theta_{\min}, \theta_{\max}$) range to protect the CAC from extreme θ values.

	Throughput (Mbps)	Delay Avg.(ms)	Delay 95% (ms)	% Data dropped
Fixed-thresh	2.93	73.6	244.1	4.18
P-control	2.8	18.5	100.1	0.06
PI control	2.66	8.26	37.5	0.004

TABLE III: Delivered CAC performance under disturbance

The designed poles for P-controller are around 0.71 with predicted settling time of 12 (84 Sec.). To test FBC performance with sustained disturbance, we introduce disturbance in the form of reduction in WLAN protocol capacity in four steps, starting at 150s, 200s, 250s and 300s, and lasting up to 500s. Other parameters of a scenario remain unchanged. The plots of Admitted load, queue length (L), and threshold value (θ) for one simulation run of Scenario-2 (Table II) are presented in Fig. 3. The disturbance (capacity reduction) for this scenario is in four steps of 4.5% each. As expected, the disturbance causes L to increase, which the FBC reacts to, by lowering θ . Nevertheless, there are considerable number of packets that are dropped. To improve this, we explore the option of proportional-integral controller.

B. Proportional-Integral Controller

We consider the Proportional-Integral (PI) control [9] that provides better steady-state response to disturbance. PI control has the control law

$$u(k) = u(k-1) + (K_P + K_I)e(k) - K_P e(k-1) \quad (4)$$

where K_P and K_I are known as proportional and integral gain respectively. As the PI controller has inbuilt filtering capability, smoothing using α_θ is not required and hence not used.

The designed poles for PI controller are around 0.81 giving a predicted settling time of 19 (133 Sec.). Significant oscillations in threshold value are observed with PI control. Figure 3 compares the performance of PI controller with P-controller and fixed-threshold scheme for Scenario-2 in presence of identical disturbance. Corresponding throughput, average delay, 95 percentile delay, and percentage data drop (due to buffer overflow at IFQ), obtained from five independent simulation runs, are compared in Table III. Both P-controller and PI controller achieve significantly better performance than a fixed threshold based CAC, even in presence of large disturbance.

V. ADAPTIVE CONTROL

A practical limitation of scenario-specific controllers is that parameters defining a scenario, i.e., average packet size, channel data rate and number of stations, keep varying and result in a large number of possible WLAN scenarios. We observe that FBC performance is more sensitive to the choice of OPs and less sensitive to the controller gains. Therefore we derive an adaptive controller with fixed controller gains, but with OPs determined *online* as per the WLAN scenario.

1) *Predicting the OP of L*: A study of Table II suggests that \bar{L} varies between 1600 to 3000 and depends only on average packet length. As a simple heuristic we set $\bar{L} = 1600 + Z_a$ where Z_a is the average packet length for the station's traffic.

Smoothing parameter α	0.85
ACU measurement interval	0.5s
Flow termination interval	1.7s
Flow arrival interval	2s
Feedback interval T_s	7s
Interface Queue (IFQ) size	30KB
Number of stations n	10

TABLE I: Simulation parameters

Scenario	Data Rate	Packet Size	a	b	Emp. \bar{L}	Predicted \bar{L}	Emp. $\bar{\theta}$	Predicted $\bar{\theta}$
1	11	100	0.63	3881	1600	1700	0.67	0.69
2	11	500	0.74	3947	2200	2100	0.7	0.73
3	11	1000	0.72	3445	2500	2600	0.72	0.76
4	11	1500	0.61	4652	3000	3100	0.72	0.77
5	2	500	0.67	3309	2200	2100	0.84	0.79

TABLE II: System Identification




Fig. 3: FBC performance with disturbance (150-500 Sec) and without

2) *Predicting the OP of θ* : To predict $\bar{\theta}$, we first use a simple analytical model to estimate the peak channel utilization ρ_{\max} . Given that T_{data} , T_{ack} and T_{cycle} are data, ACK and inter-transmission durations respectively,

$$\rho_{\max} = (T_{\text{data}} + T_{\text{ack}})/T_{\text{cycle}}$$

$$T_{\text{cycle}} = T_{\text{data}} + SIFS + T_{\text{ack}} + DIFS + BO_{\text{slots}} \cdot \sigma$$

$$BO_{\text{slots}} = CW_{\text{MIN}}/(n + 1)$$

T_{cycle} is obtained using an approach similar to that used by Tay and Chua [10], but ignoring the frame collisions. SIFS, DIFS, and σ are durations defined by the 802.11 PHY. CW_{MIN} is the minimum backoff window and n is the number of active stations. T_{data} depends on packet size and channel data rate for all frame transmissions on the WLAN. Similar to BUFFET [4], T_{data} is measured from the wireless channel.

The maximum utilization ρ_{\max} occurs at WLAN saturation. To set the operating point $\bar{\theta}$ sufficiently away from saturation, we set $\bar{\theta} = (\rho_{\max} \cdot C_{\theta})$. The value of 0.83 for the constant C_{θ} , obtained from empirical data, seems to provide accurate $\bar{\theta}$ predictions. Any small inaccuracy in $\bar{\theta}$ prediction is compensated by the disturbance rejection property of PI controller.

Table II compares predicted OPs with the OPs from empirical identification. It is seen that the OP prediction is accurate, the maximum prediction error being 6% for \bar{L} and 7% for $\bar{\theta}$. The OP prediction combined with the P or PI controllers provides an adaptive control scheme that works across WLAN scenarios without the need to redesign the controllers.

VI. CONCLUSION

WLAN CAC based on a channel utilization threshold is an appealing mechanism but the threshold needs to be determined carefully based on the variable protocol capacity. In this paper, we proposed a control-theoretic approach to determine this threshold. We formulated the WLAN interface along with

the admission controller as a controlled system, whose output measure is the interface queue length and the control input is the utilization threshold. Prototype implementation of the proposed P-controller and PI-controller in OPNET demonstrated a substantial improvement over a fixed-threshold scheme.

We also developed an adaptive controller that is able to work across WLAN scenarios by predicting operating points through heuristics and a simple model of channel utilization. With a properly configured adaptive controller, this FBC solution presents a self-regulated WLAN setup that relies only on measurements of performance metrics and traffic parameters.

The adaptive controller proposed in this paper only adapts the operating points online. Online estimation of model parameters using analytical model could provide faster and more robust control.

REFERENCES

- [1] C. T. Chou, S. Shankar, and K. Shin, "Achieving per-stream QoS with distributed airtime allocation and admission control in IEEE 802.11e wireless LANs," in *Proc. of INFOCOM*, vol. 3, 2005, pp. 1584–1595.
- [2] H. Zhai, X. Chen, and Y. Fang, "A call admission and rate control scheme for multimedia support over IEEE 802.11 wireless LANs," in *Proc. of ICST QSHINE*, 2004, pp. 76–83.
- [3] A. Bazzi, M. Diolaiti, and G. Pasolini, "Measurement based call admission control strategies in infrastructure IEEE 802.11 WLANs," in *Proc. of PIMRC*, vol. 3, 2005, pp. 2093–2098.
- [4] P. Patil and V. Apte, "An autonomous distributed admission control scheme for IEEE 802.11 DCF," in *Proc. of ICST QSHINE*, 2007.
- [5] E. Park, D. Kim, H. Kim, and C. Choi, "A cross-layer approach for per-station fairness in TCP over WLANs," *IEEE Trans. on Mobile Computing*, vol. 7, no. 7, pp. 898–911, July 2008.
- [6] G. Boggia, P. Camarda, L. Grieco, and S. Mascolo, "Feedback-based control for providing real-time services with the 802.11e MAC," *IEEE/ACM Trans. on Networking*, vol. 15, no. 2, pp. 323–333, 2007.
- [7] "OPNET Technologies, Inc," <http://www.opnet.com>.
- [8] K. Moudgalya, *Digital Control*. Chichester, UK: John Wiley & Sons, 2007.
- [9] J. Hellerstein, Y. Diao, S. Parekh, and D. Tilbury, *Feedback Control of Computing Systems*. Wiley Interscience, 2004.
- [10] Y. C. Tay and K. C. Chua, "A capacity analysis for the IEEE 802.11 MAC protocol," *Wireless Networks*, vol. 7, no. 2, pp. 159–171, 2001.