2005-01-01

A Novel Framework for Radio Resource Management in IEEE 802.11 Wireless LANs

Mark Davis
Dublin Institute of Technology, mark.davis@dit.ie

Tristan Raimondi
Dublin Institute of Technology

Recommended Citation

Davis, M. & Raimondi, T. (2005) A novel framework for radio resource management in IEEE 802.11 Wireless LANs. *International Symposium on Modeling and Optimization in Mobile, Ad Hoc, and Wireless Networks (WiOpt'05),* Riva del Garda, Trentino, Italy, April, 2005.

A Novel Framework for Radio Resource Management in IEEE 802.11 Wireless LANs

Mark Davis and Tristan Raimondi
Communications Network Research Institute
School of Electronic and Communications Engineering
Dublin Institute of Technology, IRELAND
mark.davis@dit.ie, traimondi@cnri.dit.ie

Abstract

In this paper we address the need to characterize and quantify resource usage in IEEE 802.11 WLANs in order to support radio resource management. We present a compact and intuitive framework for performance characterization and resource utilization that is based upon the concept of MAC bandwidth components. These MAC bandwidth components are directly related to the transmission rate and serve to quantify the resource requirements associated with accessing the wireless medium. We also introduce a graphical technique for presenting these MAC bandwidth components that illustrates how WLAN stations interact in contending for access to the wireless medium. We demonstrate the usefulness of this framework for radio resource management using a number of computer simulations based upon the emerging IEEE 802.11e QoS standard.

1. Introduction

In recent years there has been an explosive growth in the use of wireless LANs (WLANs) arising from the advent of the IEEE 802.11b (or Wi-Fi) standard. To date WLANs have been deployed primarily as a wireless extension to Ethernet networks and as such are suited to best effort services such as email and Internet access. However, a number of new multimedia applications such as voice over IP (VoIP) and video streaming have emerged that impose stringent requirements on network performance in order to ensure that users experience an acceptable quality of service (QoS). These new applications can be characterized by their real-time nature which requires that their data packets be delivered within strict time bounds. Specifically, these time-bounded services impose upper limits on the delay and jitter in addition to the usual performance metrics of throughput and packet loss.

The original IEEE 802.11-1999 standard [1] specifies two channel access mechanisms: A mandatory contention-based distributed coordination function (DCF) and an optional polling-based point coordination function (PCF) that has been largely ignored by the major equipment manufacturers. DCF provides a best effort service and is not capable of providing differentiation and prioritisation based upon traffic type. Consequently, neither DCF (nor PCF) has sufficient functionality to provide the QoS demanded by multimedia applications [2].

This shortcoming in the IEEE 802.11 standard is currently being addressed by the IEEE 802.11 Task Group E which is proposing a number of enhancements to the standard. The IEEE 802.11e draft [3] defines a superset of features specified in the original standard and introduces the Hybrid Coordination Function (HCF) that has two modes of operation: Enhanced Distributed Coordinated Access (EDCA) and HCF Controlled Channel Access (HCCA). The contention-based EDCA is an extension to DCF and provides for service differentiation through prioritized access to the wireless medium. Prioritization is realized through the introduction of four Access Categories (*ACs*) each with its own transmit queue and set of *AC* parameters. The differentiation in priority between *AC*s is realized by setting different values for the *AC* parameters which include the arbitration interframe spacing (*AIFS*) and minimum contention window size (CW_{min}). With proper tuning of these parameters, traffic performance can be optimized [4].

2. Radio Resource Management

It should be borne in mind that 802.11e is only a QoS *enabling* mechanism that requires some higher level management functionality in order to deliver QoS guarantees. Typically, some form of radio resource management (RRM) is required to allocate the available resources among the contending stations (STAs) in accordance with their QoS requirements and respective priorities. A critical requirement for a successful RRM scheme is the ability to accurately characterize and quantify the resource usage of the wireless medium on a per STA basis. This is a far from trivial task in 802.11 WLANs owing to the nature of the channel access protocol employed which causes the operation (and hence the resource usage) of individual STAs to become coupled. Under 802.11e operation, this situation is further complicated as the STAs contend with different sets of *AC* parameters.

In this paper we present a novel framework for characterizing the per-STA performance and resource usage of a WLAN in a compact and intuitive format. We introduce the concept of MAC bandwidth components that serve to quantify the resource usage associated with each phase of a STA's operation in contending for access to the wireless medium. The MAC bandwidth components are directly related to the transmission rate which allows for an intuitive interpretation of resource usage. Moreover, this approach captures the nature of the coupling between contending STAs and allows the resource usage of the WLAN to be described by a set of simple coupled equations.

We also introduce the MAC operating plane as a useful graphical description of WLAN operation that clearly illustrates the coupling between contending STAs. We define an access efficiency that represents the "cost" to a STA in accessing the wireless medium and we show how the access efficiency may be controlled through the *AC* parameters thereby differentiating between the STAs. We suggest how this may form the basis of an 802.11e RRM scheme where WLAN resources are allocated among the contending STAs through controlling the cost of access to the wireless medium.

3. IEEE 802.11 MAC Mechanism

The basic access scheme in 802.11 WLANs is the DCF used to support asynchronous data transfer on a best effort basis where all stations (STAs) must contend with each other to access the medium in order to transmit their data. The DCF employs a medium access control (MAC) technique known as carrier sense multiple access with collision avoidance (CSMA/CA). CSMA/CA is a "listen-before-talk" access protocol where any STA wishing to transmit a frame first invokes the carrier sense mechanism to determine the busy/idle state of the medium. If the medium is busy, the STA defers its transmission until the medium is determined to be idle without interruption for a period of time equal to *DIFS* (in 802.11b *DIFS* = 50 μs). As part of the collision avoidance mechanism, the 802.11 MAC requires STAs to delay their transmission for an additional random Backoff Interval after the medium becomes idle. The Backoff Interval is used to initialize the Backoff Timer. The Backoff Timer is decreased as long as the medium remains idle, stopped when the medium is sensed busy, and reactivated when the medium is sensed idle again for longer than *DIFS*. A STA may transmit its frame when its Backoff Timer reaches zero. The backoff time is slotted (in 802.11b *Slot_Time* = 20 μs) and a STA is only allowed to transmit at the beginning of a time slot. The Backoff Interval is randomly generated using *Backoff Interval* = $BC \times Slot\_Time$ where *BC* is a pseudorandom integer drawn from a uniform distribution over the interval [0,*CW*] and where *CW* is the Contention Window.

The effect of this procedure is that when multiple STAs are deferring and go into random backoff, the STA selecting the smallest Backoff Interval will win the contention. Fairness is promoted as each STA must recontend for access after every transmission. Occasionally, two or more STAs may choose the same *BC* value leading to a collision as the STAs involved will transmit their frames at the same time. To resolve collisions, an exponential backoff scheme is used whereby the size of the *CW* is doubled after each unsuccessful transmission.

4. MAC Bandwidth Components

From the description of the basic access mechanism above, it is possible to distinguish a number of different time intervals on the wireless medium, see Figure 1. Firstly, there are the intervals during which the medium is *busy* corresponding to the transmission of frames and their positive acknowledgments (in the case of data and management frames). This busy time on the medium is associated with the transport of the traffic load. The complementary time intervals are the *idle* intervals. A STA can make use of these idle intervals in a number of ways. If the STA has a data or management frame awaiting transmission, it uses the idle time on the medium to allow *DIFS* and *Slot_Time* intervals to elapse.

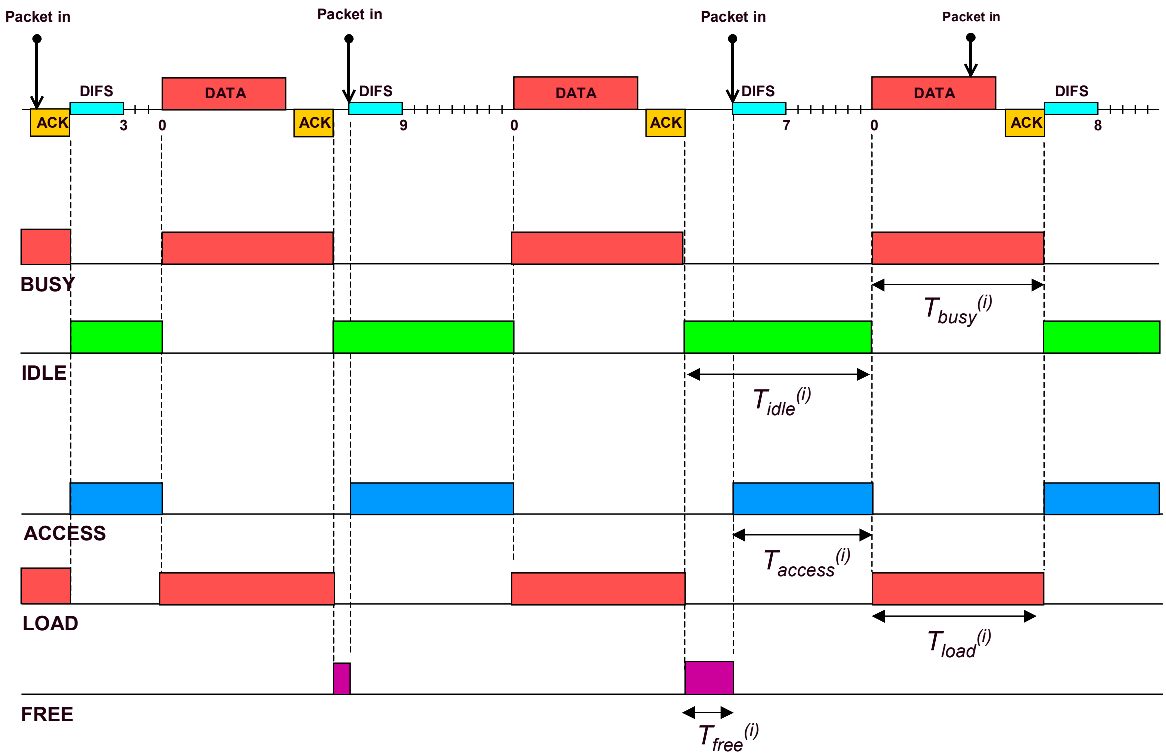

Figure 1. The various time intervals involved in accessing the medium.

This portion of the medium idle time corresponds to the time spent by a STA in contending for access to the medium. If the STA does not have a frame to transmit, the idle time is not being used and is therefore considered to be free in the sense that it is available, if required, to the STA. This free time on the medium can be viewed as spare capacity on the medium, essentially acting as a reservoir that can be drawn on when required. The amount of free time experienced by a STA is related to the level of QoS experienced by its traffic load where the greater the free bandwidth available to a STA, the better the QoS likely to be experienced. The busy and idle time intervals are summed (over some measurement interval of interest) as follows:

$$T_{busy} = \sum_i T_{busy}^{(i)} \text{ and } T_{idle} = \sum_i T_{idle}^{(i)} \qquad \textbf{(1)}$$

where $T_{busy}^{(i)}$ and $T_{idle}^{(i)}$ are the durations of the i^{th} busy and idle intervals respectively within the measurement interval of interest. A more useful and meaningful description of these quantities is to first normalize them and then convert them to a bandwidth related to the transmission rate (*TX_rate*) as follows:

$$BW_{busy} = \frac{T_{busy}}{T_{busy} + T_{idle}} \times TX\_rate$$

and

$$BW_{idle} = \frac{T_{idle}}{T_{busy} + T_{idle}} \times TX\_rate \qquad \textbf{(2)}$$

In 802.11b, *TX_rate* = 11 Mbps and obviously

$$BW_{busy} + BW_{idle} = TX\_rate \qquad \textbf{(3)}$$

Here BW_{busy} represents the portion of the transmission rate used in the transport of the total traffic load. Similarly, BW_{idle} represents the portion of the transmission rate that is idle and may be used by any STA to win access opportunities for its load. Associating the transmission of a frame with a particular STA leads to the concept of the load bandwidth $BW_{load}(k)$ which corresponds to that portion of the transmission rate used in transporting its load and is directly related to the throughput of the STA. $BW_{load}(k)$ may be calculated using:

$$BW_{load}(k) = \frac{T_{load}(k)}{T_{busy} + T_{idle}} \times TX\_rate \qquad \textbf{(4)}$$

where $T_{load}(k)$ is the busy time on the medium used by a STA k in transmitting its load (and includes collisions). In the single-station case, BW_{busy} and BW_{load} will be identical. However, in the multiple station case:

$$BW_{busy} = \sum_k BW_{load}(k) - BW_{collisions} \qquad \textbf{(5)}$$

as inevitably some bandwidth ($BW_{collisions}$) will be lost due to collisions between multiple STAs attempting to transmit at the same time.

It is worth noting here that, apart from collisions, STAs do not share their load bandwidths during their transmissions. In other words, once a STA has won access to the medium, it has exclusive use of the medium for the duration of the transmission of its frame. This is in contrast to the idle bandwidth which is shared by all STAs in the sense that any STA can make use of the idle time intervals on the medium to allow periods of *DIFS* or *Slot_Time* to elapse. Furthermore, each STA perceives the idle bandwidth as comprising two components, an access bandwidth BW_{access} used to contend for access opportunities and a free bandwidth BW_{free} corresponding to the remaining unused idle bandwidth, i.e. for any STA k the following applies:

$$\begin{aligned} BW_{access}(k) + BW_{free}(k) &= BW_{idle} \\ &= TX\_rate - BW_{busy} \end{aligned} \qquad \textbf{(6)}$$

The access time has two parts, the time spent deferring and the time spent backing off. Depending on the particular traffic conditions prevailing on a WLAN, a STA may experience several cycles of deferral (i.e. waiting *DIFS*) and backoff (i.e. decreasing its Backoff Timer) before being allowed to transmit its frame. The actual number of times a STA has to defer will depend on a number of factors, including the number of STAs currently contending for access, its own initial Backoff Interval, as well as those of all the other contending STAs.

In summary, this framework for WLAN resource usage defines three MAC bandwidth components that are coupled via equations (3), (5), and (6). These MAC bandwidth components give a compact and intuitive description of resource usage by the 802.11 MAC mechanism that is particularly suited to supporting radio resource management schemes. For example, it can be used to give an advanced warning of the on-set of saturation. The on-set of STA saturation occurs when its $BW_{free}(k)$ has been reduced to zero or when $BW_{access}(k) = BW_{idle}$, i.e. all of the idle bandwidth is being used by the STA in accessing the medium in order to service its offered load.

5. MAC Bandwidth Operating Plane

The *MAC Bandwidth Operating Plane* is essentially an extension of the MAC bandwidth components concept whereby an operating plane is formed in terms of the load and access bandwidth components, see Figure 2. A STA's operating point can be characterised by its position in this plane specified by its (BW_{load}, BW_{access}) components. The operating point of the WLAN can also be represented in this plane in terms of the (BW_{busy}, BW_{idle}) values. However, owing to the requirement given by (3), the WLAN operating point is constrained to lie along a line. This restriction does not apply to the STAs whose operating points ($BW_{load}(k)$, $BW_{access}(k)$) may lie anywhere within the (shaded) region bounded by BW_{busy} and BW_{idle}. The $BW_{free}(k)$ component may also be visualised in terms of the distance of the STA's operating point from the BW_{idle} boundary. This diagram also indicates the efficiency with which a STA is accessing the medium where the access efficiency η_a is defined as:

$$\eta_a = \frac{BW_{load}}{BW_{access}} \qquad \textbf{(7)}$$

An access efficiency angle θ_a may also be defined as

$$\theta_a = \tan^{-1}\left(\frac{BW_{load}}{BW_{access}}\right) \qquad \textbf{(8)}$$

The larger the access efficiency angle θ_a the more efficiently the STA is accessing the medium. For example, in Figure 2, STA_2 is more efficient than STA_1 in accessing the medium. Moreover despite having the larger load, it also has the larger BW_{free} owing its greater access efficiency.

As a consequence of the requirement for a STA to defer its transmission if the medium is busy, its BW_{access} requirement will depend on the load conditions of the other STAs in the WLAN. The result of this deferral feature of MAC operation is that the operating point of a STA will vary as the load presented to the WLAN varies. Therefore the impact of changes in the load of

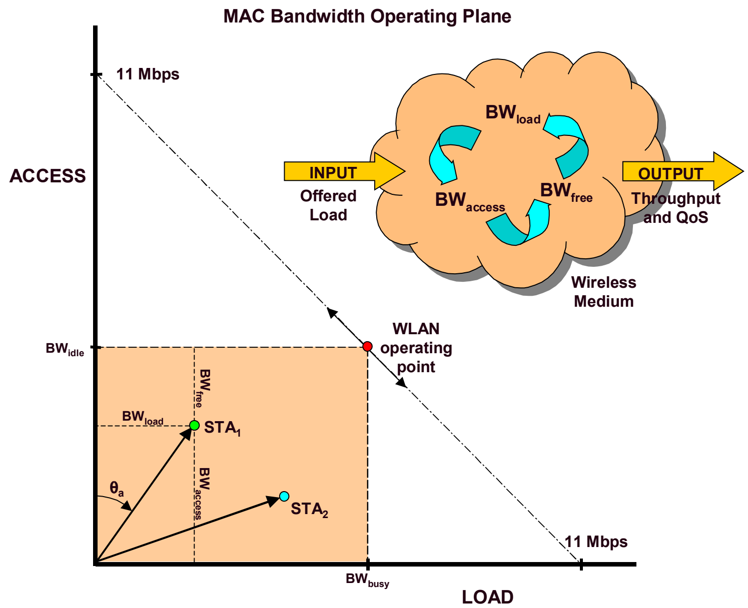

Figure 2: The MAC bandwidth operating plane description

one STA on the overall performance of the WLAN can be visualised in terms of the changes in the positions of the operating points of the other STAs.

This framework for resource utilisation based around the concept of MAC bandwidth components has been implemented in a WLAN traffic probe that is described in [5] where the results from a number of real-time traffic streaming scenarios are also presented.

6. Results

A computer simulator has been developed in C/C++ which implements the 802.11e EDCA MAC mechanism [4]. We use the 802.11b DSSS PHY standard operating at the maximum transmission rate of 11 Mbps to simulate the wireless medium. We do not consider other traffic control features such as EDCA-TXOP and the No ACK/Block ACK policy in the simulator. Any STA gaining access to the medium transmits one frame and then releases the channel to the next successful STA. We also neglect high-level management functionality such as beacon frames, association and authentication frames exchanges. In the following test scenarios, two STAs (STA_1 and STA_2) are contending for access. In the case of STA_1, its offered load comprises a Poisson traffic stream of 512 bytes packets with a mean rate of 500 pps. STA_2's offered load is also a 512 byte Poisson traffic stream, but with a mean rate that is ramped from 50 pps to 1000 pps in steps of 50 pps. The *AIFS* duration is derived from the arbitration frame spacing number *AIFSN* using $AIFS = AIFSN \times Slot\_Time + SIFS$ where $SIFS = 10$ µs in 802.11b. Three test scenarios are now considered.

6.1 Test Scenario 1: 802.11b Operation

In this first test scenario, 802.11b operation is considered where $AIFSN = 2$ (corresponding to 50 µs or *DIFS*) and $CW_{min} = 31$ for both STAs, i.e. there is no differentiation between STAs in terms of their priorities in contending for access. Figure 3 shows that as the offered load to STA_2 increases, its BW_{load} and BW_{access} requirement also increases before saturating at approximately 500 pps at which point its BW_{free} component has been reduced to zero. In the case of STA_1, its BW_{load} remains constant until STA_2's offered load exceeds 250 pps. Beyond this point there is no longer sufficient available capacity in the WLAN to meet the resource demands of both STA_1 and STA_2 and STA_1 is forced to give up some of its BW_{load} (and hence throughput) to support STA_2. The slight rise in BW_{access} for STA_1 is due to the increased number of transmission deferrals it experiences due to the increased contention for access opportunities from STA_2 increasing load.

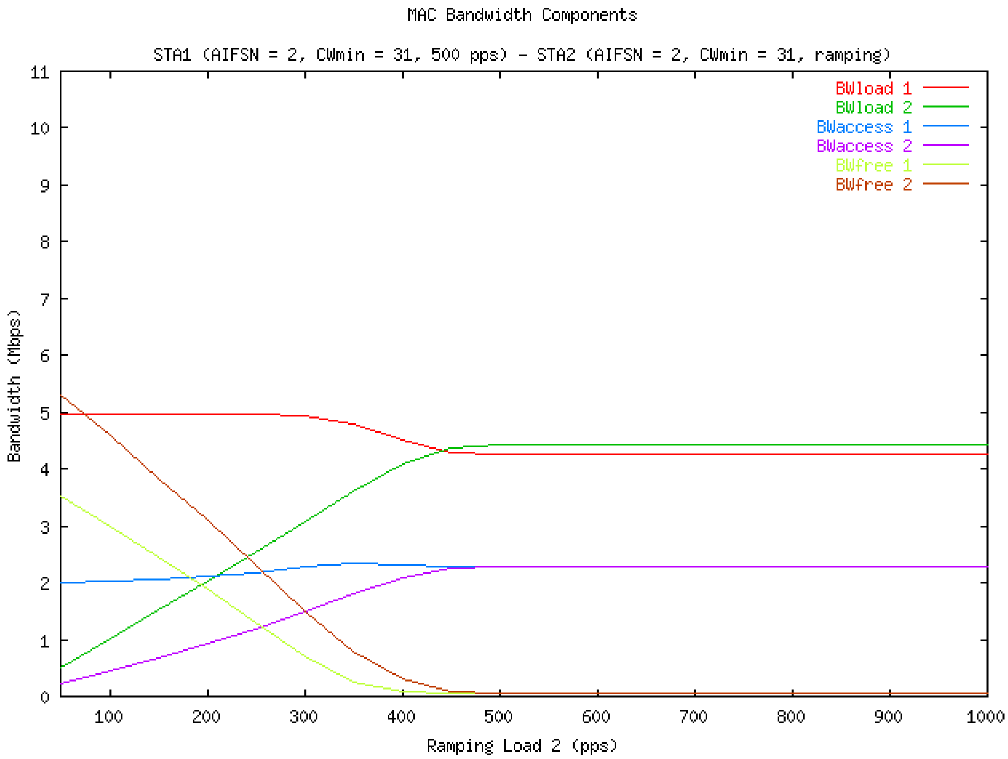

Figure 3: The MAC bandwidth components for test scenario 1 (802.11b operation).

Figure 4 presents these same results in terms of the MAC bandwidth operating plane description where the impact of the STA_2's rising load on the performance of STA_1 can be clearly observed. Initially, the impact is slight resulting in a small increase in its BW_{access} requirement, i.e. a slight reduction in its access efficiency owing to the increased number of deferrals it must undergo. However, when STA_1 is forced to give up bandwidth to STA_2, a sharp turn is observed in the characteristic denoting the reduction in its BW_{load}. This change induced in the position of STA_1's operating point manifests itself as a reduction in its access efficiency. On the other hand, STA_2's characteristic shows that its access efficiency remains constant at $\eta_a \approx 2$. This result clearly shows the interaction between contending STAs under 802.11b operation where the MAC mechanism attempts to share the access opportunities equally between STAs (i.e. both STAs enjoy the same access priority).

The change in the operating point of the WLAN can also be seen in this figure where STA_2's increasing load results in an increasing BW_{busy} and hence a decreasing BW_{idle} for the network as a whole.

6.2 Test Scenario 2: 802.11e Operation-Varying *AIFSN*

In the next test scenario 802.11e operation is considered where differentiation between the STAs is introduced through changes to the *AIFSN* parameter. For STA_1 *AIFSN* = 2, while for STA_2 its *AIFSN* parameter is increased from 2 to 10 in order to reduce its access priority relative to STA_1, both STAs use *CWmin* = 31. Figure 5 shows that the effect of increasing *AIFSN* on STA_2 is to reduce its access efficiency thereby making it increasingly expensive for STA_2 to support its load resulting in a reduction in its saturation BW_{load} value. In addition, the impact of STA_2's rising load on STA_1 is reduced as *AIFSN* for STA_2 is increased. This can be seen from the reduction in the displacement of STA_1's operating point. This result illustrates how the *AIFSN* parameter can be used to introduce differentiation between STAs in terms of their access efficiencies. This feature could be employed in a RRM scheme to manage the resources of an 802.11e WLAN through control of the "cost" to a STA in winning access opportunities to the wireless medium for its offered load.

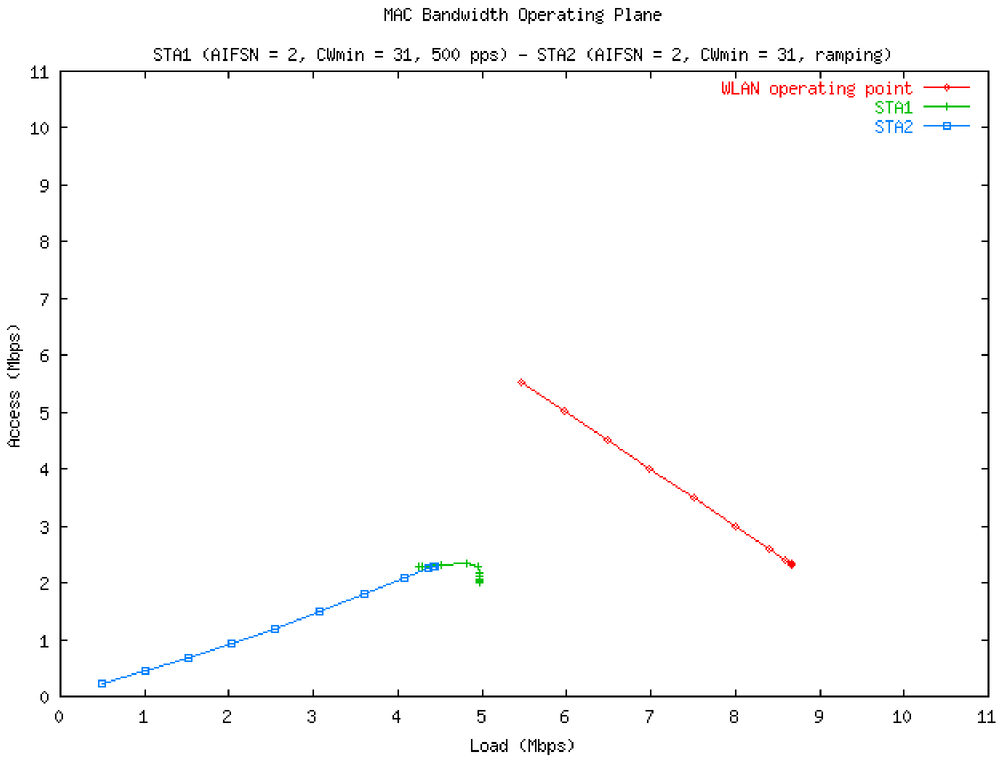

Figure 4: The MAC operating plane description for test scenario 1 (802.11 operation).

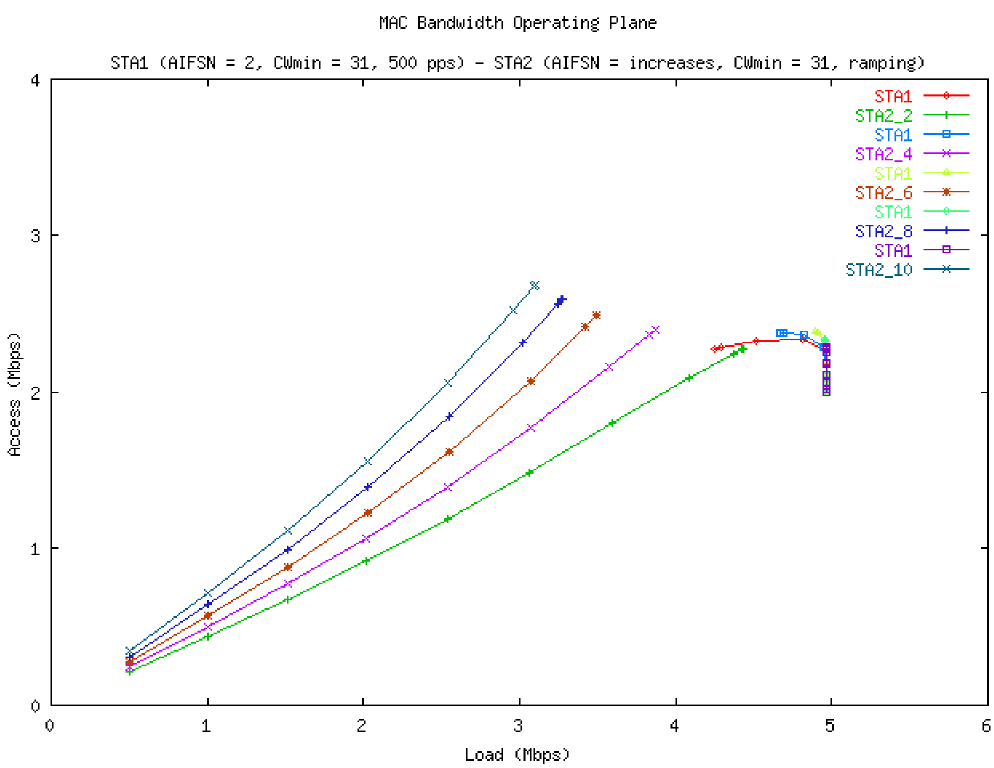

Figure 5: The MAC operating plane description for test scenario 2 (802.11e operation, varying *AIFSN*).

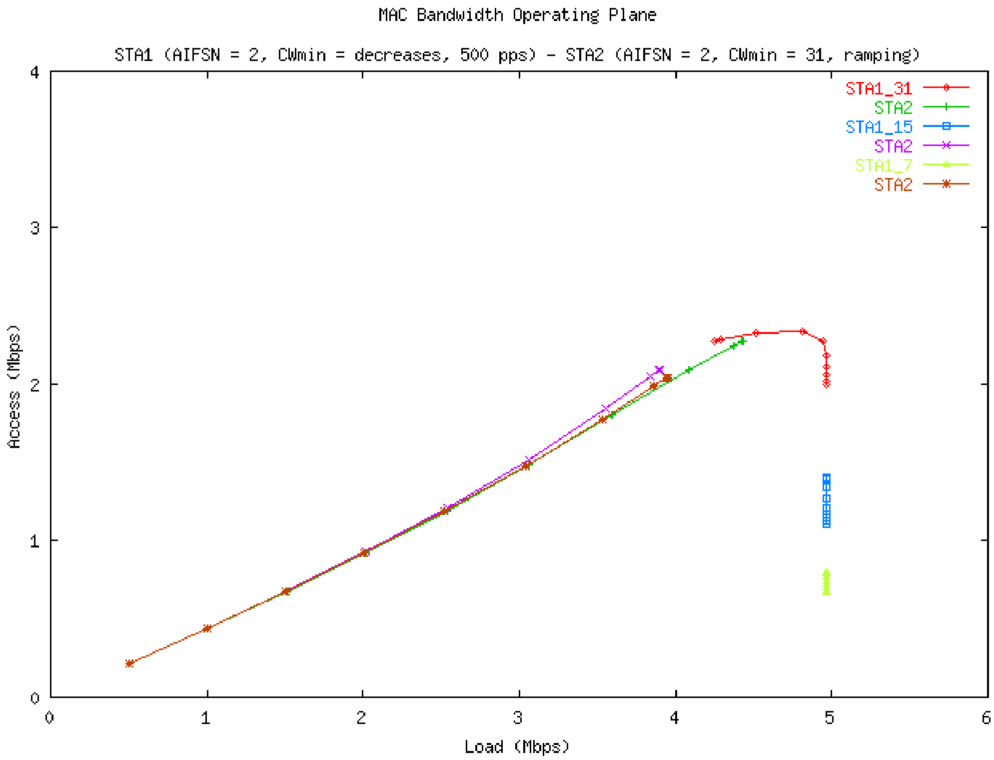

Figure 6: The MAC operating plane description for test scenario 3 (802.11e operation, varying CW_{min}).

6.3 Test Scenario 3: 802.11e Operation-Varying CW_{min}

In the final test scenario, we discriminate against STA_2 by decreasing the CW_{min} parameter of STA_1 from 31 to 15 to 7, both STAs use $AIFSN$ = 2. Figure 6 shows that there is a significant reduction in the impact of STA_2's rising load on the operation of STA_1. Moreover, there is an appreciable increase in the access efficiency for STA_1 making it less expensive (in terms of the access bandwidth requirement) to support its offered load. This illustrates how it is possible to differentiate between STAs in terms of their access efficiencies through control of the CW_{min} parameters. Again this feature could be employed in a RRM scheme to control the allocation of the WLAN resources among the competing STAs.

7. Summary

We have presented a framework for resource utilization in 802.11 WLANs that is based upon the concept of MAC bandwidth components. These MAC bandwidth components are directly related to the transmission rate and serve to quantify the resource requirements associated with accessing the wireless medium. Moreover, this approach captures the nature of the contention between STAs competing for the finite resources of the WLAN and allows for a simple model of resource usage based upon a set of coupled equations. By presenting the MAC bandwidth components in the graphical format of an operating plane, we realize a characterization of WLAN resource usage that is both compact and intuitive.

The usefulness of this approach has been demonstrated through a number of computer simulations involving the original 802.11 MAC standard and the emerging 802.11e QoS MAC standard. Presenting the results from the simulations using this framework illustrates how STAs interact in competing for resources and demonstrates how STA differentiation can be controlled through the *AC* parameters.

8. Acknowledgements

The support of the Informatics Research Initiative of Enterprise Ireland and of Science Foundation Ireland under the National Development Plan is gratefully acknowledged.

9. References:

[1] IEEE, "Part 11, Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications," *IEEE Std 802.11*, 1999 Edition.

[2] H. Zhu et al., "A Survey of Quality of Service in IEEE 802.11 Networks," *IEEE Wireless Communications*, August 2004, pp. 6-14.

[3] IEEE, "Part 11, Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications: Medium, Access control (MAC) Enhancements for Quality of Service (QoS)," *IEEE Std 802.11e/D8.0*, Feb. 2004.

[4] T. Raimondi and M. Davis, "Design Rules for a Class-based Differentiated Service QoS Scheme in IEEE 802.11e Wireless LANs," *Proc. 7th ACM MSWiM 2004*, Venezia, Italy, Oct. 2004.

[5] M. Davis, "A Wireless Traffic Probe for Radio Resource Management and QoS Provisioning in IEEE 802.11 WLANs," *Proc. 7th ACM MSWiM 2004*, Venezia, Italy, Oct. 2004.