Optimizing Internet Flows over IEEE 802.11b Wireless Local Area Networks: A Performance-Enhancing Proxy Based on Forward Error Correction

Luis Muñoz, Marta García, Johnny Choque, Ramón Agüero, University of Cantabria
Petri Mähönen, University of Oulu

ABSTRACT

The success of the IP and its associated technologies has led to new challenges as we try to use it more widely in everyday communications. In particular, the drive toward wireless and highly heterogeneous infrastructures supporting IP services transparently and independent of the underlying physical layer is a challenge. In this context, this article focuses on introducing an implementation of a generic performance enhancing proxy, called Wireless Adaptation Layer, and particularly its forward error control enhancement module. The error control module is a potentially important tool for achieving better UDP and TCP performance over the inherently unreliable wireless channels, and providing some adaptation for that. In order to assess the benefits and drawbacks of the selected design, we have also conducted some performance measurements over IEEE 802.11b WLANs.

INTRODUCTION

The exponential growth of the Internet and overall success of its main protocols in the 1990s and at the beginning of the new millennium has generated a need to design and develop a generic wireless communications paradigm to allow Internet operations over wireless links. There is also a need to support heterogeneous networks that include a large number of different wired and wireless technologies. Naturally, the emphasis is to try to use unchanged or enhanced protocols whenever possible. The often used slogan “to have everything over IP” in the same way as over reliable wireless links indicates and imposes several requirements for protocol stacks and link layers, with the aim of making interactions as transparent as possible. The existing protocol stacks should be enhanced in terms of the following criteria: compensation for throughput degradation, reliability, delay, and delay variation in conditions of high bit error rate and handoffs.

Many suggestions for making the TCP more efficient over wireless links exist, including a recent and promising TCP extension called Eifel as well as analyses of the TCP over cellular networks [1–4]. It is well known that, for example, TCP performance is severely degraded because of the high bit error rate in the wireless link. To cope with the wireless error problems either a transport layer or link layer approach could be followed in order to implement solutions. The link layer approach can often use the underlying radio equipment and channel knowledge more efficiently. Moreover, the link layer could provide some proprietary and stream-dependent solutions; for example, select between forward error correction (FEC) and automatic repeat request (ARQ) schemes, depending on channel conditions and traffic. Finally, wireless LANs (WLANs) have become very popular and will form one integral part of heterogeneous networking in the near future. This will lead to the need to increase the bit rates, reliability, and communications range of WLANs on the IP context. Because of this, we have studied the possibilities of making user-perceived performance better for IEEE 802.11 WLANs.

In this article we do not focus on particular protocol extensions, modifications, or similar schemes that might be applied to UDP or TCP. On the contrary, first, we briefly intro-

0163-6804/01/$10.00 © 2001 IEEE

duce a generic performance enhancement proxy (PEP) structure that has been tested in the European Union collaborative project called WINE [5, 6] to provide the platform for different enhancements. In particular, the main focus of our article is to describe and analyze the FEC module for PEP. The underlying idea has been to study whether the packet-level FEC introduced through PEP could provide extra performance for wireless Internet. Additionally, FEC-on-PEP could provide extra assistance to some other protocol enhancements. Hence, we do not claim that it is only a competitive method, but instead a potentially cooperative method. By introducing frame-level FEC we aim to provide:

- Better performance and protection for real-time traffic streams using UDP
- Making the situation easier for TCP if underlying channel conditions (bit error rates) are bad

We stress that FEC might be activated as a link-level enhancement with standard UDP and TCP, or transparently when other more "clever" enhancements (e.g., Snoop [2]) are operating. The basic paradigm is always to leave basic protocols unchanged as far as possible, and extensions should be transparent and used only when necessary; that is, major changes to core infrastructure should be avoided because they are difficult and slow to accept and implement to all hosts and routers.

We emphasize that our framework is WLANs, and more specifically IEEE 802.11b [7]. The WAL PEP is applicable to more generic cases, and FEC work also has, in our opinion, wider use, but this article and enhancements are limited strictly to WLANs.

WIRELESS LOCAL AREA NETWORKS

WLAN technology has started to become more and more widespread due to faster, cheaper, interoperable, and standard-compliant new-generation technologies. In particular, products based on the IEEE 802.11 standards family have gained wide acceptance, and networks are now available, for example, in hotels, airports, and coffee shops. Basically two standard families cover WLAN's scope more or less completely: IEEE 802.11 and European HIPERLAN. As mentioned previously, we focus on IEEE 802.11b, and this work is not directly and completely applicable to HiperLAN/2 since, for example, their medium access control (MAC) layers are quite different.

THE IEEE 802.11 STANDARD

The IEEE 802.11 standard defines the physical (PHY) and MAC layers for both access-point-based and ad hoc WLANs. The final version of the standard was approved in 1997. The PHY layer specification allows three transmission options that interact with a common MAC: direct sequence spread spectrum (DSSS), frequency hopping spread spectrum (FHSS), both operating at the 2.4 GHz industrial, scientific, and medical (ISM) band, and diffused infrared (DFIR) at the baseband. The DSSS specification supports mandatory rates of 1 and 2 Mb/s. For the FHSS and DFIR specifications, the 1 Mb/s rate is mandatory, the 2 Mb/s rate optional. As a high-rate (HR) extension of the standard, the IEEE 802.11 Working Group approved in 1999 the IEEE 802.11b specification that also supports 5.5 and 11 Mb/s data rates, applying a new coding scheme known as complementary code keying (CCK) at the DSSS PHY layer [7].

Regarding the access scheme to the channel, the 802.11 MAC layer is based on two methods: the distributed coordination function (DCF) for asynchronous contention-based access, and the point coordination function (PCF) for centralized contention-free access. The former is considered to be the default access mechanism, the latter an optional function envisaged to support collision-free and time-bounded services.

The basic access scheme, the DCF, is based on a carrier sense multiple access with collision avoidance (CSMA/CA) protocol. Before initiating a transmission, a station senses the channel to determine whether another station is transmitting. If the medium is found to be idle for an interval that exceeds the distributed interframe space (DIFS), the station starts its transmission. Otherwise, if the medium is busy (either immediately or during the DIFS), the station continues monitoring the channel until it is found idle for a DIFS. A random backoff interval is then selected and used to initialize the backoff timer. This timer is decreased as long as the channel is sensed idle, stopped when a transmission is detected and reactivated when the channel is idle again for more than a DIFS. The station transmits when the backoff timer reaches zero. Immediately following the successful reception of a frame, the receiver must initiate the transmission of an acknowledgment frame (ACK) after a time interval called a short interframe space (SIFS), which is shorter than a DIFS to avoid another station starting to transmit before the reception of the ACK in the transmitter. After that, if the station has more frames to transmit, in order to avoid capture of the channel for this station, it must wait for a random backoff time, even if the medium is sensed idle after a DIFS. On the other hand, if the transmitting station does not receive the ACK within a specified timeout, or detects another transmission in the channel, it retransmits the frame according to the given backoff rules. Note that the absence of the ACK could be due to either a collision episode or an error condition in the channel. In addition, the 802.11 MAC defines special functions such as fragmentation of frames or medium reservation via request to send/clear to send (RTS/CTS) interchange.

Although performance evaluation of these kinds of WLANs is extensively discussed in the current literature, only a few of these studies consider the HR extension of the standard [8-11]. In [9] the overhead introduced by the IEEE 802.11b MAC and physical procedures has been accurately described. The nominal peak throughput offered to the IP layer for a maximum transmission unit (MTU) of 1500 bytes across the four rates supported by the IEEE 802.11b is shown in Table 1.

The IEEE 802.11 standard defines the physical (PHY) and MAC layers for both access point based and ad hoc wireless LANs. The final version of the standard was approved in 1997.

Bit rate, R_b (Mb/s)	Nominal throughput (Mb/s)	Bit rate (%)
11	6.2	56
5.5	3.9	71
2	1.7	85
1	0.9	90

Table 1. *Nominal peak throughput for an IEEE 802.11b single–hop.*

As can be seen, 44 percent signaling rate consumption is observed when 11 Mb/s is used, mainly due to the overhead of the preamble and physical header in IEEE 802.11, which appears quite large relative to the payload transmission time. In order to guarantee compatibility with lower rates, these 24 bytes of PHY preamble and header overhead have to be transmitted at 1 Mb/s for the four supported rates, which consumes 192 μs for each data or acknowledgment frame. Since this overhead is constant for every data packet transmitted, the loss in throughput efficiency becomes more pronounced with shorter data packets. It has been recommended to implement the short preamble and header (which lasts only half the long one) defined as an optional feature within the standard [9].

Since the IEEE 802.11b specification was finalized in 1999, the 802.11 Working Group is also developing other specifications such as 802.11a for data rates up to 54 Mb/s (using OFDM at the 5 GHz band) and 802.11e for QoS and multimedia traffic support. In particular, the aim of this last specification is to add some new traffic management policies and error control mechanisms (e.g., FEC and selective retransmission) to the HR extensions of the 802.11 standard.

Error Control Mechanisms in IEEE 802.11

Traditionally, not all WLAN platforms have implemented FEC mechanisms. For example, original WaveLAN versions previous to the 802.11 standard did not include MAC-level retransmissions, making them more sensitive to frame loss and collisions. To cope with this, the IEEE 802.11 standard specifies an idle RQ scheme with implicit retransmission to deal with both collisions and wireless errors. When the receiving station receives a frame, it checks the cyclic redundancy check (CRC) field and sends an acknowledgment frame. If the sender does not receive an ACK within a predefined timeout, it retransmits the frame until it receives an ACK or throws it away after a given number of retransmissions. In the new Lucent WaveLAN IEEE 802.11 HR-compliant, four retransmissions are applied [9, 12]. This means that losing an IP datagram at the receiver implies the reception of four consecutive erroneous frames; therefore, the frame error rate (FER) does not match the IP packet loss.

Regarding throughput reduction in the presence of wireless errors, a complete and accurate characterization is given in [9, 11], where the effectiveness of the MAC idle ARQ scheme can be observed. Furthermore, it can be assumed that the measured throughput can be roughly upper bounded (maximum deviation is about 30 percent) by scaling the nominal throughput with a factor of (1-FER). This assumption fits better when the binary rate is shorter or the frame size is reduced (i.e., the lower the FER).

As will be shown below, both MAC and error control mechanisms design will have a direct impact, particularly for low signal-to-noise ratios (SNRs), on the TCP and UDP behavior when they work over such IEEE 802.11 WLAN.

The Experimental Testbed and Results

As follows, we show a set of baseline measurements based on an experimental setup that consists of two end-host Linux machines (Pentium III platforms using Red-Hat Linux with kernel 2.2.14), one acting as a mobile terminal (MT) and the other as an access point (AP). Each incorporates Lucent IEEE 802.11b PC cards with the promiscuous mode set on the kernel device driver of the receiver in order to log each incoming packet, even if it failed the CRC. Some different test scenarios are used to measure the performance of the wireless network in various conditions. Different SNRs are associated with each scenario, in order to force both the presence and absence of errors due to wireless impairments.

UDP Performance over IEEE 802.11b

The UDP is a connectionless transport protocol used to convey real-time traffic [13]. It is important to note that due to the fact that UDP is not a reliable protocol, when a datagram is lost no retransmission or other kind of error control occurs except the four retransmissions belonging to the ARQ technique implemented in the wireless platform.

		Good channel	Bad Channel	
R_b (Mb/s)	Payload (bytes)	Throughput (Mb/s)	Throughput (Mb/s)	FER
11	1500	6.071	1.259	0.724
	1024	5.001	1.2	0.674
	768	4.206	1.293	0.585
	512	3.172	1.548	0.387
	256	1.763	0.999	0.303

Table 2. *UDP throughputs for different payloads and channel conditions.*

In order to derive the behavior of UDP over the IEEE 802.11b platform, 10,000 UDP packets were sent from the MT to the AP in each repeated measurement series. Table 2 shows the results obtained for both ideal and bad channels (low SNR and no line of sight) working at $R_b = 11$ Mb/s.

TCP PERFORMANCE OVER IEEE 802.11B

TCP Fundamentals — TCP is a widespread transport protocol [13] used within the IP stack. It provides a reliable connection-oriented data exchange service between two hosts with support for flow and congestion control as well as error recovery. Several procedures have been included during its history to enhance the performance of the protocol, taking into account especially the congestion control point of view. Some of the most important extensions have been included in almost all widespread TCP versions. However, different proposals for making TCP more efficient with wireless links have not yet achieved any widespread implementation in operating systems.

The current situation with TCP can be very briefly summarized as follows. In order to ensure an appropriate transmission rate, the *slow start* and *congestion avoidance* methods are applied. Using the former, the TCP sender entity adapts its transmission rate to the network condition by incrementing the number of unacknowledged segments following an exponential trend (i.e., doubling its value each round-trip time, RTT) and starting with one segment. However, at some point, some of the intermediate routers might reach their limit, producing a congestion situation. To deal with it, the transmission rate should be reduced, and the congestion avoidance algorithm is devoted to that purpose. Some slight modifications were added to improve these procedures' behavior, taking advantage of the fact that TCP is required to send an immediate ACK (i.e., with no delay) when an out-of-order packet is received. This duplicate ACK (dupack) can give an idea about the loss of a particular segment. The *fast retransmit* and *fast recovery* algorithms use dupacks to retransmit lost segments (upon reception of a third dupack), thus avoiding reduction in the transmission rate, since the reception of dupacks is a strong indication that there is still flow between the two hosts.

The combination of the four procedures mentioned above results in the most widespread TCP type, known as TCP Reno. However, additional algorithms are added to enhance this protocol's performance, particularly the *Timestamp* option, which is used to perform an accurate retransmission timeout estimation, and *Selective Acknowledge* (SACK), which adds the capacity to acknowledge nonconsecutive segments, in contrast to the former cumulative scheme. Both options add some extra overhead in the TCP header.

TCP Throughput — In TCP both ends of communications share the radio channel, which is not a full-duplex link but a semi-duplex one (due to the MAC access scheme it is not possible to have information flowing from both hosts simultaneously), so the performance achieved should be lower than the nominal throughput. When placing two terminals close enough to ensure the absence of errors, the throughput offered by the TCP layer (e.g., to an FTP session) is shown in Table 3, which also indicates the percent of both the raw bit rate and the nominal throughput. It can be observed that around 10 percent of the raw bit rate is wasted due to transport-level acknowledgments [8].

R_b (Mb/s)	TCP throughput (Mb/s)	Bit rate (%)	Nominal throughput (%)
11	5.0	45	80
5.5	3.3	60	85
2	1.5	75	88
1	0.8	80	89

■ **Table 3.** *Maximum TCP throughput for IEEE 802.11b single-hop.*

R_b (Mb/s)	Test	Throughput (Mb/s)
11	1	2.906
	2	0.707
	3	4.488
	4	0.501
	5	4.586

■ **Table 4.** *TCP throughput obtained in five different and typical measurements for a bad channel.*

Proceeding with this analysis and in order to derive the influence of errors on TCP performance, a new set of measurements were taken during an exhaustive measurement campaign. In this case, the two hosts were placed so that the SNRs associated were low enough to ensure the presence of errors in the wireless link. Due to the great variability in the measurements carried out, a wide range of performances was obtained and is shown in Table 4. This reflects the real-life situation, where SNR variance is very high between locations and due to varying conditions in a fixed location.

After an in-depth analysis, it was derived that the most relevant factor in throughput reduction was the presence of long idle times, as can be seen in Fig. 1. The combinations of several factors led the transmitter, in some cases, to be idle for 57 s. In the aforementioned figure the segment interchange that led to this situation is also shown. A complete analysis of why such long idle periods exists is beyond the scope of this article. Basically, when a segment is unacknowledged within some specific timeout period (retransmission timeout, RTO), a retransmission is triggered. The value of the RTO is dynamically changed during a connection by taking measurements of both the mean value and the standard deviation of the RTT. However, when this timer expires and a retransmission is triggered, a backoff procedure is applied. In this sense the RTO is doubled constantly, until all the segments "in flight" are acknowledged. Fig-

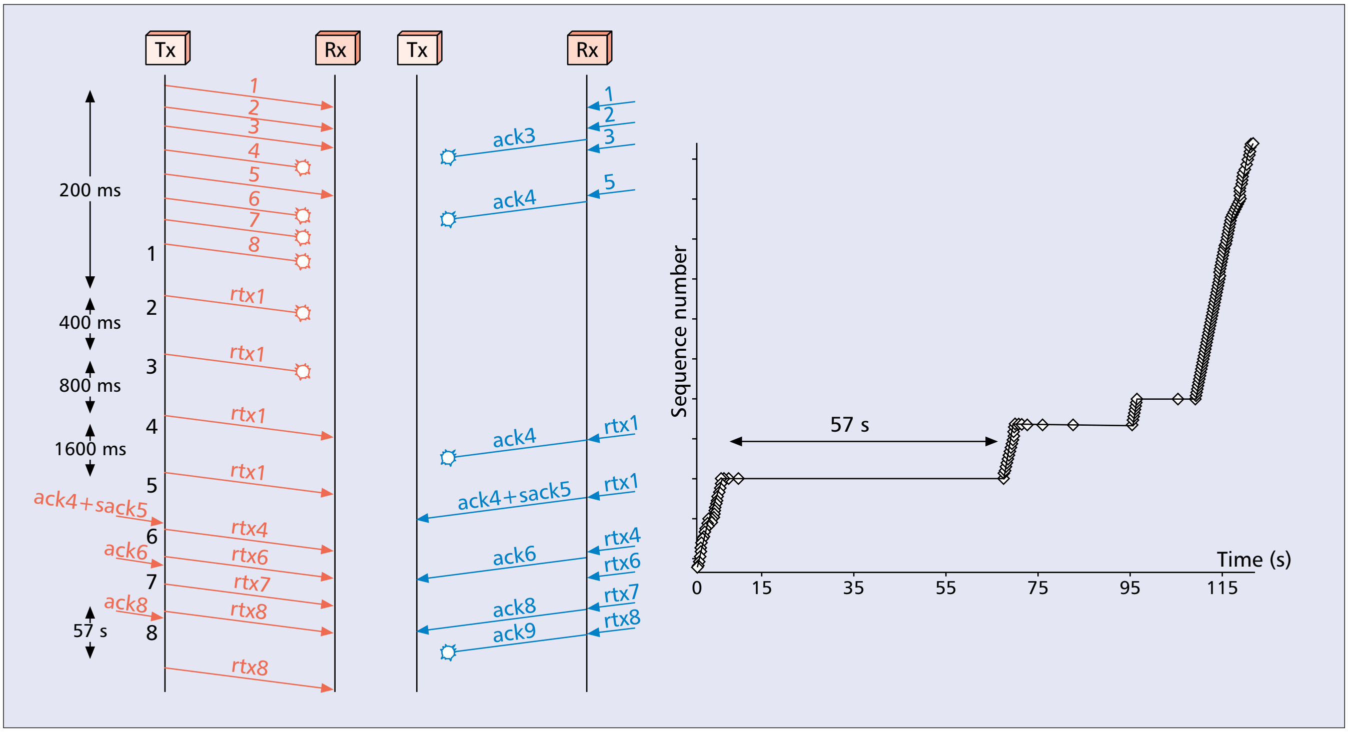

■ **Figure 1.** *TCP behavior for low SNR 802.11b WLAN.*

ure 1 shows that the first segment is retransmitted up to four times. Since at this point there are still unacknowledged segments, a backoff of five will be applied in the following retransmission, so the RTO obtained from RTT estimation will be multiplied by 32. It should be mentioned that the actual values for both the mean and the standard deviation of the RTT were quite high (310 and 580 ms, respectively). These values were obtained taking advantage of the fact that the Timestamp option was activated during the connection (a simple C program was developed to monitor RTT variation).

PERFORMANCE ENHANCING PROXIES

In the previous section the degradation of the TCP/IP stack when operating over wireless links was shown. Even though there are many proposals to increase the IP stack performance, one of the most active groups in the field of wireless Internet is the Performance Implications of Link Characteristics group (PILC).

The PILC group, belonging to the Internet Engineering Task Force (IETF), has various proposals to cope with the impairments introduced by some access networks over the Internet stack. Among them, here we are interested in the concept of the PEP [6]. In general, we can say that the PEP approach is used to improve the performance of the Internet protocols on network paths where native performance suffers due to characteristics of a link or subnetwork on the path. There are many types of PEPs that are used in different environments to overcome different link characteristics that affect protocol performance.

Below, a brief description of some key characteristics that allow for the differentiation of the types of PEPs is provided. Readers can refer to [6] where extensive examples, analysis, and description of PEPs are given:

- Layering: In principle, a PEP implementation may function at any protocol layer, but typically it functions at one or two layers only.
- Distribution: A PEP implementation may be integrated, that is, it comprises a single PEP component implemented within a single node, or distributed, that is, it comprises two or more components, typically implemented in multiple nodes. An integrated PEP implementation represents a single point at which performance enhancement is applied.
- Implementation symmetry: A PEP implementation may be symmetric or asymmetric. Symmetric PEPs use identical behavior in both directions; that is, actions taken by the PEP occur independent of in which interface a packet is received. Asymmetric PEPs operate differently in each direction. The direction can be defined in terms of the link (e.g., from a central site to a remote site) or in terms of protocol traffic (e.g., the direction of TCP data flow or the direction of the ACK flow).
- Transparency: A key characteristic of a PEP is its degree of transparency. PEPs may operate totally transparently to the end systems, transport endpoints, and/or applications involved (in a connection), requiring no modifications to any of said elements. On the other hand, a PEP implementation may require modifications to both ends in order to be used. In between, a PEP implementation may require modifications to only one of the ends involved.

It is important to remark that the PEP

paradigm allows us to introduce different PEPs depending on the requirements. Finally, PEPs can be adapted and enhanced more easily through software.

PRACTICAL IMPLEMENTATION OF A LINK LAYER PEP: THE WAL

The WAL is a layer two PEP (also distributed, symmetric, and transparent) targeted for Internet protocols when they operate on top of wireless shared access LANs. It has been studied and implemented in the European WINE project since 1999. The WAL compensates for poor wireless TCP/IP performance by hiding the channel impairments of the wireless network from the transport protocol. Rather than redesigning the existing Internet transport protocols to adapt to a wireless environment, adaptation is proposed by means of this entity residing between IP and the underlying network infrastructure. That is, the WAL, as its name indicates, adapts to the wireless network conditions to enable transport protocols to operate in the usual way. Furthermore, the WAL provides hints to the upper layers regarding the observed network conditions by SNR measurements so that its operation is adaptive to the channel state. One can get more detailed information on WAL from [5] and the WINE consortium.

In order to gain flexibility in its design, the WAL comprises a set of modules, each suitable for a particular type of traffic. Furthermore, the structure and organization of the WAL are flexible enough for both the addition of new modules and the upgrade of the algorithms the various modules are based on. One example of these modules is the forward error control enhancement module (FECEM), which will be described in the following section. As mentioned, the main focus of this article has been to use the WAL PEP as a platform for FEC and, through it, study its possibilities.

FECEM Design — The aim of a FECEM is to correct an important percentage of erroneous packets that reach the receiver to increase end-to-end throughput. After an in-depth measurement campaign, in which the error distribution for different MTUs has been obtained, a Reed-Solomon code with a maximum error correction capability of $t = 15$ symbols has been chosen. The code has been defined over the Galois Field $GF(2^{16})$, which means that each symbol is constituted by 16 bits. For an MTU of 1500 bytes, the code parameters are ($n = 750$, $k = 720$, $t = 15$); therefore, the number of parity bits added per MAC frame is 480. It is important to remark that, since the WAL is an adaptive link layer scheme, based on measurements carried out by a module devoted to that purpose, both the error correction capability and the type of code can be negotiated between the mobile terminal (MT) and the access point (AP) on the basis of the wireless channel state.

Figure 2 shows how the FEC is applied to an IP datagram coming from the upper layers. Due to the WAL design and implementation, the FEC scheme is applied to the whole WAL-PDU except for the WAL header.

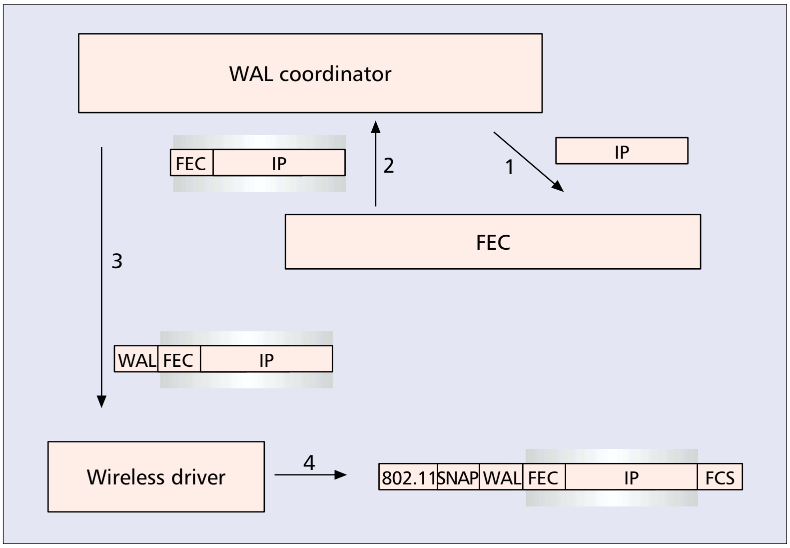

Figure 2. *IP datagram flow through an FEC module.*

The decoding process takes advantage of the CRC defined in the IEEE 802.11b standard. If no errors are detected, the FECEM will directly extract the FEC redundancy field, without calling the decoding function, thus reducing the computational overhead. However, due to the retransmission procedure implemented within the standard, the decoding function may be applied to a frame corresponding to a datagram previously corrected.

The FECEM Performance — Given that the highest IP loss has been observed using the 11 Mb/s bit rate [9], only the results obtained at this rate are presented in this article. Also, it is worth bearing in mind that FECEM performance depends on IP datagram size. The performance gain achieved for different payload sizes is shown in Table 5.

To carry out the measurements for connectionless real-time streams, once more UDP/IP traffic has been used. Besides the 28-byte UDP and IP headers, both the WAL and FEC introduce 63 bytes of extra overhead. As shown in the table, two different IP losses have been obtained. It is important to remark that both belong to the same measurement, that is, the one with FEC is extracted from the number of datagrams that arrive at the IP layer (i.e., after the FEC has been applied), and the other one is the percentage of CRC-correct frames received by the driver. In this sense, the value corresponds to the number of datagrams that would have arrived at the IP layer if the FECEM had not been activated. The difference between these IP losses is also shown in the table and represents the gain obtained by the use of the FECEM.

From these results it is derived that an important reduction of IP loss is achieved, with no degradation of throughput. For example, sending 1500-byte-length IP datagrams at the nominal rate of 6.2 Mb/s (Table 1) for a FER of 0.892, and considering that the overhead factor is (28+63)/1500, the theoretical maximum

IP datagram (bytes)	Throughput (Mb/s)	FER	IP loss (no FEC)	IP loss (FEC)	FEC gain	Correct frames	Corrected frames	Uncorrected frames	Discarded frames
1500	0.768	0.892	65.7%	41.3%	24.4%	3423	9730	13,426	7261
1024	1.138	0.751	38.6%	15.1%	23.5%	6136	9764	6424	7511
768	1.493	0.555	17.5%	0.8%	16.7%	8224	8710	1028	7128
512	1.363	0.428	12.0%	1.7%	10.3%	8798	5024	952	4119
256	0.751	0.312	6.1%	0.7%	5.4%	9390	3302	360	2942

■ **Table 5.** *A parameter report over an IEEE 802.11b single-hop with low SNR; the impact of the FECEM* (R_b=*11 Mb/s*).

throughput will be 0.629 Mb/s, which is lower than the 0.768 Mb/s measured throughput obtained with the FECEM. Furthermore, an in-depth analysis of the FECEM behavior allows verifying that besides the frames beyond FEC error correcting capability, and hence not corrected, there is an important percentage of frames that are discarded despite not having errors. As mentioned previously, there are some cases in which the FEC decoding algorithm is applied to frames associated with IP datagrams that were already corrected, due to the hardware retransmission mechanisms defined within the IEEE 802.11 standard. In this sense, a lot of the FEC efficiency is lost. A possible solution to this problem would imply accessing that retransmission procedure (i.e., firmware modifications).

Finally, the coding gain achieved by using the FECEM will be projected over the TCP throughput by correcting some of the TCP segments, and thus smoothing the effect of the idle situations identified previously.

CONCLUSIONS

The provision of IP services supported by wireless platforms with similar characteristics as over some wired platforms is one of the main aims of the IP community. Among wireless platforms, the most popular at present are those known as wireless LANs. It is foreseen that these kinds of infrastructures will play an important role in the transition toward heterogeneous IP-based networks.

In this article it has been shown that the performance of IP flows over these kinds of infrastructures is deficient when channel conditions are bad. A generic link layer PEP framework has been proposed to overcome such poor behavior. As a first step, a FECEM has been implemented and tested. The results show that the use of small MTUs with FEC improves performance, in terms of reducing the IP datagram loss. One should also note that FEC could be implemented at bit level, that is, as the part of physical layer radio implementation as is the case of many wireless networks (i.e., in GSM). However, our work with frame-level FEC is based on the fact that there are many wireless networks and WLAN standards that do not include any physical layer support for FEC at the present time. Moreover, the implementation of frame-level FEC-on-PEP enables the possibility to select different FEC schemes based on traffic conditions and link/physical layer selection. The testing of different FEC algorithms and adaptation in the WLAN context is in progress. In any case, we point out that since the processing power has increased and is relatively cheap, it is recommended to reconsider the need for FEC algorithms with WLANs.

ACKNOWLEDGMENTS

This work has been performed in the framework of the project WINE IST-1999-10028, which is funded by the European Community. The authors would like to acknowledge the contributions of their colleagues from VTT, Philips, Cefriel, University of Rome “La Sapienza,” AQL, Acorde, University of Cantabria, Intracom, University of Athens, and The Queen’s University of Belfast. We would like to thank Diego Melpignano and Zach Shelby for useful discussions and collaboration. Authors Luis Muñoz and Petri Mähönen would also like to acknowledge the partial financial support of the Spanish CICYT (project 1FD97-0960-C05-03) and Academy of Finland (Grant 50624), respectively.

REFERENCES

[1] K. Pentikousis, “TCP in Wired-Cum-Wireless Environment,” *IEEE Commun. Survey*, 4th qtr. 2000, pp. 2–14.
[2] H. Balakrishnan, S. Seshan, and R. Katz, “Improving Reliable Transport and Handoff Performance in Cellular Wireless Networks,” *Wireless Networks*, vol. 1, no. 4, Dec. 1995, pp. 469–81.
[3] R. Ludwig, A. Konrad, and A. Joseph, “Optimizing the End-to-End Performance of Reliable Flows over Wireless Links,” *5th Annual Int’l. Conf. Mobile Comp. Net.*, Seattle, WA, Aug. 1999, pp.113–19.
[4] R. Ludwig and R. H. Katz, “The Eifel Algorithm: Making TCP Robust Against Spurious Retransmissions,” *ACM Comp. Commun. Rev.*, vol. 30, no. 1, Jan. 2000, pp. 30–36.
[5] WINE Project, http://www.vtt.fi/ele/projects/wine.
[6] J. Border *et al.*, “Performance Enhancing Proxies Intended to Mitigate Link-Related Degradations,” RFC 3135, June 2001.
[7] IEEE 802.11 WG, P802.11b, “Supplement to Standard IEEE 802.11. Higher Speed Physical Layer (PHY) Extension in the 2.4 GHz Band,” Sept. 1999.
[8] A. Kamerman and G. Aben, “Net Throughput with IEEE 802.11 Wireless LANs,” *Wireless Commun. Net. Conf.*, Chicago, IL, Sept. 2000, pp. 747–52.
[9] M. G. Arranz *et al.*, “Behavior of UDP-Based Applications over IEEE 802.11 Wireless Networks,” *12th IEEE Int’l. Symp. Personal Indoor and Mobile Radio Commun.*, San Diego, CA, Sept. 2001.
[10] G. Xylomenos and G. C. Polyzos, “TCP and UDP Performance over Wireless LAN,” *Proc. IEEE INFOCOM ’99*, Mar. 1999, pp. 439–46.
[11] T. Pagtzis, P. Kirstein, and S. Hailes, “Operational and Fairness Issues with Connection-Less Traffic over IEEE 802.11b,” *IEEE ICC*, Helsinki, Finland, June 2001.
[12] S. Y. Wang, “Decoupling Control from Data for TCP Congestion Control” Ph.D. thesis, Harvard Univ., Sept. 1999.
[13] W. R. Stevens, *TCP/IP Illustrated, Volume 1, The Protocols*, Addison Wesley, 1994.

BIOGRAPHIES

LUIS MUÑOZ (luis@tlmat.unican.es) is associate professor at the University of Cantabria. He received a degree and a Ph.D. in telecommunications engineering from the Technical University of Cataluña (UPC), Spain, in 1990 and 1995, respectively. He joined the Communications Engineering Department in 1990 where he has been involved in different national and international projects, the latter within the ACTS and IST programs, related to voice and data transmission over wired and wireless infrastructures. His current research interests lie in the design and implementation of link layer control mechanisms for wireless Internet infrastructures, medium access protocols, and channel coding.

MARTA GARCIA (marta@tlmat.unican.es) is assistant professor at the University of Cantabria since 1994. She received a B.Sc. degree in physics with electronics and an M.Sc. degree in communications engineering from the University of Cantabria in 1991 and 1993, respectively. She joined the Communications Engineering Department in 1992 where she has been involved in different national and international projects, the latter within the ACTS and IST programs. Her research focuses on LLC/MAC protocol design for wireless networks as well as TCP/IP protocol performance over wireless links.

JOHNNY CHOQUE (jchoque@tlmat.unican.es) received an electronic engineering degree from the National University of Engineering, Perú, in 1995 and an M.Sc. degree in communications engineering from the University of Cantabria in 1998. Since 1998 he has been a researcher in the Communications Engineering Department of the University of Cantabria. Currently he is working toward a Ph. D. in communications engineering and carrying out research on LLC/MAC protocol design and evaluation for wireless networks.

RAMON AGÜERO (ramon@tlmat.unican.es) received a degree in telecommunications engineering from the University of Cantabria in 2001. Since 2000 he has been a researcher at the Communications Engineering Department of that university. His research focuses on LLC/MAC protocol design for wireless networks as well as TCP/IP protocol performance over wireless links.

PETRI MÄHÖNEN (Petri.Mahonen@oulu.fi) is with the University of Oulu, Telecommunications Laboratory and Center for Wireless Communications, where he leads networking research. He has also been a research professor at VTT (Technical Research Center of Finland) and founding scientific director for its Wireless Internet Laboratory. He has studied and worked in the United States, United Kingdom, and Finland in various academic posts. He has been principal investigator in different national and international projects, including large European Union research projects such as WINE and WIND-FLEX. His current research focuses on wireless Internet, low-power medium access control, ad hoc networks, broadband wireless access, and personal area networking.

Among wireless platforms, the most popular are those known as wireless LANs. It is foreseen that these kinds of infrastructures will play an important role in the transition toward heterogeneous IP based networks.