

Field trial results at DLC layer of a HiperLAN/2 prototype

Romain Rollet, Corinne Rosier, Hervé Bonneville, Christophe Mangin
Mitsubishi Electric ITE - TCL, 1 allée de Beaulieu, CS 10806, 35708 RENNES Cedex 7, FRANCE.
Tel : (+33) 2 23 45 58 58, Fax : (+33) 2 23 45 58 59
E-mail: rollet@tcl.ite.mee.com

Abstract— This paper presents the performance measurement of a Wireless LAN prototype compliant with the European ETSI BRAN HiperLAN/2 standard. HiperLAN/2 system enables bit rates from 6 Mbit/s up to 54 Mbit/s over a 5GHz wireless link. The paper focuses on the characterization of DLC parameters like delay and throughput in realistic transmission conditions with simulated data application

I. INTRODUCTION

Wireless Local Area Networks (WLANs) have recently experienced a rapid development mainly due to the increase of bandwidth they can offer, enabling their deployment in various areas. Unlike current major standards unable to guaranty specific QoS to multimedia services, HiperLAN/2 (H/2) embeds QoS features such as a connection-oriented Medium Access Control (MAC), centralized resource allocation, enhanced Link Adaptation (LA), efficient packet transmission scheduling and low latency retransmissions.

This paper presents performance measurements of a H/2 hardware prototype in realistic transmission conditions. It is organized as follows: Section II gives an overview of the standard. Section III presents the modem prototype, focusing on the specific parts characterized by the measurements, mainly Error Control (EC) retransmission scheme and MAC scheduling policy. In Section IV, performance throughput over an error-free channel is analyzed. Throughput and latency performance over a noisy channel are next evaluated, using different EC modes. Finally, Section V concludes and discusses some future work.

II. H/2 STANDARD OVERVIEW

H/2 standard specifies the air interface between an Access Point (AP) that centralizes access to the medium and Mobile Terminals (MTs). Figure 1 depicts the protocol stack composed of three layers: the Physical layer (PHY), the Data Link Control (DLC) layer, and a set of Convergence Layers (CLs).

A. Physical layer

The PHY layer [3] is based on an Orthogonal Frequency Division Multiplexing (OFDM) modulation scheme, where information is sent in PHY bursts. Each burst consists in a fixed preamble followed by data encoded in a given number of OFDM symbols. Several bit rates, resulting of the combination of different modulations (BPSK, QPSK, 16QAM

and optionally 64QAM) and coding rates, are available. Corresponding PHY modes are selected on a per connection basis depending on radio channel conditions and the quality of service required by the connection.

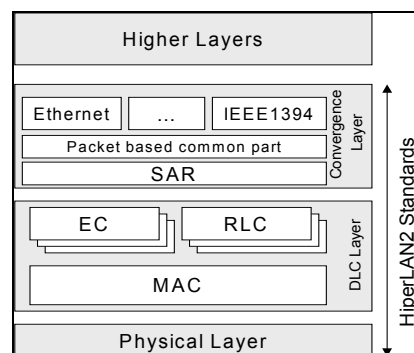


Figure 1: H/2 protocol stack overview

B. Data Link Control layer

The DLC layer [1] includes the Radio Link Control (RLC), the EC and the MAC sub-layers. The RLC, detailed in [2], manages radio resources, terminal association and DLC user connections. The MAC controls resource allocation and access to the medium. The EC performs detection and correction of transmission errors and ensures in-sequence delivery of data to the upper layers.

1) Medium Access Control

The MAC is based on a Time Division Multiple Access (TDMA) / Time Division Duplex (TDD) scheme with fixed duration frames (2 ms) built by the AP. As illustrated in Figure 2, each frame starts with a Broadcast phase including a Broadcast CHannel (BCH) that contains information about the cell, a Frame CHannel (FCH) that announces frame's composition and an Access feedback CHannel (ACH), which provides acknowledgements to access attempts made by MTs in the Random access CHannel (RCH) of the previous frame. The DownLink (DL) and UpLink (UL) phases contain control data and user data, respectively from the AP to the MTs, and from the MTs to the AP. The RCH is the only part of the frame with a contention scheme for initial communication. Access within the UL and DL phases is connection oriented and relies on two types of MAC-PDUs: 54-byte "Long PDUs" (LCH) with a 48-byte payload, and 9-byte "short PDUs" (SCH). LCHs are used to transport user data, whereas SCHs are used to convey signalling. MAC-PDUs of connections attached to the same terminal are grouped into a single so-

called “PDU train”, which in turn is mapped on a PHY burst.




Figure 2: MAC frame structure

The AP is in charge of allocating resources to the connections. The MTs are informed via the FCH of the position of their allocated time intervals in the MAC frame.

2) Error Control

The measurements focus on two EC modes provided by H/2: unacknowledged and acknowledged modes. The first mode is transparent with no delivery guaranty, providing an unreliable link with short transfer delay. The second mode implements an Automatic Repeat reQuest (ARQ) protocol providing reliable connections using retransmissions. An optional discard mechanism may be implemented to bound transfer delays.

C. Convergence Layer

The CL, through its Common Part Convergence Sub-layer (CPCS) and Segmentation And Re-assembly (SAR) sub-layer, is responsible for mapping service requests and data from upper layers onto the service offered by the DLC.

The CPCS is in charge of formatting the variable size upper layer packets into Service Data Units (the CPCS-SDUs) handled by the SAR sub-layer. As shown in Figure 3, it mainly adds packet length information in a trailer and inserts the required padding so that the obtained CPCS-PDU has a length multiple of 48 bytes. Upon emission, the SAR sub-layer segments the CPCS-PDU into units of 48 bytes that become MAC-PDU payloads. Upon reception, the SAR collects the MAC-PDUs delivered by the DLC, drops padding and checks CPCS-SDU consistency using the trailer before passing it to the upper layer.

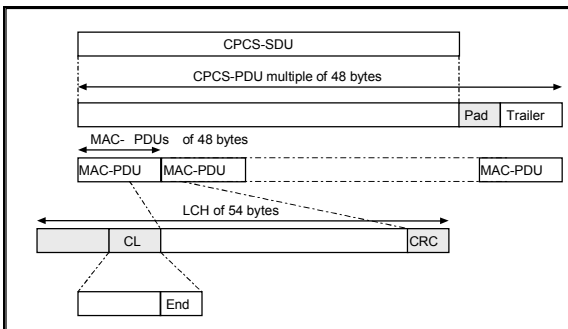


Figure 3: Segmentation And Reassembly

III. DLC PROTOTYPE PRESENTATION

We developed a full H/2 prototype for AP and MT. Both prototypes use the same hardware platform constituted of four boards: Radio Frequency (RF) at 5GHz, Intermediate Frequency (IF) at 140MHz, OFDM (Base Band), and DLC, as shown in Figure 4.

The DLC board implements MAC, EC (acknowledged and

unacknowledged modes), RLC and SAR functions. Upper layers and non real-time functions are embedded into a H/2 driver running on a Linux PC host. MAC and EC are detailed hereafter.

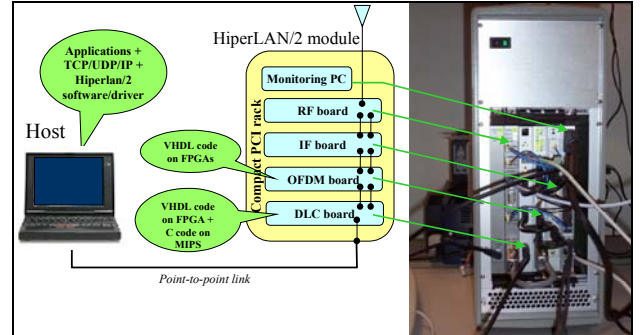


Figure 4: HiperLAN/2 prototype

A. MAC Scheduling

The aim of the scheduler is to share the physical link among DLC connections. Two types of schedulers are implemented depending on the QoS required by the connections:

- a “rigid” scheduler based on a Virtual Scheduling Algorithm [4] guarantees a peak Constant Bit Rate to the connection.
- an “elastic” scheduler based on a Self Clocked Virtual Scheduling Algorithm [5] allocates resources to connections proportionally to their weight.

When allocating resources, priority is given to rigid connections over elastic ones.

LCH resource scheduling is fully dynamic and depends on each connection’s QoS parameters, on the state of connection queues, and on available radio resources. In UL, the terminal advertises its connection queue state via Resource Request (RR) messages. SCHs used to transport control information such as ARQ FeedBack (FB) or RR are also dynamically scheduled. SCH scheduling has precedence over LCH scheduling.

As the frame composition has to be announced in the FCH, scheduling is performed one frame ahead, which introduces a 1 ms mean delay between an RR and the corresponding resource grant. Scheduler inputs are updated each time resource is requested for a connection. However, no resource is allocated for a connection when its queue is empty or when its ARQ window stalls. Furthermore, DL and UL RRs are performed in a different way.

For DL connections, LCH resources are requested internally in the AP when packets are received from the upper layer and ARQ window progresses, or when retransmission is needed. For initial transmission, only resources for PDUs that can enter the ARQ window are requested in order to avoid unused allocated slots due to stalled window. One FB SCH for the reverse direction is allocated each time n LCHs have been sent.

For elastic UL connections, the terminal indicates to the AP (via RRs) the amount of LCHs requested for new PDUs that can enter the ARQ window plus LCHs for retransmission.

For rigid UL connections, the AP allocates the guaranteed amount of LCHs except when the ARQ window has stalled.

The terminal sends RRs only when extra LCHs are requested for retransmission. This mechanism decreases connection delay and reduces signalling overhead. For each UL connection, one or more FB SCHs for the reverse direction are dynamically allocated, based on the distance between the bottom and the top PDU in the ARQ receiver window.

B. ARQ Scheme

The ARQ protocol relies on a Selective Repeat ARQ with Partial Bitmap (SR-PB). In this scheme, both the sender and the receiver maintain a window to keep track of MAC-PDU numbers. Five different window sizes are possible: 32, 64, 128, 256 and 512 PDUs. Retransmissions of errored PDUs are triggered upon reception of FB messages sent by the receiver. Each FB message may contain up to three different bitmap blocks of positive or negative acknowledgements (ACK / NAK). The receiver may also send a cumulative ACK that makes the transmitter window progress.

At the receiver side, the policy consists in generating feedback advertising in priority the oldest errored PDUs to ease the progression of the transmitter window. Acknowledgements sent in a given phase (UL or DL of frame n) may cover MAC-PDUs received in the previous phase (DL of frame n or UL of frame $n-1$). Therefore, the transmitter is able to optimize the use of the granted slots in the next frame according to the most up-to-date information.

At the transmitter side, priority is given to retransmissions rather than initial transmissions. PDUs are retransmitted in increasing order of their sequence number, so that the most aged PDUs are re-sent first. Extra resources for retransmission are allocated upon feedback reception, based on the number of NAK notifications. In addition, a timer protects each connection from possible lack of feedback.

IV. PERFORMANCE RESULTS

All our measurements were performed using Ethernet flows mapped onto a single DLC connection of different types (rigid or elastic, acknowledged or unacknowledged). Data was sent either in DL or UL using the configuration described in IV.A. In sub-section IV.C, we first analyse the maximum throughput of a single connection over an error-free channel. Then throughput and latency performances are evaluated for an elastic UL connection over a noisy channel, using acknowledged and unacknowledged modes.

A. Measurement test bed and protocol

System configuration used for measurements is composed of one MT and one AP, as shown in Figure 5. For the sake of simplification, all measurements have been performed at the IF level (RF is bypassed) by using a H/2 channel simulator coupled to a noise generator. This configuration reproduces realistic, yet controlled conditions. Only one direction (DL or UL) can be affected at a time by the simulator.

Pseudo-BRAN fading channel C (restricted to six taps) described in [6] has been used to perform measurements in erroneous channel conditions.

Packet Error Rate (PER), at DLC layer level, was chosen to characterize the DLC performance. PER is averaged over a 60

second time period for a given level of noise.

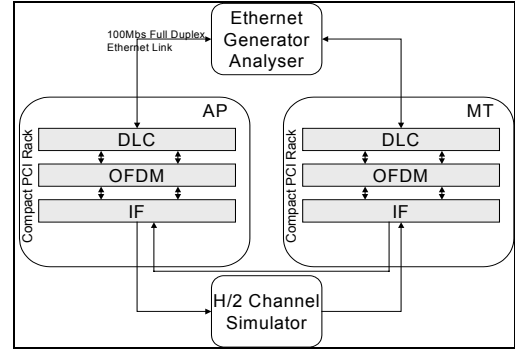


Figure 5: Test bed configuration

Traffic is emitted by an Ethernet generator/analyser at a constant packet rate and a fixed packet size of 1514 bytes. To reach the maximum PHY rate of 54 Mbit/s, LCHs are encoded in 64-QAM. SCHs are encoded in BPSK $\frac{1}{2}$, so that errors on FB and RR messages are negligible. Re-assembled packets are collected above the SAR by the Ethernet analyser. We assume that DLC connections are opened with suitable parameters and that no error occurs in the SAR. In addition, ARQ PDU discarding is not enabled.

To determine system's throughput performance and its limitations under ideal conditions, a first campaign of measurements was performed over an error-free link for every ARQ window size. Throughput was compared to a theoretical limit analytically determined as follows: $Min(K_s * p / RTT, Th_m)$ (1), with K_s the window size, p the fixed-size payload (48-byte) of a PDU, and Th_m the available throughput with modulation m . In 64-QAM, $Min(K_s * p / RTT, Th_{64}) = Th_{64}$ for $K_s \geq 256$. Performances in erroneous channel conditions were next measured by running the elastic scheduler applied to a single DLC connection.

B. Implementation & standard constraints

ARQ performance is widely influenced by the scheduler behaviour. As explained in sections above, the scheduler was designed to optimise resource allocation and therefore only schedules LCHs in the limit of the ARQ window capacity.

As shown in Figure 6, for DL connections, resources are allocated for feedback in UL once the ARQ transmitter has sent a preset number of LCHs. As the frame is scheduled one frame ahead, the ARQ receiver inserts feedback with a 2 ms delay. Upon reception of an FB SCH that allows window progression, the transmitter requests resources for PDU transmission/retransmission. The scheduler then grants the corresponding LCHs in the next frame.




Figure 6: LCH and SCH allocation in downlink

Such a mechanism may lead to a 4 ms mean RTT. For example, when K_s LCHs are granted in the same frame, the

transmit window stalls at the end of the DL burst and the scheduler can not allocate any LCH for first transmission until an FB message cumulatively acknowledges some PDUs. As FB resources are requested only during LCH transmission, K_S LCHs may be emitted only every two frames (see Figure 6). For UL connections, FB SCH is requested upon PDU reception and is allocated in the next frame. As the DL phase precedes the UL phase, a 2 ms delay is theoretically reachable. However, for elastic connections, RRs are needed to notify resource requirements and as H/2 standard imposes to place SCHs after LCHs in an UL burst, RRs are transmitted after user data, as depicted in Figure 7. If the channel is error-free, the RR contains only the number of PDUs waiting for initial transmission that can enter the ARQ window.




Figure 7: LCH and SCH allocation in uplink

Assuming the connection queue is always full, the amount of requested LCHs should then be equal to $K_S \cdot N_{Tx,LCH}^{(n)}$, with $N_{Tx,LCH}^{(n)}$ the number of PDUs sent in frame n . For window sizes smaller than 256, as the scheduler grants all the requested LCHs in the next frame, we have $N_{Tx,LCH}^{(n+1)} = K_S \cdot N_{Tx,LCH}^{(n)}$ and thus $E[N_{Tx,LCH}] = K_S/2$, which leads to an average apparent RTT of 4 ms. For ARQ windows greater or equal to 256 PDUs and high bit rate traffic, large UL bursts are used to transport LCHs, and RR SCHs end up in the end of the frame. In that case, when RR is analyzed by the AP, the time remaining to the scheduler to prepare next frame's allocation is likely too short to schedule all the requested capacity, and a part is deferred one frame later, resulting in an RTT ranging from 4 to 6 ms.

C. Measurement results

Figures 8 and 9 present the maximum throughput available for DL and UL connections over an error-free link for different ARQ window sizes, using respectively the elastic and the rigid schedulers.

In both histograms, a grey bar is added for reference that indicates the maximum theoretical rate computed with (1). Computed UL and DL rates are identical for ARQ window sizes < 256 . Above 256 PDUs, rates are actually limited by the frame's capacity and differ due to the frame structure and turn-around times. Maximum MAC throughput has been evaluated to 44,5 Mbit/s in DL versus 41 Mbit/s in UL. Only the highest capacity in DL is represented.

The maximum throughput in unacknowledged mode (window size = 0) is also provided, to characterize limitations introduced by ARQ in perfect conditions.

Unacknowledged mode performances show the high efficiency of H/2 standard in terms of frame's occupancy. Note that thanks to the fixed-size nature of the MAC-PDUs, maximum throughputs do not depend on upper layer's packet size.

In acknowledged mode, measurement results do not reach so

closely the theoretical rates, except for the largest window size. Overall throughput performances are better for rigid connections than for elastic ones, especially in UL. This is mainly due to the scheduling process, as explained below.

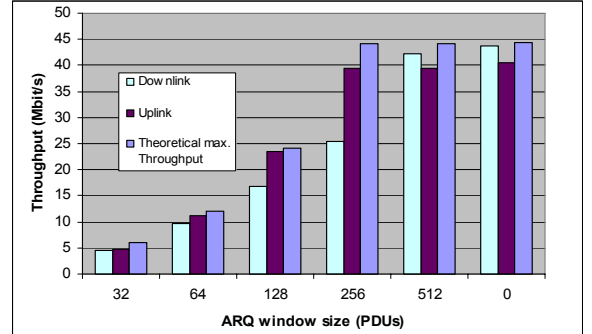


Figure 8: SAR Throughput vs Window Size with rigid scheduler




Figure 9: SAR Throughput vs Window Size with elastic scheduler

For elastic DL connections with a window size less than 256, scheduling process introduces a 4 ms RTT (see IV.B). For rigid DL connections, three cases can be distinguished: small windows (< 256 PDUs) experience an RTT close to 2 ms because the scheduler allocates a number of resources in accordance with connection's bandwidth parameter, so that the transmit window never stalls and FB SCHs are requested every frame. Large windows face an RTT between 4 and 6 ms (see IV.B), but, contrary to small windows, throughput for the 512-PDU window does not drop to half the theoretical one, thanks to the oversized window. Besides, the particular case of the 256-PDU window illustrates a side effect of optimisation: during the large DL burst period, only a limited number of LCHs can be pre-allocated due to the relative small window capacity. Since the FB-SCH is received and analysed at the end of the frame, the scheduler is not able to allocate as many LCHs as the window progression could allow. This explains why the results are not as good as for a 512-PDU window, where pre-allocation of LCHs is always possible, as the window never stalls. The elastic connection with 256-PDU window faces the same issue, and so similar results are obtained in DL.

For rigid UL connections, maximum throughput is reached that almost matches a 2 ms RTT as expected. The difference with theoretical values is due to the suspension of resource allocation when the transmitter window stalls. For elastic UL connections, the terminal has to wait to be polled for its bandwidth requirements before being able to transmit any data

and thus mechanism depicted in IV.B also occurs, leading to a 4 ms RTT. Moreover, when window size reaches 256 PDUs and injected traffic increases accordingly, the lack of time for allocation computation occurs from time to time like in DL and leads to a mean RTT of nearly 6 ms.

Figure 10 illustrates the maximum UL and DL throughput versus PER in acknowledged mode for channel C. Results show that DL and UL performance degrades almost in the same proportion as PER increases.




Figure 10: SAR Throughput vs PER, for 1514-byte frames

Since ARQ does not include PDU discarding, the observed throughput degradation is due to Ethernet frame losses, resulting from input queue overflows prior to ARQ. It has to be noted that this graph does not illustrate ARQ efficiency but demonstrates system's behaviour in hardly stressed conditions. It provides the maximum data rate that can be supported by the transmission system considering a given PER. Results show that the system still offers good performances with PERs up to 10%.

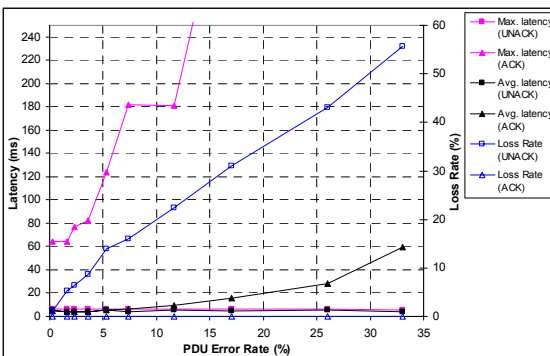


Figure 11: Transfer delays / Loss rates vs PER, for 1514-byte frames

Figure 11 shows transfer delays and loss rates in UL versus PER over channel C, for both implemented EC modes. The two connections emit data at 6 Mbit/s. In addition, the acknowledged connection uses a 256-PDU ARQ window size. Since ARQ does not bound the retransmission number, the system is able to provide perfectly reliable data to such connections as long as the traffic bit rate is compatible with the window capacity, resulting in no loss whatever the PER is. It is however at the expense of increased delays.

The unacknowledged mode on the contrary fulfils its purpose of transparency. Indeed, the two plots of maximum and average latency are linear and almost superimposed. Even at very high PER, the maximum delays remain constant and

bounded within 6 and 7 ms, while average latency remains close to 4,5 ms. Note that the unacknowledged connection suffers relatively low loss rates compared to the ones that would be obtained if errors were randomly distributed, which is consistent with the bursty behaviour of the radio channel.

However, the unacknowledged mode can not compete with an ARQ able to offer perfect reliability to loss-sensitive connections, even in very bad radio conditions. Furthermore, we can consider that latency in acknowledged mode remains acceptable even for real-time applications. Indeed, for a maximum PER up to 11,6%, delay values remain under 200 ms, while average delays are inferior to 10 ms and 50% of the frames are delivered in less than 5 ms. This latter ratio reaches 85% when the delay bound is brought to 10 ms.

V. CONCLUSION AND FUTURE WORK

For all the connection configurations considered in this paper, i.e. UL, DL, rigid, elastic, acknowledged and unacknowledged, the results in error-free channel conditions conform to the assumptions initially made (refer to IV.B). They highlight the influence of a dynamic SCH and LCH resource allocation policy on DLC performances, especially for ARQ connections, and also show some limitations introduced by the rigid frame organisation defined in H/2 standard (specific ordering of DL, UL, SCHs and LCHs transport channels). In acknowledged mode, the performance limitations mainly come from the fact that resources are granted in the limit of ARQ window capacity, to optimize the link utilization. However, this side effect can be corrected by using a larger ARQ window.

Over an error-prone link, measurements made in the context of a BRAN-type fading channel C show that ARQ such as the one implemented can be successfully applied to a wide range of applications, including real-time traffic with strong delay constraints. Additional mechanisms such as ARQ Discard may be used to limit the delay in accordance with connection's QoS requirements.

ACKNOWLEDGEMENTS

Thanks to Isabelle Moreau, Gwillerm Froc and David Mentre for their review of this contribution, and to Arnaud Bouttier and Bertrand Penther for helping out with the measurements.

REFERENCES

- [1] ETSI TS 101 761-1 V1.3.1, "Broadband Radio Access Networks (BRAN); HIPERLAN Type 2; Data Link Control (DLC) Layer; Part 1: Basic Data Transport Functions", January 2002
- [2] ETSI TS 101 761-2 V1.3.1, "Broadband Radio Access Networks (BRAN); HIPERLAN Type 2; Data Link Control (DLC) Layer; Part 2: Radio Link Control (RLC) sub-layer", January 2002
- [3] ETSI TS 101 475-2 V1.3.1, "Broadband Radio Access Networks (BRAN); HIPERLAN Type 2; Physical (PHY) Layer; October 2001
- [4] P. Boyer, F. Guillemin, M. Serval, J.-P. Coudreuse, "Spacing cells protects and enhances utilization of ATM network links," in *IEEE Network Magazine*, vol. 6, no. 5, pp. 38-49, September 1992.
- [5] J. Roberts, "Virtual spacing for flexible traffic control," in *Int'l Journal of Communication Systems*, vol. 7, pp. 307-318, 1994.
- [6] ETSI EP BRAN#9 July 1998, "Criteria for Comparison"; 30701F – WG3 PHY subgroup; Mai 1998