

Operational and Fairness Issues with Connection-less Traffic over IEEE802.11b

Theo Pagtzis, Peter Kirstein and Steve Hailes
Department of Computer Science
University College London
Gower Str. WC1E 6BT
London, U.K.
{t.pagtzis,p.kirstein,s.hailes}@cs.ucl.ac.uk

Abstract—The IEEE802.11 group has recently ratified high rate (HR) extensions to enable high speed wireless communications over WLANs. The HR extensions specified in revision IEEE802.11b, encompass mainly new RF modulation schemes. This paper attempts an experimental evaluation of the performance characteristics of 802.11b in terms of throughput and loss over high speed transmission rates with respect to connection-less network traffic. We present a simple analysis of the protocol's throughput capacity over high speed rates while we reveal fundamental design considerations that prevent 802.11b from reaching its true throughput potentials in the light of rate adaptivity. We further recommend some extensions to the Medium Access Control (MAC) protocol sub-layer that reconsider the multi-rate compatibility requirement while maintaining fairness in throughput between nodes at short or long distances within range from a Base Station (BS). The recommendations subject to simulations as work in-progress, expect to effect improvements in *throughput* over a proportionally-fair rate fallback scheme, in the order of 15% for transmission rates of 11 Mbps. We also provide some key observations to enable efficient protocol design for adaptive mobile environments.

Keywords—802.11b, Medium Access Control, Point Coordination Function (PCF), Distributed Coordination Function DCF), Mobility, Dynamic Rate Fallback, Fairness.

I. INTRODUCTION

Wireless Local Area Networks (WLANs) are becoming a significant part of the Internet constituency as a popular alternative to high installation and maintenance costs incurred over wired LAN infrastructures. The license de-regulation of the ISM band has allowed WLANs to facilitate ubiquitous communication, coupled by location independent end-to-end network services in restricted spatial domains such as military settings as well as campuses, offices or enterprise facilities.

However, integrating WLANs in the Internet is not quite transparent as expected in contrast to their wire-line counterparts. Limited bandwidth availability and higher error rates induce a noticeable presence of wireless to the mobile Internet user. The long-enjoyed assumption of wire-line networks that attributes losses primarily to congestion [1], as physical medium bit error rates (BER) lie in the region of $10E-9$, is rendered in wireless networks partially invalid.

The nature of the physical medium itself, accounts for breaking this fundamental assumption. While each network segment in a LAN is electromagnetically isolated from the rest, neighboring segments in a Wireless LAN (WLAN) are prone to interference from each other. The inherent property of radio networks [2], **mobility**, accentuates further the problem; propagation and

This work was supported by the PIMMS research initiative, funded by the British EPSRC research council.

attenuation effects, like path loss and multi-path fading are incurred as a result of varying mobility characteristics such as velocity and distance with respect to the associating Base Station. Wireless BER of the order of $10E-3$ to $10E-6$ [3], [4] dictate requirements for corruption control in addition to congestion control, in view of the restricted bandwidth resources available for service provisioning in diverse application domains like voice and interactive multimedia [5]. We, thus, deem essential that integrating WLAN technologies with the Internet requires a concrete understanding of the underlying medium access control (MAC) protocols deployed over real WLAN vendor implementations.

There has been a significant interest in the design [6], [7], [8], and standardization [9], [10] of local area communication protocols for wireless networks. Amongst them the study group 802.11 formed under IEEE802 project, and the ETSI High Performance European Radio LAN (HIPERLAN) [11], [12], [13] both commissioned to provide recommendations for an interoperable WLAN standard. Both standards target the physical (PHY) and medium access control sub-layer. A number of vendor implementation have adopted the 802.11 standard while recently the IEEE802.11 study group ratified high rate (HR) extensions over the protocol standard to effect higher transmission rates.

This work concentrates on experimental evaluation of the 802.11b protocol [14]. We investigate the performance profile of 802.11b over proven vendor implementations. This work also aims in verifying and extending results in [15], [16], [17] over the new HR extensions. For the sake of conciseness, the paper presents conclusive summaries of our findings; these assume a basic understanding of the DCF and PCF functions of 802.11b. For more comprehensive discussions and results, the interested reader is invited to refer to [18] or for detailed background information, the IEEE 802.11 specification [14].

This paper is organized as follows: Section II details the environment and methodology of our experimental setup. Section III presents results over the experimental scenarios as well as transmission overheads accounted for the MAC and PHY layers of 802.11b. Section IV presents fairness issues with regard to rate compatibility for 802.11b. We conclude with a summary of our findings in Section V.

II. ENVIRONMENT AND METHODOLOGY

The experiments employed the Lucent IEEE/802.11b wireless network interfaces¹, operating at 2.4GHz over Direct Sequence Spread Spectrum (DSSS) and delivering signaling rates up to 11 Mbps.

To the best of authors' knowledge, most of the published results² have been considering first generation (WAVELAN I) network interfaces; these have been 802.11 implementations employing low bandwidth RF modulation schemes with maximum signaling rates up to 2 Mbps. To this end, this paper aims to contribute fresh observations about the protocol's performance over it's HR extensions.

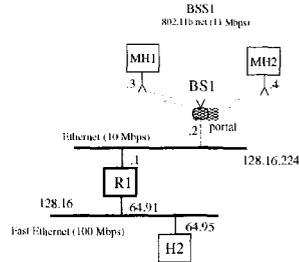


Fig. 1. Experimental environment employed

The experimental topology of Figure 1 featured a set of four FreeBSD (4.1), Pentium class, PC hosts. Both MH1 and MH2 were equipped with PCCARD 802.11b WNICs. These were driven by the device drivers included in the above OS release. On the Ethernet side H1 and R1 featured Fast Ethernet network interfaces. The Access Point BS, was an 802.11b (L2) bridge available from the same vendor; its wired segment featured a 10 Mbps Ethernet interface³. A PC-based Access Point was not considered for reasons of efficiency and performance described in [18]. The private wireless subnet routed by R1 ensured no bypassing traffic during the course of the experiments.

For the generation of traffic patterns we engineered *TrafSim* [18]. Through *TrafSim*, the tests featured variable synchronisation over configured client permutation test sets, variable isochronous transmission rates and signal strength tracing in unattended mode. Our traffic patterns generation borrowed implementation ideas from other well-known measurement tools like *rtptools* [19] and *DBS* [20].

We increased the clock frequency of the FreeBSD kernel to 1000Hz⁴ effecting isochronous transmission rates up to 12 Mbps. This would be sufficient for simulating congestion conditions over the 802.11b link. We ensured that increasing the kernel clock resolution bears no adverse effects to the performance of the OS over the available set of processor speeds [18]. Test packet flows consisted of 10K packets, while test permutations generated by the (*send rate*, *packet size*) test space tuple, were repeated 15 times to ensure reduced variance in observed values.

¹referred as WNICs henceforth

²to the date of writing

³there was no Fast Ethernet interface on any Access Point implementation at the time of writing

⁴default value is 100Hz

Experiments were isolated⁵ from potential sources of radio interference with SNR readings captured throughout each test on the WNIC. Post-measurement processing and analysis was carried out on traffic traces captured near the receiver. All tests were performed in both directions between Ethernet and 802.11b peers. Our results were tested against a confidence interval of 95% to ensure statistical validity.

III. EXPERIMENTAL SCENARIOS

A. Peak Throughput and Error Rate

Performance of the 802.11b WNICs in terms of peak throughput and error rates was first tested with respect to wireline LAN segment. The wireless host, MH1 was stationary with very small distance from the BS1. Figure 2 shows, the user throughput at the wireless receiver reaching a maximum of 6.22 Mbps (528.91 packets/sec) for packet sizes close to the MTU⁶. This is 56.36% of the total bandwidth capacity of the 11 Mbps 802.11b link. 44% of the signaling rate is consumed on transmission overheads and latencies.

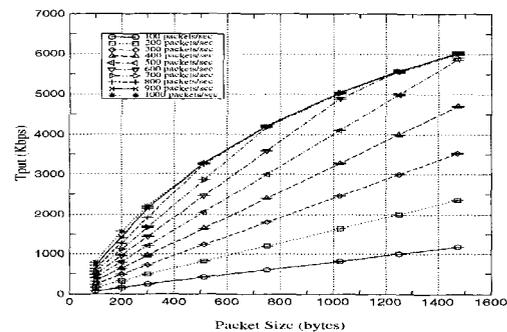


Fig. 2. Throughput on Wireless receiver (varying packet size)

Throughput performance over different send rates exhibits a somewhat different behavior. This is shown in Figure 3. It can be seen that above a certain send rate threshold, when packet size remains constant, user throughput remains constant also. This may seem counter-intuitive when compared against throughput performance of Fig.2. However, it is clearly explained by looking at the corresponding frame error rates for constant throughput when the send rate increases, as well as considering that the Ethernet link supported by BS is 10Mbps.

From the throughput peaks of Figure 3 and the send rate points of Figure 5, we deduce that increasing the send rate beyond a certain threshold, for constant packet size, saturates the 10 Mbps Ethernet segment between R1 and BS. This implies at the points where the throughput remains constant, BS exhausts its receive buffers on its 10 Mbps Ethernet interface as it receives packets at a rate faster than its 10 Mbps link can cope with, during reception. As such, the wired segment of BS will drop packets before these can be transmitted over the wireless segment. In the reverse direction throughput shows identical behaviour for send rates up to 500 pps (and thus omitted from the graph), while it drops with higher variance at offer loads of

⁵conducted during the night

⁶assumed to be 1500 bytes

600 pps or above. Such behaviour may be attributed to both the driver efficiency the WNIC in MH1 as well as the limitation imposed by the 10 Mbps Ethernet link effected by the BS.

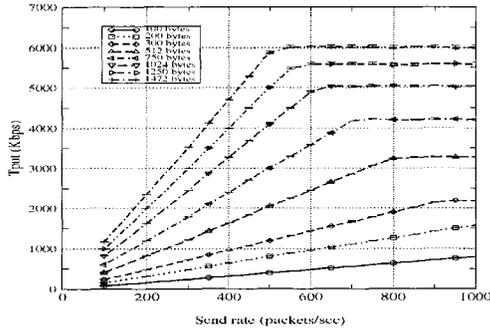


Fig. 3. Throughput on Wireless receiver (varying send rate)

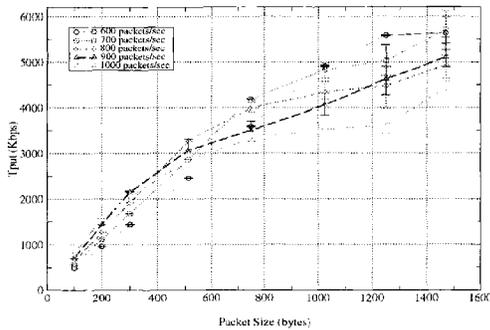


Fig. 4. Wired receiver throughput (varying packet size)

With respect to packet loss rate, Figure 5 shows that for transmission rates up to 550 pps frame error rates (FER) are in the region of $4E-3$, implying that at peak rates 1 in 250 frames on average is lost for packet sizes close to the Ethernet MTU. Beyond 550 pps the FER become significantly high (way above 1%) implying a saturated link beyond the transmission capacity of the protocol.

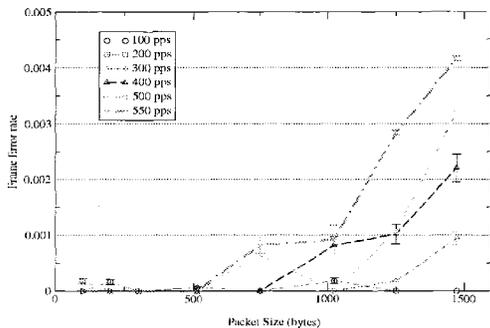


Fig. 5. Packet Loss Rate of Wireless receiver

On the *wired* receiver, the throughput becomes noisy and effectively oscillates for send rates above 600 pps, while the max-

imum user throughput is observed at 526.36 pps (6.19Mbps) for packet size of 1500 bytes. Loss rate curves on the wired receiver are similar to Figure 5 and, thus, omitted. However, in this case it is the wireless sender that experiences buffer overflows⁷ as transmission rates above 550 pps overrun the transmission buffer capacity of the wireless interface. This implies that the output queue `if_output()` of the wireless interface gradually fills up to the point where the transmission buffers of the wireless interface are exhausted and frames in the output queue get dropped.

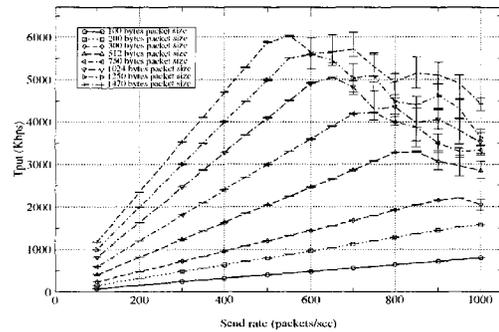


Fig. 6. Wired receiver throughput (varying send rate)

A.1 Transmission Overheads

In 802.11 DCF/PCF function, transmission time is allocated according to the CFP repetition interval $T_{CFP_{ri}}$ and is equal to:

$$T_{CFP_{ri}} = T_{CP} + T_{CFP}$$

where T_{CP} is the time allocated to DCF (contention period) and T_{CFP} the time allocated to PCF (contention-free pollable period).

A.2 DC Function Overheads

In contention-based (CP) mode, the wireless node has to (re)contend for the medium for *every* MPDU⁸ scheduled for transmission (Tx). The Tx cost of a single MPDU during CP mode (including the cost induced by the binary exponential backoff, as the *contention window size*), assuming RTS/CTS handshake is:

$$C_{CP_{Txdata}} = C_{difs} + 3C_{sifs} + C_{RTS} + C_{CTS} + C_{MPDU_{data}} + C_{ACK} + C_{Backoff} \quad (1)$$

$$C_{Backoff} = \text{int}(C_{win} \cdot \text{rnd}()) \cdot T_{slot} \quad (2)$$

$$C_{win} = 2^{2+\rho} - 1, \wedge \text{rnd}() \in [0, 1], \rho \in \mathbb{N} \wedge \rho \in [1, 6]$$

with $\text{rnd}()$ the pseudo-random number generator and C_{win} the contention window size randomized. ρ is the number of attempts to capture the medium after the expiry of the backoff counter and a subsequent collision with another node that attempts to do

⁷the output queue is typically 50 packets long

⁸in the 802.X terminology a MAC protocol data unit (MPDU), is the unit that encapsulated the data payload

so, T_{slot} has a nominal value of $20 \mu sec$ and identifies a unit measure of C_{win} . The C_{win} value increases exponentially on a re-capture attempt of the medium following a collision, with a freeze after six consecutive collisions. C_{difs} and C_{sifs} represents inter-frame intervals that account for Tx latencies such as propagation and Rx/Tx turnaround times. Table I lists the header cost at different contention modes and layers of both encapsulating and encapsulated framing. It can be seen that control and framing overheads total to 178 bytes for Tx of an MPDU during CP. For MPDU payload below the RTS_Threshold, the RTS/CTS mechanism is not employed; that brings the Tx cost down to 96 bytes of framing for both MAC and PHY.

TABLE I
FRAME TYPES AND ASSOCIATED HEADER LENGTH

Mode	frame type	length (bytes)
any	PLCP - PHY	24
CP	MSDU (RTS) - MAC	20
CP	MSDU (CTS) - MAC	14
CP	MSDU (ACK) - MAC	14
any	MPDU - MAC	34
CFP	MMPDU (Beacon) - MAC	33-323
CFP	MSDU (CFend) - MAC	20
CFP	MPDU CFpoll,ack,poll+ack - MAC	34

A.3 PC Function Overheads

We are interested in the lower bound of transmission costs affected by a single host. Such bound is described for both uni-/bi-directional traffic as the initialization time of CFP and a single data frame exchange between the AP and the node separated by the respective IFS intervals. This is:

$$\begin{aligned}
 C_{CFP_{Txdata}} &= C_{CFP_{init}} + C_{MDPU_{poll}} + \\
 &\quad C_{MPDU_{D1out+ack}} + \\
 &\quad C_{MPDU_{D2in+ack}} + 2C_{sifs} \quad (3) \\
 C_{CFP_{init}} &= C_{pifs} + C_{beacon}
 \end{aligned}$$

Table II gives a summary of header overheads per user data packet. Both PCF and DCF functions incur similar transmission overheads if RTS/CTS is employed. Exception to this is the setup cost for PCF which iterates per CFP repetition interval. This implies that a longer CFP amortizes the setup cost over more MPDUs in either direction.

A.4 High Rates and Signaling Time Cost

The 64-code set [21] at the 11 Mbps signaling rate, allows a total of 8 bits of information to be encoded per CCK symbol [22]. At 5.5 Mbps the processing reduces the amount of information bits per codeword to 4. Both 5.5 and 11 Mbps rate support a codeword transmission rate of 1.375 MHz as the PN-sequence is now 8 chips instead of 11 supported in 1 and 2

Mbps. As such, each bit is Tx-cd at 11 Mbps between a wireless node and the Base Station in 90,909 nsec.

Furthermore, depending on the coordination mode (DCF or PCF), the per MPDU Tx cost⁹ comprises of the following constituent components:

Tx time of MAC frame payload. The size of the frame body may vary between 0-2312 bytes. Subsequently, the time slot required for transmitting such payload C_{pload} varies in the interval of $[0, 1681.44] \mu sec$.

CP-mode Tx with RTS/CTS. The cost of a single MPDU transmission (dominated by $C_{win} \in [0, 255]$) becomes:

$$907.62 + C_{pload} \leq C_{CP_{Txdata}} \leq 6007.62 + C_{pload} \mu sec \quad (4)$$

CP-mode Tx without RTS/CTS. A case for multi-cast/broadcast data frames or for MPDU frames below the RTS threshold; an ACK is only required in the case of unicast MPDUs transmitted during CP. This yields:

$$478.9 + C_{pload} \leq C_{CP_{Txdata}} \leq 5708 + C_{pload} \mu sec \quad (5)$$

During the experiments C_{pload} varied between 146 and 1518 bytes (including Ethernet headers). That is, the lower time bound ($C_{win} = 0$) required for a successful packet transmission is in $[585.08, 1581.9] \mu sec$.

CFP-mode transmission (PCF). This encompasses the CFP initialization time prior to any time spent in actual data transmissions. Depending on the size of the Beacon frame, the cost is bound by $[236.36, 447.27] \mu sec$. The beacon size and, thus, Tx time depends to a small extent on the size of the Service Set ID (SSID)¹⁰. It's main dependency is the size of the traffic indicator map (TIM)¹¹, which indicates outstanding traffic for delivery to stations. As the CF period commences after PIFS time, the total $C_{CFP_{init}}$ cost increases by $30 \mu sec$ at the bounds of the above interval. Thus, $C_{CFP_{Txdata}}$, as in (3), for a single station during unidirectional CFP transmission of an MPDU is $680.16 + C_{pload} \mu sec$.

In the presence of N stations, the cost of $C_{CFP_{init}}$ is spread amongst them, becoming:

$$\begin{aligned}
 &\frac{266.36 + 680.16N}{N} + C_{pload} \leq \\
 &\leq C_{CFP_{Txdata}} \leq C_{pload} + \frac{477.27 + 680.16N}{N} \quad (6)
 \end{aligned}$$

⁹Details on the individual figures can be found in [18]

¹⁰that may vary between 2-34 bytes

¹¹TIM varies between 1-251 bytes

TABLE II
CONTENTION MODES AND ASSOCIATED OVERHEADS

Mode	header o/heads (bytes/MPDU)
$CP_{RTS/CTS}$	178
CP	96
CP_{init}	n/a
CFP	174
CFP_{init}	57-347

For bidirectional traffic, the Tx cost of a single MPDU increases by an extra data frame of potentially different size. For each frame sent by the wireless host, the AP is expected to Tx a frame pending for delivery to that station. That multiplies C_{load} by 2 for bidirectional traffic. The above figures assume that CF-Polls are not lost and that stations ACK the poll message. At this stage they do not include the case where a node does not respond to the CF-Poll. This is described in [18].

Figure 7 shows how Tx cost increases with the MPDU payload during DCF and PCF, for a single station associations with the AP. It can be seen that PCF with a single wireless source incurs almost the same cost as DCF with RTS/CTS enabled for the entire size range of an MPDU payload, when the contention window lies between counter values of a single backoff attempt (i.e. no collisions). This is not the case for the basic DCF access mode with no RTS/CTS control frames. Such observation doesn't hold if the number of stations associated with the AP is $N > 1$.

Transmission times are considerably shorter when no RTS/CTS access control is enabled. However, this is true only for a small number of stations and with small offered load (low contention). As the number of associated nodes together with their offered load grows, contention for medium capture increases; this induces a higher probability of collisions [23] which throttles the contention window and number of Backoff attempts. This is the primary cause for growing transfer delays during DCF as shown by Chhaya [24] and Chen [25] which may be reduced through the use of RTS/CTS control. The latter is, however, not a panacea as it renders beneficial only above a certain frame length [26].

We further evaluated the validity of our results with respect to throughput and loss by approximating numerically the protocol's capacity under both PCF and DCF [18]. The estimate approximation of the transmission capacity was found somewhat higher (around 6.54 Mbps) than observed in the throughput results. Since we did not experience collisions for a single station (since the AP does not deliver any packet, while ACKs arrive in response to transmissions and before the renewal period set by the next DIFS interval) we concluded that such difference is largely due to inefficiencies of the driver.

DCF offers a *fairer* chance of acquiring the medium by having all nodes contending for channel access after *each* MPDU transmission. However, the probability of collisions during a backoff is bound to increase with the number of stations associating with the AP; this can increase the backoff time for a single node up to 5.3 ms. As shown in (4), C_{win} is dominant during an MPDU transmission over DCF and can, thus, induce significant variance per frame payload (packet). This reasons for DCF being *unable* to offer distributed time-bounded services (DTBS). On the contrary, PCF, through its CF polling scheme sets specific upper bounds of delay.

Measurements showed that user throughput can increase up to 7.8 Mbps for packet sizes larger than 1500 bytes and up to 2266 bytes. This figure is, however, non-realistic considering that packet size mostly bound by the MTU size; the standard MTU can be as low as 586 bytes, while specific application classes, such as interactive multimedia or VoIP, strive for small packet sizes. For instance, audio packets range between 50-500

bytes depending on the codec and the error resilience scheme employed. Video packets peak around 1024 bytes. It is, thus, clear that a MAC protocol should strive to increase efficiency and capacity over small packet sizes rather than emphasizing on large packet size for protocol capacity.

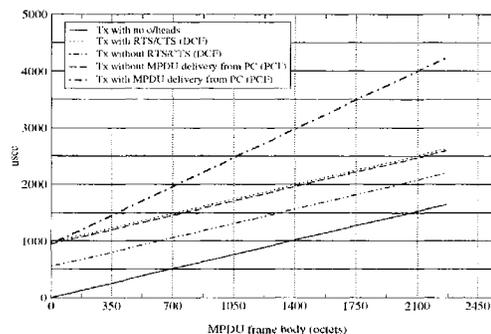


Fig. 7. Transmission cost of single data frame under DCF and PCF

B. Bi-directional streaming over the wireless link

This scenario looked at MAC behaviour during simultaneous transmission and reception. The tests simulated overload conditions over send rates of up to 400 pps (4.7 Mbps) in each direction. Measurements were taken only during times that the WNIC was interleaving between *both* Tx and Rx of streamed UDP flows. In this manner the throughput and loss results were prevented from being *inflated* by single transmission fragments if one of the two peers finished its Tx earlier.

TABLE III
BI-DIRECTIONAL THROUGHPUT AT SIGNIFICANCE LEVEL: 0.95

Packet Size (bytes)	Throughput (Kbps)					
	200pps		300pps		400pps	
	Tx	Rx	Tx	Rx	Tx	Rx
150	221	240	316	360	411	479
250	367	400	527	600	689	799
350	516	560	739	840	947	1120
450	670	720	949	1080	1213	1439
550	807	880	1139	1320	1470	1759
650	953	1040	1338	1560	1723	2077
750	1091	1200	1543	1800	1941	2397
950	1365	1520	1945	2280	2476	2689
1050	1497	1680	2098	2520	2722	2711
1250	1786	1999	2413	2998	3001	2900
1350	1921	2161	2569	3234	3097	2995
1470	2026	2353	2785	3470	3230	3073

We observed that for Tx rates up to 300 pps (3.52 Mbps) in each direction, the MH2 receiver experienced higher throughput than the H2 receiver. While for small packet sizes the increase in the observed throughput values at the 802.11 station was only marginal, as packet size grows the increase gets slightly bigger (see Fig. III). This was true for all packet sizes up to 1470 bytes of user data, even though the wireless link was running into mild

congestion at a combined send rate of 7.05 Mbps. The above is independent of which host initiated Tx first.

The higher throughput experienced at the wireless receiver largely justifies the original observation that the wired source would complete transmission faster than the wireless host. While a detailed discussion is provided in [18] we inferred from the sequence numbers of frames exchanged between BS1 and the 802.11b station, that the latter is a slow sender compared to BS1. This can only be attributed to different implementations between the FreeBSD WNIC driver and the kernel firmware of BS1. This make clear that the performance of the driver induces its own limitations over representative inference about the protocol's performance.

While the 802.11 receiver experienced higher throughput than the fixed receiver, it also witnessed higher loss compared to the fixed host. In fact, it is suprizing to note that the fixed receiver did not experience any loss at all even for send rates up to 400pps in each direction. On the contrary FER at the 802.11 station varied between $8E-4$ and $8E-3$ for packet sizes above 1150 bytes and 300 pps send rate. This is shown in table IV. For lower send rates (100 and 200pps) we observed only single burst losses that varied between 1-6 packets. Longer loss runs (6-10 packets) were observed but were extremely sparse ($1.25E-05$) for aggregate send rates below the observed peak throughput of the 802.11b link.

TABLE IV
LOSS RATE ON HIGH-THROUGHPUT WIRELESS RECEIVER

Packet Size	Loss Rate (@300pps)	CI 95%
1150	$8.61E-4$	$1.33E-4$
1250	$5.29E-3$	$8.13E-4$
1350	$2.12E-3$	$4.26E-4$
1470	$8.36E-3$	$1.84E-4$

C. Throughput fairness between different wireless hosts

Fairness tests considered identical offered loads between two 802.11b hosts (MH1 and MH2), that transmit to a single wired host from initially stationary positions and a distance less than 2 meters from BS1. Transmission rate and packet size were kept constant and within the capacity of the MAC protocol (up to 300 pps) for each wireless sender.

Figure 8 shows the reduction in throughput¹², when at time T MH2 attempts to send traffic during an in-progress transmission by MH1. As expected, there was no significant difference in the observed throughput of these hosts. The average throughput for the two 802.11b hosts, was 3.15 Mbps and 3.17 Mbps respectively. This implies that for small number of 802.11b nodes the MAC sublayer allocates a fair share of its bandwidth capacity for stations **with similar distances from the base station**. In addition, the MAC layer offered the same aggregate throughput irrespective of the number of associated 802.11b stations (peak of 6.22 Mbps).

Distance of MH1 from BS1 was then increased, while the distance of transmitter MH2 was kept constant. As shown in Fig-

¹²throughput of the incoming transmission is shifted along the X axis so as to observe potential differences

ure 9, bandwidth sharing between the two 802.11b hosts at different distances from BS1 was no more the same. This is not unreasonable since the propagation delay increases with no increase in the transmitter's power. Propagation effects instigate poor channel conditions over distance. This induces increased corruption in the received data or control or management frames which triggers retransmissions through ARQ. Without ARQ at this stage, corruption in received frames escalates to subsequent collisions which if sustained at high signaling rates, may render the MAC protocol unstable. Thus, during poor channel conditions, retransmissions use up the bandwidth and thus reduce the effective data rate¹³.

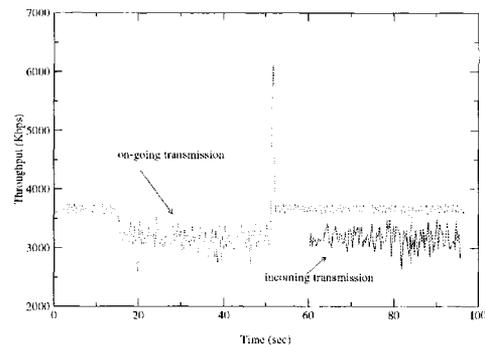


Fig. 8. Throughput for two simultaneous transmitters with symmetric distances of 2 meters

From the above one can now justify the observed drop in throughput for the on-going transmission of Figure 9. That, may be erroneously interpreted as lack of fairness in the throughput experienced by two stations at different distances, but within range from the AP. We observe that **the measure of distance of the wireless host from the Base Station behaves naturally as a service differentiator**. The farther a wireless node moves away from its Access Point, the more its bandwidth and subsequently throughput is bound to degrade. Of course the measure of distance is based explicitly on the measure of the underlying *Signal to Noise Ratio* (SNR) that the wireless node is experiencing.

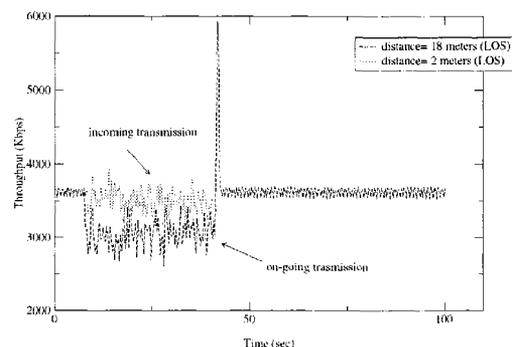


Fig. 9. Throughput for two simultaneous transmitters with assymmetric distances of 18 meters

¹³the effective data rate may be defined a the rate of the information that is correctly transmitted

We, thus, argue that in the face of differentiated services over wireless networks, a provisioning model must consider some measure of *proportional fairness in bandwidth allocation*, where bandwidth may be traded-off with certain SNR levels to maintain connectivity. The SNR, however, may also be affected by other factor orthogonal to distance, such obstructions in the line of sight, or bypassing interference. Nevertheless, in all of these case the measure of SNR may discribe some measure of mobility for the wireless host. Furthermore we argue the measure of SNR should also act as a **link adaptation notification** signal that need be propagated to the IP layer so as to allow faster L3 response times in the face bandwidth fluctuation.

IV. RATE COMPATIBILITY AND FAIRNESS

A basic requirement in the selection of the high-rate modulation scheme has been inter-operability with the base 802.11 protocol. For that purpose, the signal acquisition sequence (PLCP) of 802.11b was designed to be tranceived¹⁴ at 1 Mbps across *all* supported rates, while sending its frame *payload* at different rates. As such, any 802.11b station can sense and defer¹⁵ access to the medium if they cannot process frames at these rates. Table V shows the calculated transmission times/bit for all Tx rates supported by the 802.11b specification.

TABLE V
SUPPORTED RATES AND TX TIME/BIT

rate (Mbps)	bits per symbol	symbol Tx rate	Tx time/bit
1	1	1 Mhz	1 μ sec
2	2	1 Mhz	500 nsec
5.5	2+2	1.375 Mhz	181.8175 nsec
11	6+2	1.375 Mhz	90.90875 nsec

While this serves nicely the rate compatibility objectives, it trades off bandwidth from high signaling rates for the purpose of maintaining connectivity for entire set of associated stations. As such, the 802.11b protocol *sets a lowest common denominator of transmission rate for all stations associated with the AP of a single cell*. This is inherently **unfair** for any 802.11b station operating at signaling rate above 1 Mbps.

We have shown [18] that the PHY preamble and header acquisition sequence consume 192 μ sec over 1 Mbps modulation. If the PHY aquisition sequence was modulated over the 2, 5.5 or 11 Mbps signaling rate, it would require only 96, 34.9 and 17.45 μ sec. This is 2, 5.6 and 11-fold reduction in transmission cost of the PHY layer acquisition sequence! Such savings are quite significant if we consider the transmission cost of a single data frame, especially for small frame payloads over real-time traffic.

Figure 10 illustrates the cost of acquisition sequence C_{as} in proportion to the transmitted packet. C_{as} is considerably larger than the cost of the transmitted payload at 1 Mbps, especially for packet sizes up to 264 bytes. At this point C_{as} is equal to the cost of transmitting the 264-byte payload sustaining a poor

¹⁴constant Tx cost of 192 μ sec

¹⁵NAV update

utilization of the medium. Utilization, however, improves significantly *if* the higher transmission rates (modulation schemes) are effected over the acquisition sequence.

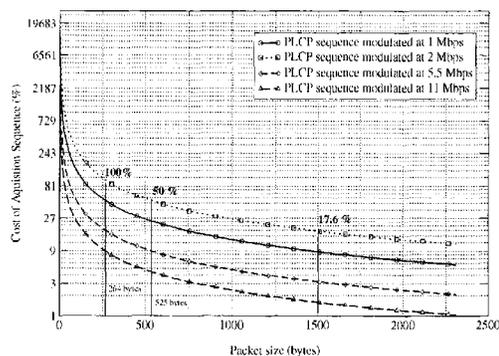


Fig. 10. Transmission cost of the PHY layer acquisition sequence in proportion to the transported payload

At the rate of 11 Mbps the reductions rendered over the PLCP sequence would imply a minimum increase of 15.5% (1 in 8 frames at 1500 bytes) in the throughput rates *currently* experienced over 802.11b at 11 Mbps. We say a 'minimum' because the figures given above do not encompass the PHY layer cost reductions on control (SPDU) and management (MMPDU) frames exchanged prior any data frame transmission.

Of course the design requirement for continuous connectivity beyond the SNR levels offered at the 11 Mbps signaling rate could be accommodated by the straightforward solution of populating the required coverage area with more APs. This is, however, a naive solution towards the robustness of the 802.11b protocol itself.

To cater for the resolution of this fairness issue over the critical requirement of rate compatibility we propose an extension to the 802.11b protocol that makes use of the aforementioned observations of Section III-C. In particular, this extension scheme instigates traffic differentiation for the purpose of resolving the fairness issue. A brief description is provided in the following section.

A. Resolving Fairness through Traffic Differentiation

In this scheme we assume that each station modulates at its own operational signaling rate both for its PLCP sequence and frame payload. However, we require that the association/reassociation function, if solicited by the 802.11 station, is effected with a PLCP sequence modulated at 1 Mbps. The primary requirement for the stability of the protocol is that the stations do not attempt to contend for channel access in the event they operate at a different signaling rates from the operational one currently supported by the AP.

Our scheme proposes that on the side of the Access Point, the Point Coordinator maintains *multiple* PCF polling lists, instead of the single standard one. For that we define the notion of a **Polling List Set** (PLS) denoted as Ψ . Such set will comprise the set of all *active* polling lists and should be expandable to accommodate new ones. Each polling list $\psi_i \in \Psi$, will map to a single element of the set Q , which is defined as the set of all

supported discrete bandwidth classes. We also define the set of all supported modulation schemes in current or future versions of 802.11 as M . Each element q_i , may identify one or more elements m_i from M . That is, a single bandwidth class may encompass a range of discrete supported bandwidths.

The union of all Ψ_n constitutes the set of all stations N_{AP_m} found in a single 802.11b cell for that AP_m . The scheme initially considers two bandwidth classes for 802.11b stations:

- a bandwidth class that supports associations for **highest** operational signaling rate (i.e. 11 Mbps); this is denoted as Q_h .
- an aggregate bandwidth class that supports associations for all **lower** supported signaling rates (such as 5.5, 2 or 1 Mbps)

Each polling list will map naturally to a single contention-free pollable period CFP. Thus, we define further the set of all CF periods as Ω . The size of both Ψ and Ω sets would be dictated by the number of the stations associated with each list as well as their supported and operational signaling rates. The ordering of elements in Ψ at the moment dictates a relative priority according to its bandwidth class. This implies that depending on the order high bandwidth classes may be treated first or last. That puts some extra significance on the size of the individual CF period of ω_n , but most important the total length of all CF periods, i.e. $\sum_{i=1}^n \omega_i$. We are currently investigating optimal ordering schemes with CFP repetition intervals of around 50 ms for this purpose.

In set Ω , for each following ω_l the CF-End message of ω_{l-1} will signal to the stations in psi_{l-1} the duration of ω_l , so that for each station $STA_{k,\psi_{l-1}}$ access to the medium is deferred during that period. This is effected by adjustment of the network allocation vector $NAV_{k,\psi_{l-1}}$ for the k th station. Alternatively, and in the event that such action does not hinder the performance of stations in ψ_i , a station from psi_l may consider re-associating with the BS for ψ_n , where $l \neq n$ for some different bandwidth class q_n .

The 802.11 protocol is required to maintain rate compatibility for two reasons: one is the possibility that a wireless station cannot support certain HR modulation schemes; the second is that even if a station supports all available modulation schemes, the dynamic rate-shifting mechanism of 802.11 will up/downgrade its signaling rate according to its SNR levels from the AP. Both cases, however, rely on the capture of the acquisition sequence.

As such in both situations¹⁶ the station (re)associates with the Point Coordinator as a rate-limited CF-pollable station that will operate only under PCF for a particular CFP identifier. The DCF function for that polling list psi_{l-1} of associated stations is **prohibited**. To complement for the removal of the DCF function from that class of stations, we consider the *optional resizing of the CFP period per ψ_i* to accommodate for repeated polls of stations of the low-bandwidth service class that has expressed such interest/request. This may be conveyed over standard traffic indication map (TIM) type that can be configurable not only by the BS but also by the station of that Polling list.

B. Propagating the Bandwidth Differentiation at the IP layer

Availability of different signaling rates, through either the dynamic rate-shifting mechanism of 802.11 or restricted support

¹⁶stationary or mobile within the range of the AP

for high-rate modulation schemes by stations effect a change in link's characteristics **without** notification to the IP layer and subsequently the sender or receiver of IP traffic. This is critical in a wireless network since the link conditions overload the semantics of congestion without notification to any IP congestion control mechanisms. This is not as critical in the case where the link changes to higher bandwidth (higher signaling rates). It is, though, severe in the case where the link degrades to lower *aggregate* bandwidths. This is further exacerbated when the wireless node roves, not only within the AP coverage area, but between different APs of potentially different domains over IP mobility-aware networks.

We thus, consider essential that L2 rate conditioning should be propagated to the upper layers of the network. Such signal is bound to reduce reaction times of the IP (and further application) layer by informing the required peer entities (routers or hosts) about effected link status. While such L2 *hint* may break the model of independence between different layers of the network, they prevent the cascading effect of *temporal link conditioning* induced to the IP layer when transmitting over dynamically conditioned links. We argue that this is bound to have important implications on the sustainment of service quality over wireless and in particular 802.11 links.

As such, our scheme further proposes a *link status notification* (LSN) effected at the MAC sublayer of the 802.11b protocol. The LSN signal would convey knowledge of the link's dynamic characteristics from the MAC layer to the IP layer, and signal the peer entities for the purposes of triggering some IP flow conditioning rule. We view that such notification should be quite useful further during handoff decisions.

LSN would be implemented in the driver of the wireless interface. It would further propagate as an extension on the IP header that would be interpretable by LSN-aware entities for the purposes of conditioning/adapting their IP flows.

The above scheme is currently work in progress. For it we postpone more detailed discussion and analysis as the object of a separate paper.

V. CONCLUSIONS AND FURTHER WORK

This paper presented a set of measurements as well as fairness-related issues of connection-less traffic over the IEEE802.11 MAC protocol vendor implementations. It has further attempted to provide with some detailed insight in the mode of operation of WLAN protocols such as the 802.11b.

This work has looked into the issue of fairness in view of the original specification requirements of maintaining connectivity as well as rate compatibility in 802.11b WLAN networks. We have shown that the real performance of the HR extensions of the protocol is inhibited by the above requirements when a common acquisition sequence is adopted for all supported signaling rates. To this end, we have devised an extension to the current dynamic rate fallback scheme by enforcing some aspect of proportional fairness in bandwidth allocation effected over the discrete set of supported signaling rates for 802.11b. We argue that such fairness should be instigated at the lowest of the communication layers (i.e. physical) if the supported signaling rates are to provide their maximum performance.

The devised extensions allow us to argue further about a

bandwidth-(and subsequent service-)differentiated wireless Internet. We feel that bandwidth-differentiation should not be primarily supported at the IP layer but at the MAC/PHY layer over wireless. A two-tier differentiation should be the fundamental characteristic in bandwidth provisioning over wireless, and in particular WLANs. The 802.11b protocol provides an excellent candidate where such differentiation scheme can be employed. The discrete range of modulation schemes supported allow effectively the definition of discrete bandwidth allocation classes over which proportional fairness can be effected as a function of the received signal strength (RSS). Clearly, such function would manifest itself as inversely proportional to the distance between the wireless node and the base station.

We are currently devising the appropriate simulations as well as a detailed analysis of the proposed extensions over 802.11b entailing link-adaptation notifications. This will provide us with results on the performance and viability of such scheme but also act as a substrate for further service qualitative considerations that may allow PCF to act in a QoS-aware fashion.

VI. ACKNOWLEDGMENTS

This work have been sponsored by EPSRC, London, UK. We would like to thank Jon Crowcroft for their invaluable discussions and feedback during the writing of this paper. We would further like to thank Charlie Perkins and Jari Malinen from Nokia Research Center for their comments as well as the Hannu Flinck for the provision of OPNET simulation resources in NOKIAUS. We would also like to thank the UCL/MICE (aka G11) research group that has provided essential insight in aspects relative to this work. Last we would like to thank the FreeBSD community for being there and shedding light during our problem-solving endeavours.

REFERENCES

- [1] V. Jacobson, "Congestion Avoidance and Control," in *Proceedings of ACM SIGCOMM '88*, 1988, pp. 314-329.
- [2] J. Jubin and J. D. Tornow, "The DARPA Packet Radio Network Protocol," *Proceedings of IEEE*, vol. 75, pp. 21-32, 1987.
- [3] S. Seshan H. Balakrishnan, V. N Padmanabhan and R. H. Katz, "A Comparison of Mechanisms for improving TCP Performance over Wireless Links," *Proceeding of the ACM SIGCOMM '96*, Aug. 1996.
- [4] J. Kim B. Crow, I. Widjaja and P. Sakai, "IEEE 802.11 Wireless Local Area Networks," *IEEE Communications*, September 1997.
- [5] J. Sobrinho and A. Krishnakumar, "Real-Time Traffic over the IEEE 802.11 Medium Access Control Layer," *Bell Labs Technical Journal*, 1996.
- [6] P. Karn, "MACA: A new Channel Method for Packet Radio," in *ARRL/CRRL Amateur Radio 9th Computer Networking Conference*, 1990.
- [7] V. Bharghavan, "MACAW: A Media Access Protocol for Wireless LANs," in *Proceedings of ACM SIGCOMM '94*, Aug. 1994.
- [8] C. Fullmer and J. Garcia-Luna-Aceves, "Floor Acquisition Multiple Access (FAMA) for Packet-Radio Networks," in *Proceedings of ACM SIGCOMM '95*, 1995.
- [9] M. Chuah A. DeSimone and O. Yue, "Throughput Performance of Transport-layer Protocols over Wireless LANs," in *Proceedings of IEEE GLOBECOME '93*, Dec. 1993.
- [10] D. Bantz and F. Bauchot, "Wireless LAN Design Alternatives," *IEEE Network*, pp. 43-53, Apr. 1994.
- [11] R. Lamaire et al., "Wireless LANs and Mobile Networking: Standards and future Directions," *IEEE Communications*, vol. 34, no. 8, pp. 86-94, Aug. 1996.
- [12] ETSI 300 652, *ETSI TC-RES "Radio Equipment and Systems (RES); High Performance Radio Local Area Network (HIPERLAN); Functional Specification*, ETSI, July 1995.
- [13] A. Festag J. Weinmiller, M. Schallager and A. Wolisz, "Performance Study of Access Control in Wireless LANs - IEEE 802.11 DFWMAC and ETSI RES 10 Hiperlan," in *Journal of Mobile Networks and Applications*, 1997, pp. 55-67.
- [14] 802.11 WG, *P802.11 Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specification Standard*, IEEE, Oct. 1997.
- [15] G. Xylomenos and G. Polyzos, "TCP and UDP Performance over a Wireless LAN," in *Proceedings of IEEE INFOCOM '99*, 1999.
- [16] B. Noble G. Nguyen, R. Katz and M. Satyanarayanan, "A Trace-based Approach for Modeling Wireless Channel Behaviour," in *Proceeding of the Winter Simulation Conference '96*, 1996.
- [17] D. Duchamp and N. Reynolds, "Measured Performance of a wireless LAN," in *Proceedings of the 17th IEEE Conference on Local Computer Networks*, Sept. 1992.
- [18] T. Pagtzis, "Operational semantics and analysis of IEEE802.11b for connection-less IP traffic," Tech. Rep., University College London, Aug. 2000.
- [19] "Rtp tools," <ftp://ftp.cs.columbia.edu/pub/schulzrine/rtpools>.
- [20] Y. Murayama S. Yamaguchi, "DBS: a powerful tool for TCP performance evaluation," in *Proceedings of Performance and Control of Network Systems (SPIE '97@*, Nov. 1997, vol. 3231, pp. 570-581.
- [21] D. van Wykt and L. Linde, "A Turbo-coded DS/CDMA System with Embedded Walsh-Hadamard Codewords: Coder Design and Performance Evaluation," in *Proceedings of 1998 IEEE 5th International Symposium on Spread Spectrum Techniques and Applications*, 1998, vol. 2, pp. 359-363.
- [22] M. Webster K. Halford, S. Halford and C. Andren, "Complementary Code Keying for RAKE-base Indoor Wireless Communication," in *Proceedings of the 1999 IEEE International Symposium on Circuits and Systems*, 1999.
- [23] H. Chhaya and S. Gupta, "Throughput and Fairness Properties of Asynchronous Data Transfer Methods in the IEEE 802.11 MAC Protocol," in *Proceedings of IEEE PIMRC '95*, September 1995, pp. 613-617.
- [24] H. Chhaya, "Performance Evaluation of the IEEE 802.11 MAC Protocol for Wireless LANs," Msc. thesis, Graduate School of the Illinois Institute of Technology, 1996.
- [25] G. Bianchi, "Performance Analysis of the IEEE 802.11 Distributed Coordination Function," *Journal on Selected Areas on Communications*, 2000.
- [26] B. Crow et al., "Investigation of the IEEE 802.11 Medium Access Control (MAC) Sublayer Functions," in *Proceedings of INFOCOM '97*, 1997.
- [27] "The freesd project," <http://www.freesd.org/>.
- [28] L. Fratt G. Bianchi and M. Oliveri, "Performance Evaluation and Enhancement of the CSMA/CA MAC Protocol for 802.11 WLAN," in *Proceedings of PIMRC '96*, 1996.
- [29] A. Phonphoem A. Ganz and Z. Ganz, "Robust Superpoll Protocol for IEEE Wireless LANs," in *Proceedings of MILCOM '98*, 1998.
- [30] C. Andren and M. Webster, "CCK Modulation Delivers 11 Mbps for High Rate IEEE 802.11 Extension," Tech. Rep., Harris Semiconductors, May 1999.
- [31] C. Andren, "Short PN Sequences for Direct Sequence Spread Spectrum Radios," Tech. Rep., Harris Semiconductors, 1997.
- [32] D. Eckhardt and P. Steenkiste, "Measurement and Analysis of the Error Characteristics of an In-Building Wireless Network," in *Proceedings of ACM SIGCOMM '96*, Aug. 1996.
- [33] L. Lenzini G. Anastasi, "QoS provided by the IEEE 802.11 Wireless LAN to Advanced Data Applications: a Simulation Analysis," *Journal of Wireless Networks*, 2000.
- [34] H. Chhaya and S. Gupta, "Performance Modeling of Asynchronous Data Transfer Methods of IEEE 802.11 MAC Protocol," *Journal of Wireless Networks*, pp. 217-234, 1997.
- [35] M. Conti L. Bononi and L. Donatiello, "Design and Performance Evaluation of a Distributed Contention Control (DCC) Mechanism for IEEE 802.11 Wireless LANs," in *Proceedings of WOWMOM'98*, 1998.
- [36] R. van Nee et al., "New High-Rate Wireless LAN Standards," *Communications Interactive*, Dec. 1999.
- [37] Z. Velkov and L. Gavrilovska, "Performance of the IEEE 802.11 Wireless LANs under Influence of Hidden Terminals and Pareto Distributed Packet Traffic," in *Proceedings of ICPWC '99*, 1999.