

2004-01-01

A Wireless Traffic Probe for Radio Resource Management and QoS Provisioning in IEEE 802.11 WLANs

Mark Davis

Technological University Dublin, mark.davis@tudublin.ie

Follow this and additional works at: <https://arrow.tudublin.ie/commcon>

 Part of the [Electrical and Electronics Commons](#)

Recommended Citation

Davis, Mark (2005) A wireless traffic probe for radio resource management and QoS provisioning in IEEE 802.11 WLANs. *ACM Symposium on Modeling, Analysis and Simulation of Wireless and Mobile Systems (MSWiM'04)*, Venezia, Italy, October, 2004.

This Conference Paper is brought to you for free and open access by the Communications Network Research Institute at ARROW@TU Dublin. It has been accepted for inclusion in Conference papers by an authorized administrator of ARROW@TU Dublin. For more information, please contact arrow.admin@tudublin.ie, aisling.coyne@tudublin.ie, gerard.connolly@tudublin.ie, vera.kilshaw@tudublin.ie.

A Wireless Traffic Probe* for Radio Resource Management and QoS Provisioning in IEEE 802.11 WLANs

Mark Davis

Communications Network Research Institute

School of Electronic and Communications Engineering

Dublin Institute of Technology, IRELAND.

mark.davis@dit.ie

ABSTRACT

The emergence of real-time services such as voice over IP (VoIP) and video streaming, imposes stringent requirements on the performance of a network if quality of service (QoS) targets are to be achieved. In the case of wireless networks, some form of radio resource management (RRM) is typically required to allocate the available resources among the contending stations in accordance with their needs and respective priorities. A critical aspect of any RRM scheme is the ability to monitor resource usage and to determine the resource requirements on a per-station basis. In this paper we describe a wireless traffic probe for IEEE 802.11 WLANs capable of obtaining this information and presenting it in a compact and intuitive format. The probe also shows how the wireless stations interact with one another in competing for the resources of the WLAN in a clear and quantifiable way. The results from the WLAN traffic probe obtained in a series of video streaming test scenarios are also presented that clearly demonstrate its usefulness and importance as a network tool in RRM and QoS provisioning schemes.

Categories and Subject Descriptors

C2.1 [Network Architecture and Design]: *Wireless Communication*

General Terms

Measurement, Performance

Keywords

Wi-Fi, Traffic Probe, Radio resource management

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 these networks have been deployed primarily as a wireless extension to Ethernet networks and as such are suited to best effort services such as Internet

access, email, and file transfers. However, this period of time has also witnessed the emergence of new network applications such as voice over IP (VoIP) and video streaming that impose stringent requirements on network performance in order to ensure that users experience an acceptable quality of service (QoS). Typically, these new network services tend to 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.

In wired networks the QoS targets for real-time data traffic can usually be met by over-provisioning, however such an approach cannot be adopted with wireless networks due to the limited availability of radio spectrum. Support for traffic with QoS requirements is currently being addressed by the IEEE 802.11e Task Group and this work is expected to be completed in late 2004. However, 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 users in accordance with their needs and respective priorities.

A critical aspect of any RRM scheme is the ability to monitor resource usage and to determine the resource requirements on a per-station basis. In this paper we describe a wireless traffic probe for IEEE 802.11 WLANs capable of obtaining this information and presenting it in a compact and intuitive format. The commercial WLAN analysers currently available are essentially protocol analysers in that they operate by extracting and processing information from the various protocol headers found in a wireless frame. None of these protocol analysers explicitly consider the operation of the medium access control (MAC) mechanism and as such cannot give the type of detailed information on radio resource usage required for RRM and QoS provisioning schemes.

The WLAN traffic probe described here operates at the MAC layer and is capable of producing (in real time) detailed and accurate information on the resource usage on a per-station basis. The probe also shows how the stations interact with one another in competing for the resources of the WLAN in a clear and quantifiable way. The results from the WLAN traffic probe obtained in a series of video streaming test scenarios are also presented that clearly demonstrate its usefulness and importance as a network tool in RRM and QoS provisioning schemes.

* The wireless traffic probe described herein is patent pending.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

MSWiM'04, October 4–6, 2004, Venezia, Italy.

Copyright 2004 ACM 1-58113-953-5/04/0010...\$5.00.

2. THE IEEE 802.11 WLAN STANDARD

The IEEE 802.11 WLAN standard [1] is by far the most popular and widely deployed wireless LAN (WLAN) technology. The original IEEE 802.11 standard was published in June 1997 and specifies the physical layer (L1/PHY) and medium access control layer (L2/MAC) for interoperable WLAN operation. The original standard operates in the unlicensed Industrial, Scientific, and Medical (ISM) band at 2.4 GHz and supports a mandatory bit rate of 1 Mbps and an optional higher rate of 2 Mbps. In September 1999 the IEEE approved the HR or “high rate” extension to the standard, known as the IEEE 802.11b, which supports data rates up to 11 Mbps.

The basic access scheme in 802.11 WLANs is the distributed coordination function (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 allows multiple STAs access the medium without the need for central control and employs a technique known as carrier sense multiple access with collision avoidance (CSMA/CA). Access priority to the medium is controlled through the use of Inter Frame Space (IFS) time intervals between the frame transmissions. The IFS intervals are mandatory periods of idle time on the medium. The 802.11 standard defines four different IFS intervals as follows:

- Short Inter Frame Space (SIFS)
- PCF Inter Frame Space (PIFS)
- DCF Inter Frame Space (DIFS)
- Extended Inter Frame Space (EIFS)

The Short IFS (SIFS) is used for the highest priority transmissions (i.e. control frames), such as ACK and RTS/CTS frames. In 802.11b, SIFS = 10 μ s. The PCF IFS (PIFS) is used by the point coordination function (PCF) during contention-free operation. STAs with data to transmit in the contention-free period can transmit after PIFS has elapsed and preempt any contention-based traffic. In 802.11b, PIFS = 30 μ s. The DCF IFS (DIFS) is the minimum idle time for contention-based (i.e. DCF) services and is used for the transmission of data and management frames. In 802.11b, DIFS = 50 μ s. The Extended IFS (EIFS) is used to recover from a failed transmission attempt. It is derived from the SIFS, DIFS, and the time required to transmit an ACK frame at the basic rate of 1 Mbps.

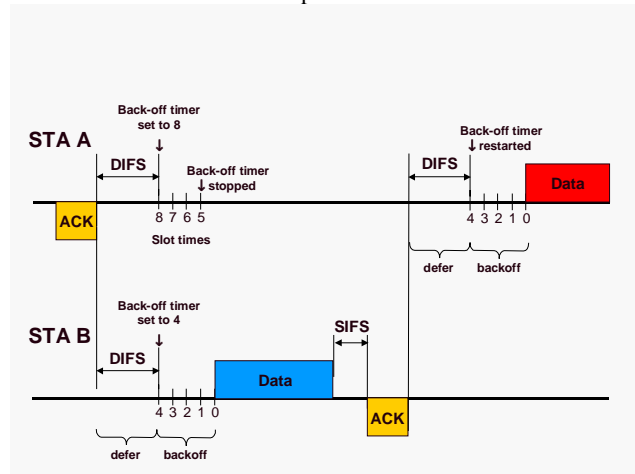


Figure 1: The 802.11 Basic Access Mechanism

CSMA/CA is a “listen-before-talk” access protocol where any STA wishing to transmit a frame first invokes its 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 (or EIFS in the case of an incorrectly received frame). 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 (or EIFS as appropriate). A STA may transmit its frame when its Backoff Timer reaches zero. The backoff time is slotted and a STA is only allowed to transmit at the beginning of a time slot. The duration of the time slot in 802.11b has been defined as Slot_Time = 20 μ s.

The Backoff interval is randomly generated using the following:

$$\text{Backoff interval} = \text{BC} \times \text{Slot_Time} \quad (1)$$

where BC = pseudorandom integer (backoff counter) drawn from a uniform distribution over the interval [0,CW] where CW is an integer referred to as 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 time will win the contention. It promotes fairness among the STAs as each STA must contend for access after every transmission. Occasionally, two or more STAs may choose the same BC value and will subsequently lead to a collision whereby the STAs involved will transmit their frames at the same time. In order to resolve collisions between STAs, an exponential Backoff scheme is adopted whereby the size of the CW is doubled after each unsuccessful transmission. It is worth noting here that the access mechanism employed by the 802.11 MAC is inherently stochastic both from the point of view of the randomly selected Backoff Interval but also from the number of times a STA may have to defer to another STA.

3. THE 802.11 MAC BANDWIDTH COMPONENTS

The origin of the 802.11 WLAN traffic probe lies in a particularly useful descriptive framework for identifying network resource usage on WLANs that is based around the concept of *MAC bandwidth components*. In particular, three MAC bandwidth components have been identified: A *load bandwidth* (BW_{load}) that is associated with the transmission of the data frames, an *access bandwidth* (BW_{access}) associated with the contention mechanism (whereby a station wins access to the wireless medium) and a *free bandwidth* (BW_{free}) that is associated with the QoS. This framework results in a compact and intuitive description of MAC resource usage that is particularly suited to radio resource management schemes.

From the description of the basic access mechanism above, it is possible to distinguish a number of different critical time intervals on the wireless medium. Firstly, there are the intervals during which the medium is *busy* corresponding to the transmission of frames and their positive acknowledgments (at least 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 (or EIFS as appropriate) and Slot_Time intervals to elapse. This portion of the medium idle time corresponds to the time spent by a STA in contending for access to the medium.

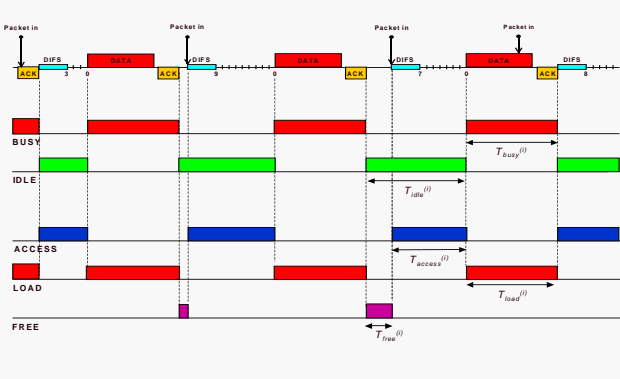


Figure 2: The various time intervals involved in accessing the medium under CSMA/CA

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 closely related to the level of QoS experienced by its traffic load where the greater the free capacity available to a STA, the better the QoS likely to be experienced.

The busy and idle time intervals may be summed (over some measurement interval of interest) as follows:

$$T_{busy} = \sum_i T_{busy}^{(i)} \quad (2)$$

$$T_{idle} = \sum_i T_{idle}^{(i)} \quad (3)$$

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 line rate as follows:

$$BW_{busy} = \frac{T_{busy}}{T_{busy} + T_{idle}} \times Line_Rate \quad (4)$$

$$BW_{idle} = \frac{T_{idle}}{T_{busy} + T_{idle}} \times Line_Rate \quad (5)$$

where obviously,

$$BW_{busy} + BW_{idle} = Line_Rate \quad (6)$$

Here BW_{busy} represents the portion of the line rate bandwidth used for the transport of the total traffic load. Similarly, BW_{idle} represents the portion of the line rate bandwidth that is idle and may be used by a STA to win access opportunities for its load. In 802.11b WLANs, the $Line_Rate = 11$ Mbps. By examining the address fields contained in the MAC header, it is possible to

identify the sender of a frame. This permits the transmission of a frame to be associated with a particular STA and leads to the concept of the load bandwidth $BW_{load}(k)$ which corresponds to that portion of the line rate bandwidth used by a particular STA k in transporting its traffic load. The load bandwidth is directly related to the throughput of the STA. The busy time on the medium used by STA k in transmitting its load is

$$T_{load}(k) = \sum_i T_{load}^{(i)}(k) \quad (7)$$

Again, this can be converted to a line rate bandwidth using

$$BW_{load}(k) = \frac{T_{load}(k)}{T_{busy} + T_{idle}} \times Line_Rate \quad (8)$$

In the single-station case, BW_{busy} and BW_{load} will be identical. However, in the multiple station case,

$$BW_{busy} = \sum_k BW_{load}(k) - BW_{collisions} \quad (9)$$

Inevitably, some bandwidth 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 the line rate bandwidth 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 its transmission. 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:

$$BW_{access}(k) + BW_{free}(k) = BW_{idle} \quad (10)$$

or

$$BW_{busy} + BW_{access}(k) + BW_{free}(k) = Line_Rate \quad (11)$$

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 or EIFS) 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 large number of factors, including the number of STAs currently contending for access, its own initial Backoff Interval, as well as the Backoff Intervals of all the other contending STAs. As both of these intervals are essentially random, it makes sense to consider the average time spent in deferring \bar{T}_{defer} and the average time spent in backing off $\bar{T}_{backoff}$. In other words, the average idle time on the medium needed by a STA in order to access the medium is

$$\bar{T}_{access} = \bar{T}_{defer} + \bar{T}_{backoff} \quad (12)$$

where

$$\bar{T}_{defer} = \bar{\#}_{defers} \times \bar{T}_{IFS} \quad (13)$$

and

$$\bar{T}_{backoff} = \bar{BC} \times Slot_Time \quad (14)$$

Here $\bar{\#}_{defers}$ is the average number of times that a STA has to defer to a busy medium condition, \bar{T}_{IFS} is the average duration of the deferral interval, and \bar{BC} is the average initial BC value. The access time T_{access} is obtained by multiplying \bar{T}_{access} by the total number of frames $frame_cnt$ transmitted within the measurement interval of interest:

$$T_{access} = \bar{T}_{access} \times frame_cnt \quad (15)$$

Obtaining T_{access} allows BW_{access} (and hence BW_{free}) to be calculated using

$$BW_{access}(k) = \frac{T_{access}(k)}{T_{busy} + T_{idle}} \times Line_Rate \quad (16)$$

It is not possible to measure the parameters $\bar{\#}_{defers}$, \bar{T}_{IFS} , and \bar{BC} directly from the medium. Instead, an indirect approach has been adopted based upon measuring the average contention experienced by a STA. The average contention is defined as the average number of STAs contending for each access. Specifically, the average contention will determine the average number of deferrals $\bar{\#}_{defers}$ and also the probability of a collision which in turn will determine \bar{T}_{IFS} , and \bar{BC} . The procedure adopted here is to compute the values of these parameters offline (through computer simulation) as a function of the average contention. A “look up table” approach is then used to obtain the values of these parameters from measured values of the average contention.

4. MEASURING THE MAC BANDWIDTH COMPONENTS

This descriptive framework of WLAN resource usage based upon MAC bandwidth components (i.e. BW_{access} , BW_{load} , and BW_{free}) forms the basis of the WLAN traffic probe. The guiding principle behind the development of the WLAN traffic probe is that it should be possible to obtain (or at least infer) all of the various traffic metrics by passively “sniffing” the wireless frames on the medium. At the heart of the WLAN traffic probe is an 802.11b WLAN card operating in the promiscuous mode. In this mode, the WLAN card is capable of reception only, but is nevertheless capable of receiving all frames transmitted on the medium. The various operations and procedures comprising the probe can be divided into three categories: Capturing the frame, processing the captured frame, and processing the information gathered over the duration of the measurement interval.

The processes concerned with capturing the frames relate to the computer platform hosting the WLAN traffic probe application. The 802.11b WLAN network interface card (NIC) usually takes the form of a PCMCIA card and is configured to operate in the promiscuous mode. The operating system used is Linux and is responsible for managing the hardware interactions between the PCMCIA WLAN card and the rest of the computer hardware. *Libpcap* provides for implementation-independent access to the underlying packet capture facility provided by Linux and the

PCMCIA WLAN card. Essentially, *libpcap* is a packet capture library that provides a high level interface to packet capture systems. Once a frame has been captured on the wireless medium. The captured frame is handed over to the WLAN traffic probe application in its entirety together with a time stamp giving the time of capture. The first task is to parse the PHY and MAC headers in order to obtain the relevant information regarding the frame. This information includes the transmission rate, the type of PLCP preamble used, the frame type and size, the sender and intended recipient of the frame, and various status flags.

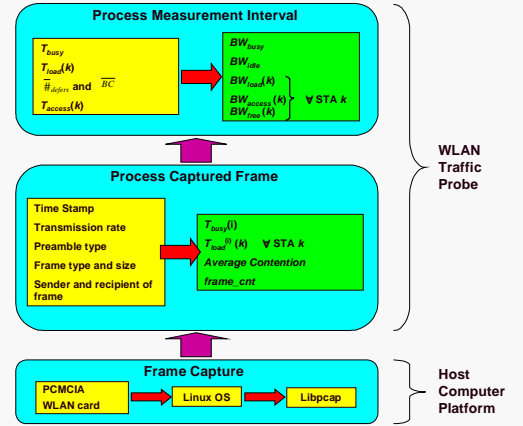


Figure 3: Structure of the 802.11 WLAN Traffic Probe

The transmission rate, preamble type, and frame size and type are used to calculate the duration of the busy interval $T_{busy(i)}$. Summing over all the busy intervals $T_{busy(i)}$ within the measurement interval gives T_{busy} which allows BW_{busy} and hence BW_{idle} to be obtained. Sorting the captured frames by their sender address allows $T_{load}^{(i)}(k)$ to be measured which when summed over all the transmissions from a particular STA k allows $T_{load}(k)$ and hence $BW_{load}(k)$ to be calculated. Similarly, the average contention may be measured. The relationship between the average contention and the $\bar{\#}_{defers}$, \bar{T}_{IFS} , and \bar{BC} parameters is computed off-line through computer simulation.

4.1 Performance Characterization Using the MAC Bandwidth Components




Figure 4: The 802.11 MAC Bandwidth Components Concept
This descriptive framework for WLAN resource usage defines three MAC bandwidth components that are tightly coupled via

equations (6), (9), (10), and (11). This framework gives a compact and intuitive description of the operation of the 802.11 MAC mechanism that is ideally suited to supporting radio resource management schemes such as admission control and QoS provisioning on 802.11 WLANs.




Figure 5: Performance Characterization using MAC Bandwidth Components

For example, it can be used to give an advanced warning of the on-set of saturation, see Figure 5. Consider the case of a single STA where the offered traffic load is linearly increasing. As the load increases, BW_{load} also increases as more and more wireless frames are transmitted. Similarly, BW_{access} must also increase as more transmission opportunities need to be obtained in order to support BW_{load} . As BW_{load} increases, BW_{busy} increases and BW_{idle} decreases. The effect of an increasing BW_{access} and a decreasing BW_{idle} causes BW_{free} to decrease until a point is reached where BW_{free} has been reduced to zero or where there is just a sufficient amount of idle time on the medium to support the access requirements for the transmitted load. Increasing the offered load beyond this point does not result in any further increase in either BW_{load} or BW_{access} as saturation has been reached. Severe packet loss results at this point. The on-set of STA saturation occurs when $BW_{free}(k)$ has been reduced to zero, i.e.

$$\left. \begin{array}{l} BW_{free}(k) = 0 \\ \text{or } BW_{access}(k) = BW_{idle} \end{array} \right\} \text{ under saturation} \quad (17)$$

In other words, all of the idle bandwidth is being used by the STA in accessing the medium in order to support its load. There is no additional capacity remaining to support an increased load, i.e. the STA is saturated and cannot support any more traffic.

5. EXPERIMENTAL RESULTS

The experimental set up involved sending video streams between stations on the WLAN and on the wired backbone. The video sessions were hosted using the Microsoft NetMeeting application which employs the ITU H.263 video codec for compressing and encoding the webcam images. The WLAN was based around IEEE 802.11b equipment, namely Cisco Aironet 350 series AP and PCMCIA cards. The traffic generator used (PCEN TG) was obtained from the Postel Center for Experimental Networking [2] and was used to create one way UDP streams between a source

and a sink. The TG was set up to generate 512 byte packets with exponentially distributed inter-arrival times (i.e. Poisson traffic). The tests were run for 300 seconds. The WLAN traffic probe was used to measure and record the MAC bandwidth components BW_{load} , BW_{access} , and BW_{free} at 5 seconds intervals for each wireless station including the AP. The BW_{busy} and BW_{idle} components for the WLAN were also measured and recorded. A number of different test scenarios were investigated and the following results were obtained.

5.1 Streaming Uni-directional Video Traffic 1

In this test scenario, Figure 6, we send a single video stream from a wireless STA (EE) to a station on the wired network via the AP. Under this configuration, there is a single up-link (UL) traffic stream accessing the wireless medium. Figures 7-10 show the results from the WLAN traffic probe recorded for this test scenario.




Figure 6: Streaming video traffic from the WLAN to the wired network

Figure 7 shows the BW_{busy} and BW_{idle} measurements, while Figure 8 shows the BW_{load} measured for both the wireless STA and the AP. The BW_{load} consumed by the video stream is approximately 650 kbps. The small BW_{load} measured for the AP is due to management frames such as beacons and signaling packets from the NetMeeting application.




Figure 7: BW_{busy} and BW_{idle}




Figure 8: BW_{load}

Figure 9 shows the access bandwidth requirement for the wireless STA and the AP. The access bandwidth requirement for the wireless STA to support the video stream is $BW_{access} \approx 270$ kbps, while for the AP the access bandwidth requirement for its load is $BW_{access} \approx 50$ kbps. As a consequence of its smaller access bandwidth requirement, the AP also experiences a larger free bandwidth, see Figure 10.



Figure 9: BW_{access}




Figure 10: BW_{free}

5.2 Streaming Uni-directional Video Traffic 2

In this test scenario we stream a single video stream from a wireless STA (EE) to another wireless STA (4C). Under this configuration, there are now two traffic streams accessing the wireless medium: an up-link (UL) stream to the AP and a

downlink (DL) stream from the AP. Essentially, the AP is relaying the video traffic stream between the two wireless STAs.




Figure 11: Streaming video traffic from the WLAN to the WLAN

Figure 12 shows the BW_{busy} and BW_{idle} measurements, while Figure 13 shows the BW_{load} measured for both the wireless STA and the AP.




Figure 12: BW_{busy} and BW_{idle}




Figure 13: BW_{load}

The BW_{busy} has increased compared to the previous scenario as both the wireless STA and AP are transmitting the video stream (i.e. the UL and DL respectively). The BW_{load} is essentially the same for both the wireless STA (EE) and the AP as they are both carrying the same video stream. Again $BW_{load} \approx 650$ kbps. The AP has a slightly larger load as it is also transmitting management

frames and NetMeeting signalling frames. The access bandwidth requirements are essentially identical where $BW_{access} \approx 270$ kbps, see Figure 14, as both the wireless STA (EE) and the AP are carrying the same video stream. Similarly for the free bandwidth in Figure 15. The slight difference is due to the additional traffic carried by the AP.




Figure 14: BW_{access}




Figure 15: BW_{free}

5.3 Streaming Bi-directional Video Traffic 1

In this test scenario, Figure 16, we stream bi-directional video traffic between a wireless STA (EE) and a station on the wired network via the AP. Under this configuration, there are now two traffic streams accessing the wireless medium: an up-link (UL) stream to the AP and a downlink (DL) stream from the AP.




Figure 16: Streaming bi-directional video traffic from the WLAN to the wired network

Figure 17 shows that the BW_{busy} for this scenario is similar to the previous case as there are two video streams being carried on the wireless medium. The BW_{load} for the two streams is shown in Figure 18 where again $BW_{load} \approx 650$ kbps. The reason for the difference between the characteristics is that here the two video streams are different unlike in the previous case where the AP was essentially relaying the same video stream.




Figure 17: BW_{busy} and BW_{idle}




Figure 18: BW_{load}

Figure 19 shows the access bandwidth requirements for the two streams where $BW_{access} \approx 270$ kbps. The access bandwidth requirement for the AP is somewhat larger owing to the additional traffic being transmitted. Both streams experience a similar free bandwidth as they have similar access bandwidth requirements, see Figure 20.




Figure 19: BW_{access}

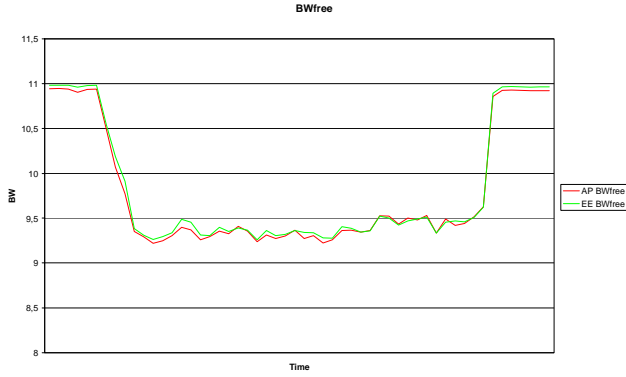


Figure 20: BW_{free}

5.4 Streaming Bi-Directional Video Traffic 2

In this test scenario, Figure 21, we stream bi-directional video traffic between two wireless STAs (EE and 4C) with the AP essentially acting as a repeater. Under this configuration, there are now four traffic streams accessing the wireless medium: two up-link (UL) streams from the two STAs and the two corresponding downlink (DL) streams from the AP. As there are four traffic streams on the WLAN, one would expect this scenario to have the largest resource usage of the cases considered so far.




Figure 21: Streaming bi-directional video traffic from the WLAN to the WLAN




Figure 22: BW_{busy} and BW_{idle}

Figure 22 shows that the $BW_{busy} \approx 3$ Mbps for this scenario which is the largest of the cases so far considered, as expected. Figure 23 shows BW_{load} from the two wireless STAs and the AP. For each of the wireless STAs, $BW_{load} \approx 650$ kbps as in the previous cases, however $BW_{load} \approx 1.3$ Mbps for the AP as it is transmitting the two video streams on its DL.




Figure 23: BW_{load}



Figure 24: BW_{access}

The increased traffic load of the AP is also reflected in its access bandwidth requirement, see Figure 24, where it is approximately double that of the individual STAs. This in turn is reflected in the reduced free bandwidth, see Figure 25.




Figure 25: BW_{free}

5.5 Introducing Background Data Traffic

In this test scenario, Figure 26, we stream bi-directional video traffic between two wireless STAs with the AP essentially acting as a relaying station. However, we now introduce some background TG data traffic from a wireless STA (TG src) to the wired network. Under this configuration, there are now five traffic streams accessing the wireless medium: three up-link (UL) streams comprising the two video streams and the data stream and the two corresponding downlink (DL) video streams from the AP. In this test, the load from the traffic generator (TG) is increased in steps of 0.5 Mbps. The TG load is applied for 30 seconds and then removed for 30 seconds in order to allow any buffers in the system to clear.




Figure 26: As Figure 21, except background data traffic introduced




Figure 27: BW_{busy} and BW_{idle}

Figure 27 shows the step increments in BW_{busy} as a result of the TG's increasing load. However, a point is reached when BW_{busy} saturates, this point corresponds to a $BW_{load} \approx 4.25$ Mbps for the TG and a corresponding $BW_{access} \approx 1.4$ Mbps. The access requirements for the two video stream remains unchanged at $BW_{access} \approx 270$ kbps each and approximately double this value for the AP relaying the two streams, see Figure 29.

Figure 30 shows the BW_{free} for all of the streams. The TG is experiencing less BW_{free} owing to its larger access requirement. One might expect $BW_{free} = 0$ for the TG as its load is saturated, however $BW_{free} \approx 3$ Mbps which suggests that it is the TG

application that is saturated rather than the wireless station transmitting its load.




Figure 28: BW_{load}

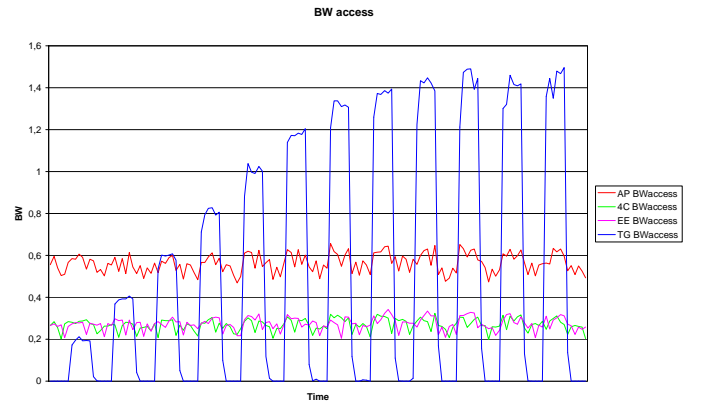


Figure 29: BW_{access}




Figure 30: BW_{free}

However, the most useful result from this test scenario was that the quality of the transmitted video began to noticeably deteriorate when the TG $BW_{load} \geq 4$ Mbps. At this point the AP $BW_{free} \approx 4.25$ Mbps which would suggest that the minimum BW_{free} requirement for a single video stream is approximately 2.1 Mbps. While the NetMeeting H.263 video stream requires only 270 kbps of idle bandwidth in order to access the medium, it never the less requires a significantly greater amount (approximately 8 times greater) of idle bandwidth in order to guarantee QoS. In terms of transmitting

NetMeeting H.263 video stream on an 802.11b WLAN, we have the following:

$$\begin{aligned} BW_{load} &\approx 0.65 \text{ Mbps} \\ BW_{access} &\approx 0.27 \text{ Mbps} \\ BW_{free} &\geq 2.1 \text{ Mbps} \end{aligned} \quad (18)$$

This information can be used to calculate the maximum number of NetMeeting sessions, N_{max} , that can be supported on an 802.11b WLAN. As each session comprises two H.263 video streams, the maximum number of streams will be $2N_{max}$. If both streams are relayed through the AP, then there will be a total of $4N_{max}$ streams attempting to access the medium. Therefore

$$\begin{aligned} BW_{idle} &= 11 - \sum_k BW_{load}(k) \\ &= 11 - 4N_{max} \times 0.65 \\ &= 11 - 2.6N_{max} \end{aligned} \quad (19)$$

The AP downlink is the critical stream as it carries the aggregate of the video streams, i.e. $2N_{max}$ video streams, therefore its idle bandwidth requirements to support access and QoS is:

$$\begin{aligned} BW_{idle} &= BW_{access}(AP) + BW_{free}(AP) \\ &\geq 2N_{max} \times (0.27 + 2.1) \\ &\geq 4.74N_{max} \end{aligned} \quad (20)$$

Equating (19) and (20) gives $N_{max} \leq 1.50$ or $N_{max} = 1$ since it must be an integer. Therefore, an 802.11b WLAN can support a most one NetMeeting session (with an acceptable QoS) between two wireless stations.

6. SUMMARY AND CONCLUSIONS

We have described a novel 802.11 WLAN traffic probe for measuring the resource usage on a per-station basis. The WLAN probe is non-intrusive and operates by passively “sniffing” the wireless frames on the medium. Moreover the probe specifically addresses operation at the L2/MAC layer unlike current commercially available WLAN analysers which are essentially network protocol analysers and tend to operate at the network layer and above. Consequently, the probe is capable of producing the type of resource information required for radio resource management and QoS provisioning schemes.

The probe is based around a particularly useful descriptive framework for identifying network resource usage on WLANs involving the concept of *MAC bandwidth components*. In particular, three MAC bandwidth components have been identified: A *load bandwidth* (BW_{load}) that is associated with the transmission of the data frames, an *access bandwidth* (BW_{access}) associated with the contention mechanism (whereby a station wins access to the wireless medium) and a *free bandwidth* (BW_{free}) that is associated with the QoS. This framework results in a compact and intuitive description of MAC resource usage on a per-station basis that is particularly suited to radio resource management schemes.

The WLAN traffic probe was used to analyse resource usage in a number of experimental set ups involving streaming video data between wireless stations. The results show the probe to be useful in identifying and quantifying both resource usage and resource requirements on a per-station basis. The probe also allows the interaction (and corresponding impact on performance) between stations to be observed and quantified. For example, we have shown that the transmission of video streams requires a significant amount of free resources to be reserved if QoS is to be ensured. Most significantly, the probe allows this amount of free resources to be quantified and to be monitored on a real time basis.

It is expected that the WLAN traffic probe described here will be an important element in any QoS provisioning scheme for 802.11 WLANs.

7. ACKNOWLEDGEMENTS

The support of the Informatics Research Initiative of Enterprise Ireland is gratefully acknowledged.

8. REFERENCES

- [1] IEEE, “Part 11, Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications”, *IEEE Std 802.11*, 1999 Edition.
- [2] Postel Center: www.postel.org