Design and Performance Evaluation of an Hybrid Reservation-Polling MAC Protocol for Power-Line Communications

G. Bianchi[1], I. Tinnirello[1], G. Conigliaro[2]

[1] Università di Palermo, Dipartimento di Ingegneria Elettrica

Viale Delle Scienze, Parco D'Orleans, 90128 Palermo, Italy

bianchi@elet.polimi.it, ilenia.tinnirello@tti.unipa.it

[2] Fiat Auto S.p.A. – Ingegneria di Prodotto – Sistemi Elettrici Elettronici

Corso Settembrini 40, 10100 Torino, Italy

giconigliaro@libero.it

Abstract

This paper presents an hybrid reservation-polling MAC protocol able to effectively deal with the channel disturbances typical of Power-Line Communication systems. Channel access is controlled by a scheduler running on a centralized station. Fast retransmission of corrupted packets is employed to achieve high performance on Power-line channels characterized by frequent error bursts due to impulsive noise. The performance of our proposed mechanism are evaluated by means of simulation, for the more challenging case of time-sensitive traffic sources (voice calls). Results show that the system efficiency is virtually not affected by lightly to moderately disturbed PLC channel. Moreover, by slightly reducing the system utilization, a target 1% packet loss ratio can be met even in the presence of severe channel disturbances.

I. INTRODUCTION

Power Line Communication (PLC) is a quickly developing technology, aiming at the utilization of low-voltage power lines for the transmission of data. Since wires exist to every household connected to the low-voltage grid, PLC systems can provide new opportunities for mass-market provisioning of local access at a reasonable cost. Moreover, PLC is an effective playground for the deployment of new value added services for the utilities, (e.g. automatic remote electricity meter reading, energy management, appliance control and maintenance, home automation, etc.), which are difficult to be deployed on by other technologies. The liberalization in both electricity and telecommunication areas, currently being carried out in Europe, represents a dominant driving force in the development and exploitation of PCL technology.

The design of a suitable Medium Access Control (MAC) scheme for PLC is a crucial issue. In PLC systems, a possibly large number of users (Access Units) in a low-voltage power grid, share the transmission capacity of the power-line network. To make PLC systems competitive with other access technologies is necessary to ensure Quality of Service (QoS) control, meanwhile providing a good network utilization.

We propose an Hybrid Reservation-Polling (HRP) MAC protocol, where access to the channel, for stations that have reserved via a random access procedures, is dynamically managed by a centralized controller. The dynamic resource allocation allows to fast-retransmit corrupted packets, therefore improving the packet error ratio in the presence of highly disturbed channels.

Our proposed solution has several motivations and advantages. First, a PLC network access structure is inherently a centralized system, where a privileged station (usually

placed in the transformer station MV/LV) already performs a plethora of functions [1], including Gateway to the backbone. Moreover, most of the communication needs occur between the Access Units and the Gateway. Therefore, a centralized channel access mechanism appears to be the natural choice for PLC.

Second, future PLC networks will support services ranging from best effort data transmission to real-time communication such as voice and compressed video. Traditional random access protocols appear unable of providing integration of heterogeneous services with widely different Quality of Service Requirements, especially since the PLC network size is fairly large with respect to the available channel rate (e.g. 250 users over a 2 Mbps LAN [2]) .

Third, an accurate analysis of the impact of noise in PLC [3, 4, 5] shows a large influence of short duration impulsive disturbances (up to 40 dB above the background noise). Rather than Forward Error Correction, a much more effective mean to cope with impulsive disturbances is to rely on fast retransmission of corrupted packets.

The rest of this paper is organized as follows. Section 2 describes the hybrid reservation-polling MAC scheme. Section 3 details the dynamic resource management and the fast-retransmission mechanism. The performance evaluation of HRP is carried out in section 4, for different PLC noise models. Closing remarks are given in Section 5.

II. MAC PROTOCOL DESCRIPTION

The physical architecture of a PLC network is depicted in Fig. 1 . A small number of PLC links (typically 3-5) depart from the MV/LV transformer station, and expand in a tree-like topology toward the final residential customers. Transmissions occurring on different sub-trees are received at the MV/LV transformer station via different links, and thus do not compete for channel access. Conversely, simultaneous transmission

within a sub-tree may occur, and thus a suitable Medium Access Control (MAC) protocol needs to be employed.
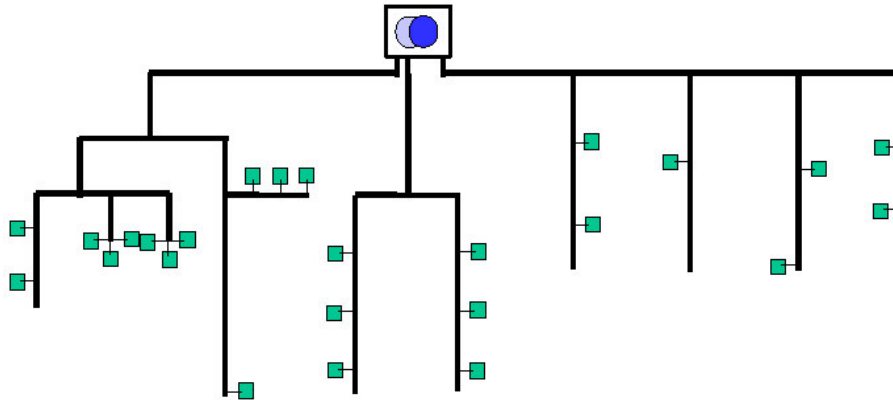

Fig. 1 - PLC network lay-out

The general ideas underlying the Hybrid Reservation-Polling protocol proposed in this paper originate from a MAC protocol suggested for wireless communications [6]. As in Time Division Multiple Access (TDMA) schemes, the power-line channel is divided into fixed-size slots. Slots are further organized into frames (Fig. 2-a). This organization allows to include, in the bytes composing the frame delimiter F, information useful for synchronization and other physical layer purposes.

In traditional TDMA systems, the frame structure has a fundamental role, as it allows to assign an index to every slot, based on its position within the frame. A given slot is allocated to a data transfer session (connection) upon setup, and it remains reserved for the whole duration of the connection. Conversely, in our scenario, the framing structure does not provide any additional information at the MAC layer. In fact, unlike the fixed slot allocation approach of traditional TDMA systems, in our scheme slot assignment within a frame is dynamically managed by a central scheduling algorithm placed at the

transformer station (MV/LV), hereafter referred to as Base Station (BS). Not only any slot, regardless of its position in the frame, can be assigned to any data transfer session, but, as shown in Fig. 2-a, the central scheduler running at the BS is in charge of deciding whether a slot is used to transfer data (and in this case, in which direction), or it is used for other purposes (specifically, reservations).

Three different types of slot can be specified: Downlink Transmission slots, to transmit data from the BS to AUs; Uplink Transmission slots, to transmit data from an AU to the BS or to another AU within the PLC network, and Reservation slots. The type of slot is specified via a "command" C, transmitted at the begin of the slot (see Fig. 2-b,c,d). The command specifies the Access Unit (AU) identifier, to which the command is addressed, and the type of subsequent data exchange within the slot. A number of bits in the command C are reserved for acknowledgement purposes. We now discuss into details the usage of each of the three slot types.
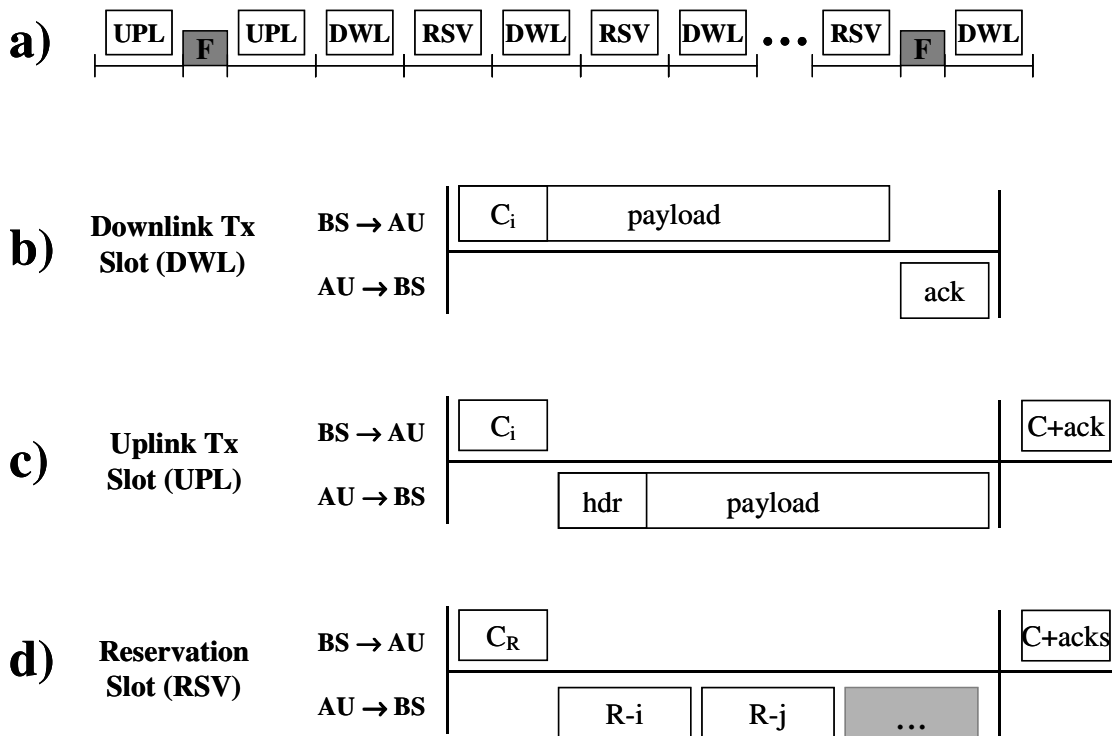
Fig. 2 - Slot organization and structure

A. Downlink transmission slot

Among the three slot types, the structure of a downlink transmission slot is the simplest, and is depicted in Fig. 2-b. The subscript i for C_i means that, in the command field, the BS specifies the identifier i of the recipient AU. Transmission of the information payload immediately follows the command. At the end of the packet reception, the AU reacts with an immediate acknowledge. By detecting a successful acknowledge, the BS is able to determine whether the transmitted packet was correctly received, and take the related actions (schedule transmission for the next packet or reschedule retransmission of the corrupted packet). As shown in the following, the capability to fast retransmit a corrupted packet (capability made possible by immediate acknowledgements) is the key to achieve high performance even in the presence of severely disturbed channel conditions and services with strict delay requirements. The price to pay for an immediate acknowledgement is a reduced slot efficiency[1].

B. Uplink transmission slot

Uplink transmission slots are used to grant transmission opportunities to the AUs, among those reserved (see next section), which have more urgent need. Within an Uplink Transmission slot, packet transmission takes place in command/response way. The BS transmits the command to the selected AU, and the "commanded AU" responds transmitting its packet, along with a MAC header which specifies the identifier of the

---

1. Because of a number of factors (among which propagation delay, receiver to transmitter turnaround time, internal computational time) some time is wasted between the end of the information payload and the begin of the ack transmission. Similar inefficiencies are present also in the case of uplink and reservation slots. As these inefficiencies depend on the technology employed in the device implementation, their analysis is out of the scopes of the present paper. It is worth to remark that this problem can be greatly reduced by designing more complex MAC details (e.g. piggybacking the ack for reservation and uplink slots not in the immediately following command, but in a one-slot delayed one,

transmitting station, as well as the identifier of the recipient station (which can be either the BS or another AU on the same PLC network sub-tree – see Fig. 1).

This operation requires the BS to be explicitly involved not only in the transmissions addressed to the BS itself, but also in the local transmissions within the PLC network. However, this additional complexity in the access management allows a contention-free channel access. It is also worth to note that the presence of the reservation phase before the issue of a transmission grant guarantees that each "commanded AU" will always have a packet available for transmission. This avoids to waste channel capacity, which is a critical issue in traditional polling schemes.

The management of acknowledges differs, depending on the type of communication. In the case a packet is sent from the AU to the BS (by far the most frequent case), the acknowledgement is piggybacked in the next command issued by the BS (in general for a different AU). Conversely, in the case of packet transmission within the PLC network (i.e. between two different AUs), an explicit acknowledgement, similar to the downlink case, must be necessarily transmitted at the end of the slot (this case is not illustrated in Fig. 2).

C. Reservation slot

Reservation slots provide a random access channel that allow AUs to "reserve", i.e. to notify the scheduler running at the BS that there is data, at the AU, that needs to be transmitted. As soon as an AU becomes active, it enters a "contending state". Reservation slots are subdivided into $m>1$ mini-slots (Fig. 2-d illustrates the case $m=3$), each capable of accommodating a reservation packet. The subdivision into mini-slots is made possible by the fact that the reservation packet payload is short. At each

---

delaying/aggregating downlink slot acks, etc). Clearly, MAC flexibility and performance will be slightly

Reservation slot available on the channel, AUs in the contending state transmit their reservation packet in a randomly selected mini-slot. In case of collision, this process is repeated until the reservation packet is successfully received by the BS. Once a reservation mini-packet is correctly received, the corresponding AU is inserted in the polling list at the BS. In the command following the reservation slot, the BS issues a reservation acknowledge. This acknowledge is efficiently implemented as a pattern of *m* bits, where bits set to 1 indicate that a successful reservation has occurred in the relevant mini-slot.

The rationale for using mini-slots is to increase the efficiency of the reservation channel. In fact, the reservation traffic load may eventually become relevant, especially if (as we recommend, to improve statistical multiplexing – see results in section 4), reservations occur on a per-traffic-burst basis rather than on a per-session basis. Clearly, the greater the reservation traffic load, the greater the value *m* required to achieve stable conditions.

III. DYNAMIC RESOURCE MANAGEMENT AND FAST RETRANSMISSION

The proposed HRP scheme is characterized by a flexible channel management. The central scheduler is in charge of deciding whether a channel slot is dedicated to uplink or downlink transmission, or reservation. Such a decision is dynamically taken on the basis of the actual traffic load; for example, in the presence of traffic asymmetries typical of access networks, more downlink slots are delivered than uplink transmission grants.

A very important feature of the proposed scheme is the ability to allow fast retransmission of corrupted packets. Usually, in communications systems, error

---

reduced with respect to the scheme described in this paper.

detection and retransmission are taken care in the upper level of the Data Link Layer. When this layer discover a corrupted packet, it sends a retransmission request to the MAC level, undergoing the same access delay as a new packet. As a consequence, the retransmission delay is high and such a recovery mechanism cannot be considered for delay-sensitive services, such as voice or video-conferencing. Transmission is generally protected from channel noise via forward error correction (FEC). Unfortunately, PLC noise is impulsive [3, 4, 5]: in the presence of a noise burst (which can last up to a couple of ms), all the signal is destroyed beyond the capabilities of FEC mechanisms. Although this problem can be mitigated by using bit-interleaving schemes, the necessarily high interleaving depth adds further delay which may impair performance of delay sensitive services.

In our HRP scheme, we protect transmission via fast retransmission rather than via FEC. As a consequence of the dynamic slot assignment, a corrupted packet can be recovered immediately through a retransmission attempt. Our approach allows a more efficient use of the available bandwidth, as it is adaptive to the channel noise conditions. In fact retransmissions use bandwidth only when needed, i.e. after a packet is corrupted. Conversely, FEC requires an a priori bandwidth consumption due to the fixed coding overhead. Moreover, fast retransmission is of paramount importance when dealing with delay sensitive traffic. We'll show in the performance section that, for voice sources with extremely tight delay requirements (32 ms), a packet loss ratio lower than 1% can be achieved on a highly disturbed channel, where more than 15% of the transmitted packets are corrupted.
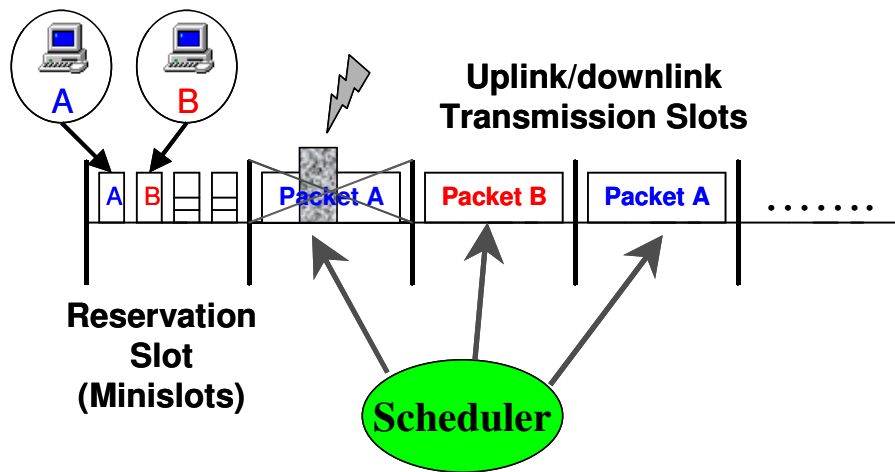
Fig. 3 - Dynamic slot assignment operation

The fast retransmission capability is schematically illustrated in Fig. 3. In the first reservation slot, two reservation packets are received for two different uplink connections generated by two AUs, referred to as A and B. In the next two consecutive slots, the scheduler grants a transmission opportunity for each reserved AU. In the figure, we assume that packet A, while transmitted on the channel during the second slot, is corrupted by a burst of impulsive noise. The scheduler may grant an immediate retransmission attempt at the first opportunity (in the example, the fourth slot in the figure – but the scheduler might decide to assign the third slot to A and the fourth to B, depending on the signaled transmission urgency for each session). This allows to recovery the packet with very little additional delay. In what follows, we detail the scheduler operation adopted in the simulation program. In the following description, we focus on the more complex case of uplink transmission (the downlink case is similar, but with the fundamental simplifying difference that no reservation is necessary).

A. Statistical Multiplexing

In order to be able to access the channel, a connection needs to pass a set-up phase,

whose success may depend upon an admission control decision. If accepted, the connection is assigned an identifier. In what follows, for simplicity of presentation, we'll use the words connection and AU as synonymous, i.e. we assume that an AU handles exactly one connection, and that an identifier is univocally assigned to each AU.

An uplink connection is a stream of fixed-size packets, each accommodated in a channel slot, that need to be delivered from the AU to the BS. In general, in the case of real-time services, data to be transmitted is not stored beforehand in the AU transmission buffer, but is dynamically generated during the connection lifetime (e.g. think to a voice source). Moreover, the emission rate can vary during time. In the case of voice sources, a typical emission pattern is ON-OFF, i.e. the traffic source alternates between talkspurts, in which traffic is emitted at peak rate, and silence suppression periods, in which no traffic is generated.

Our HRP scheme provides the scheduler with the ability to poll an AU only when the relevant traffic source is effectively emitting packets. To this purpose, an AU can be found in one of the following states: "idle", "active" and "contending". In the case of "idle" state, the AU is not generating traffic (e.g. a voice call during a silence suppression period). When "active", the AU is granted uplink slots according to the scheduling rule. In the transition from active to idle state, the AU, via in-band signaling (a bit in the packet header), notifies the scheduler to stop sending transmission grants. This allows the scheduler to dedicate channel resources only to active stations, and thus to increase the number of calls that can be simultaneously supported on the PLC channel (statistical multiplexing). However, as soon as the AU re-starts to emit packets, e.g. at the beginning of a voice talkspurt, the scheduler needs to be promptly informed

of this fact, in order to re-start issuing transmission grants. This information is conveyed via the reservation channel. Specifically, when reactivating, the AU enters a "contending" state, and starts sending reservation minipackets in the reservation slots periodically made available by the scheduler. When a reservation packet is successfully received, the station moves into active state, and waits for dedicated uplink transmission slots to be issued.

B. Management of reservation slots

In the scheduling design, one issue to consider is how frequently a reservation slot needs to be scheduled. Too many reservation slots consume channel capacity which might be used for packet transmission. Conversely, a low frequency of reservation slots leads to higher delay in the reservation phase. In our implementation, we have adopted the following heuristic trade-off. One reservation slot is issued at least every T ms. In the simulation program, we have set T=16 ms, being this value the inter-arrival time among two voice packets. This allows voice sources (and, in general, delay sensitive sources), to rapidly find opportunities to reserve the channel, and allows the scheduler to exploit statistical multiplexing at the voice talkspurt time scale (hundreds of ms).

If collision is detected in at least a mini-slot composing the reservation slot, the scheduler enters a contention resolution phase. In this phase, the scheduler issues consecutive reservation slots until a reservation slot without collision is encountered. As an exception to this rule, the scheduler may intertwine reservation slots with transmission slots issued to AUs whose head of line packet stored in the transmission buffer has reached the maximum lifetime (i.e. the considered slot is the last transmission opportunity for the considered packet).

C. Scheduler parameters

The scheduler is the core of the MAC protocol. It manages channel resources in order to optimize the system performance and to integrate traffic sources with different transmission rates, priorities, delays and packet-loss requirements. The scheduler operates on the basis of service parameters acquired for each connection during the setup and during the reservation. In particular, during the setup, each AU specifies the traffic source descriptor, namely the parameters:

- *MD*: the maximum tolerable packet delay;

- *IT*: the inter-arrival time, measured in slots, of packets at the AU, i.e. the source emission rate when active[2].

This traffic source descriptor is stored at the BS. This descriptor is complemented by additional information included in the reservation packet, and specifically the time *W*, in slots, spent by the first packet of the burst in the transmission buffer before the reservation. Based on the above parameters *MD*, *IT*, and *W*, the scheduler can construct a list of AUs, keeping track of the following dynamic information:

- Head-of-Line (HOL) packet remaining lifetime: this counter is initialized as *MD-W* (i.e. the maximum packet lifetime *MD* minus the time already spent by the same packet in the AU buffer), and is dynamically updated on a per-slot basis (decreased at every slot, and increased of a quantity *IT* after the successful transmission of a packet);

- Number of retransmission attempts for the considered AU. This value is initialized to 0, meaning that the current packet has never been transmitted; it is increased after every unsuccessful retransmission, and it is reset after a

successful transmission.

D. Basic management of data slots

The above parameters allow to implement an earliest deadline first scheduler. The scheduler maintains a dynamic list of reserved AUs. For convenience of implementation, this list is organized into a polling register of size equal to the maximum *MD* value declared by the offered traffic sources. Each position of the register is numbered starting from 0, and is either empty or filled with the identifier of an AU waiting to be polled. An AU placed at position *k* in the polling register implies that the relevant packet must be transmitted on the channel in at most *k* slots, otherwise it expires (i.e. the packet remotely buffered at the AU reaches the maximum tolerable delay).

In each position of the polling register, at most one AU can be stored. When an AU successfully reserves, the polling register attempts to store the AU at the position *MD-W*, representing the remaining lifetime of the current HOL packet. If this position is not empty, the new reservation is stored in the highest available position smaller than *MD-W*.

Slot by slot, the scheduler issues a grant for the AU in the lowest register position (actually, the implemented rule is a bit more complex, as explained below). This implies that, at each slot, the transmission grant is assigned to the AU whose HOL packet has the lowest remaining lifetime. At the end of the slot, the scheduler analyzes the outcome of the transmission grant. On of the following four cases may occur.

1) If the transmission is successful and the transmitted packet is not the last in a talkspurt (this information is indicated in a flag in the packet header), the BS

---

2. Although not considered in this paper, note that the parameter IT can be dynamically varied when the AU is in the

scheduler a successive transmission for the next packet of the same AU. The arrival time of the next packet as well as its deadline is trivially computed by knowing the inter-arrival time *IT*.

2) If the transmission is successful and the packet is the last in a talkspurt, the AU is removed from the polling register. It is now duty of the AU to reserve again once a new talkspurt begins.

3) If no transmission is detected on the channel (e.g. all the packets stored in the AU has expired), the AU is removed from the polling register.

4) If the transmission is corrupted, the BS leaves the reservation in the same position of the polling register, and updates the relevant retransmission counter, whose usage is explained below.

Finally, the scheduler shifts the register, and schedules transmission for the new slot. Note that in the special case of corrupted transmission for an AU reserved in position 0 of the register, the relevant packet is lost (after the register shift, this position would become -1), and the next packet for the AU is rescheduled.

Note that the described operation implies that an eventually large number of consecutive transmission grants can be issued for the same packet, since a corrupted packet remains in the lowest polling register position. This is on one side unfair, and on the other side it can be a critical problem when packet corruption is caused by an abrupt AU malfunctioning rather than channel disturbance. To overcome this problem, we have added to each reserved AU a retransmission counter (see subsection *C* and case 4 described above). The scheduling rule is modified as follows: the transmission grant is assigned to the station in the lowest register position, among the ones with lowest

---

active station by means of signaling carried in the packet header. This feature allows to easily support variable rate traffic with emission patterns different than ON-OFF.

retransmission counter value (with the exception of an AU reserved in the polling register position 0, for which a last transmission opportunity is always granted). This implies a somehow cyclic operation: a first transmission opportunity attempt is given to all AUs reserved, then a second one is given to AUs for which one transmission failed, and so on.
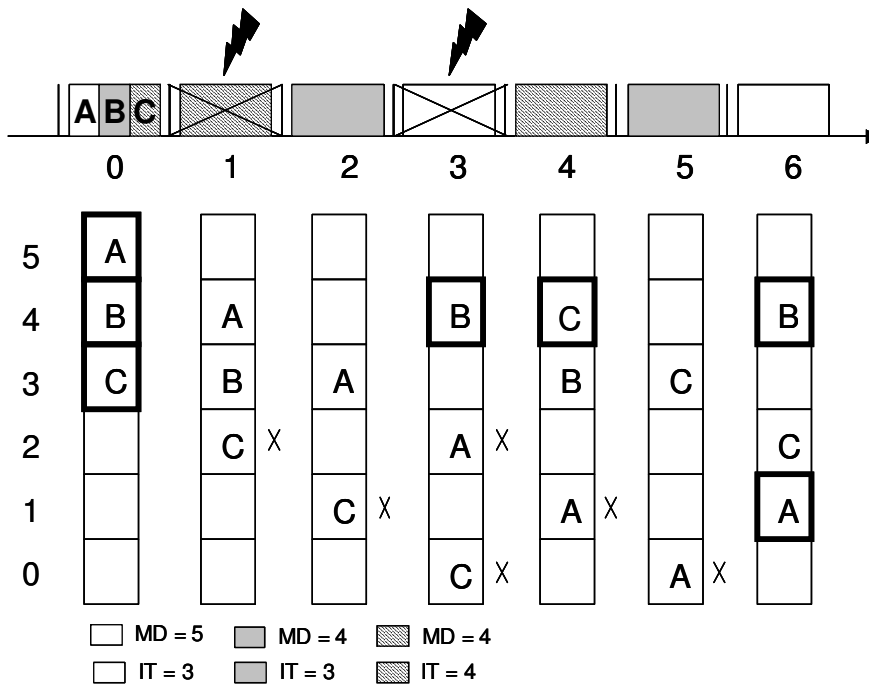

Fig. 4 - Example of polling register operation

For convenience of the reader, Fig. 4 illustrates a detailed example of scheduler operation. The figure shows, slot by slot, the transmission grant and the corresponding polling register status. Horizontal numeration refers to time slot scale, while vertical numeration refers to polling register positions. Assume that three active stations need to share the channel. The figure reports, for each station, the traffic source descriptors *MD* and *IT*. After the reservation slot, each AU is inserted in the polling register in the highest available position of order not superior to *MD-W*. Note that, since the

reservation is immediately successful for all the stations, $W_A = W_B = W_C = 0$. Reservations are stored in the register on the basis of the transmission order. This implies that reservation for station C is stored in the third position, since the location number 4 is already filled by station B. In the figure, the bold lines of the register indicate the storage of a new packet reservation. According to the previous scheduler description, the first station to be served is station C as its reservation is located in the smallest register position. In slot 1, transmission by station C fails. Hence, the station C reservation is left in the same polling register position. At the end of the slot, the polling register is shifted and the retransmission counter for station C is incremented (the mark *X* in the figure indicates a failed transmission). In slot 2, a successful transmission grant is given to station B because it has a lower retransmission counter value than C. To reschedule the next packet transmission for B, it must be considered that B has not yet received the next packet. This packet will arrive in the buffer at time 3 (since the previous packet was arrived at time 0 with W=0, and IT_B =3), and will be scheduled in position 4 since $MD_B = 4$. In slot 3, the packet transmitted by station A is corrupted. At this time, each station has received a first transmission grant. Moreover, at slot 3 station B and station A have generated a new packet ($IT_B = IT_A$ =3), but while the reservation for B appears in the polling register, the reservation for A is not inserted until the previous reservation is served or expired. In slot 4, a transmission grant is issued to C (instead of B, which has a lower reservation counter value) only because it is stored in polling register position 0 and cannot wait anymore in the transmission buffer. After a successful transmission, a new reservation for C is rescheduled in position 4 of the polling register. In slot 5, a transmission grant is issued to B, while, similar to the case of slot 4, in slot 6 a transmission grant is issued to A because it is placed in position 0.

Finally, at the end of slot 6, a new reservation for station A can be stored in position 2, but since this position is already used by C, it is stored in position 1.

E. Management of different traffic classes

The central scheduler allows different traffic sources to share the channel, according to different rate and delays requirements. In order to further differentiate the service provided to each user, it is possible to exploit the intrinsic flexibility of the dynamic slot assignment, by introducing other forms of priority in the scheduling algorithm.

We assume that each traffic class is characterized by two different constraints: the maximum tolerable packet loss and the maximum tolerable delay. For example, as we detail in the following section, we assume that for a voice source the values of such parameters are respectively 1% and 32 ms. While delay differentiation is natively provided by the reservation scheme, packet loss differentiation requires the introduction of other mechanisms.

Packet loss phenomenon is due to two different contributions: temporary network congestion and noise corruption. The first contribution depends on the statistical multiplexing and cannot be avoided even in absence of noise on the channel. If a great number of stations are contemporaneously active, some transmission requests cannot be served before they expire. The second contribution is due to the retransmissions scheduling. Because of bursts of impulsive noise, several successive packets are corrupted and channel resources can be not sufficient to recover all of them. To cope with these phenomena, two different priority strategies can be applied, separately or in conjunction.

Multiplexing Priority: This form of priority operates on the reservation requests. High priority *HP* stations have the right to substitute low priority *LP* stations in the polling

register, if they find the location corresponding to the HOL remaining life-time filled. Deleted *LP* reservations are rescheduled. In other words, a register position is considered available for a new reservation not only if it is empty, but also if it contains an AU identifier with a lower priority. This operation implies that, in the case of congestion (i.e. available polling positions not sufficient to accommodate all the reservations), *LP* requests are discarded for first. On the other hand, high priority reservations are, on average, stored in higher register positions. Consequently, they experience a lower loss probability but an higher average service delay with respect to lower priority stations.

Transmission Priority: This form of priority operates on the transmission grants for already reserved stations. High priority and low priority stations divide the same polling register, but *LP* stations are polled only after that the last *HP* reservation is served . In other words, the previous scheduling algorithm is separately applied to each station class, and *LP* class is considered only at the end of *HP* class service. This operation implies that reserved *HP* stations do not see the channel shared with *LP* stations, and *LP* stations can only use excess channel resources. Exceeding resources depend on the number of *HP* stations, but also on the channel quality. In fact, as the channel conditions degrade, resource consumption due to *HP* retransmissions increases and bandwidth is subtracted to *LP* stations.

IV.   PERFORMANCE EVALUATION

The performance of the proposed HRP MAC protocol have been evaluated using a C++ simulator that implements all the procedures specified in the previous sections. Results are obtained for the more critical case of uplink transmission and delay-sensitive ON-OFF voice sources. Talkspurts (ON periods) and silence periods (OFF)

are exponentially distributed random variables. Following the traditional model [7], we have set the average talkspurt duration equal to 1 s, and the average silence duration equal to 1.35 s. During a talkspurt, the emission rate is constant and set to 32 Kbps. After packetization, a source is assumed to emit a 512 bit packet payload every 16 ms. A very stringent delay requirement of 32 ms maximum packet lifetime (i.e. equal to two inter-arrival packet periods) is enforced: a packet not successfully transmitted within this deadline is considered lost. The quality of service requirement considered in this paper for a voice source is a packet loss ratio not greater than 1%

Regarding the channel, we have considered a slot size equal to 576 bits. This slot size is assumed to be able to accommodate the 512 bits packet payload, plus command, packet header and guard times (see Fig. 2-c). We have considered three different channel rate scenarios: 720 kbps, 1.44 Mbps, and 2.88 Mbps. Given a slot size of 576 bits, and an inter-arrival packet time IT=16 ms, a channel with rate R can accommodate up to $F = R \cdot IT / 576$ simultaneously active voice call (i.e. F=20,40,80 for the three channel rate scenarios considered).

		Hardly disturbed	Moderately Disturbed	Lightly disturbed
Mean IAT	**Milliseconds**	**15**	**476**	**1903**
	Slots @ 720 kbps	18.75	595	2378.75
	Slots @ 1.44 Mbps	37.5	1190	4757.5
	Slots @ 2.88 Mbps	75	2380	9515
Mean error burst duration	**Milliseconds**	**2.08**	**0.87**	**1.82**
	Slots @ 720 kbps	2.6	1.09	2.28
	Slots @ 1.44 Mbps	5.2	2.18	4.55
	Slots @ 2.88 Mbps	10.4	4.35	9.10

% corrupted	720 kbps	16.86 %	0.35 %	0.14 %
slots	1.44 Mbps	14.52 %	0.27 %	0.12 %
	2.88 Mbps	13.35 %	0.22 %	0.11 %

Tab. 1 - Considered noise scenarios

The impulsive disturbance of the power-line channel is modeled as a GOOD/BAD Process, with exponentially distributed GOOD and BAD periods. The BAD state represents the duration of an impulse, during which we consider the channel as disturbed and assume that no transmission data is possible. The GOOD state represents the absence of noise impulses; during this state the channel is available for information transmission and no transmission error occurs.

In addition to the ideal reference case of no channel errors, three disturbance models have been simulated: lightly disturbed, moderately disturbed and hardly disturbed. These models, proposed in [5], are specified in Tab. 1. This table reports, for each disturbance model, the values in ms of the mean interarrival time (IAT) between two error bursts (i.e. the mean duration of the GOOD period), and the mean duration of the impulsive noise (BAD period). The hardly disturbed case is particularly critical, as error bursts lasting, in average, 2.08 ms occur with a mean interarrival time of as low as 15 ms.

To better visualize how critical each disturbance model is, Tab. 1 reports the mean IAT and error burst duration expressed in channel slots, for the three channel rates considered in the paper. In addition, the table reports the percentage of corrupted channel slots. This value depends on the channel rate, and is greater than that computed starting from the values IAT and error burst duration expressed in ms. In fact, an error burst generally starts and ends at any time inside a channel slot. Therefore, it is

immediate to prove that an error burst lasting, in average, *x* slots, will result in corrupting an average number *x*+1 of slots. As shown in Tab. 1, the percentage of corrupted slots is close to 15% for hardly disturbed channels, while it drops to about 0,3% and 0,1% for moderately and lightly disturbed channels.
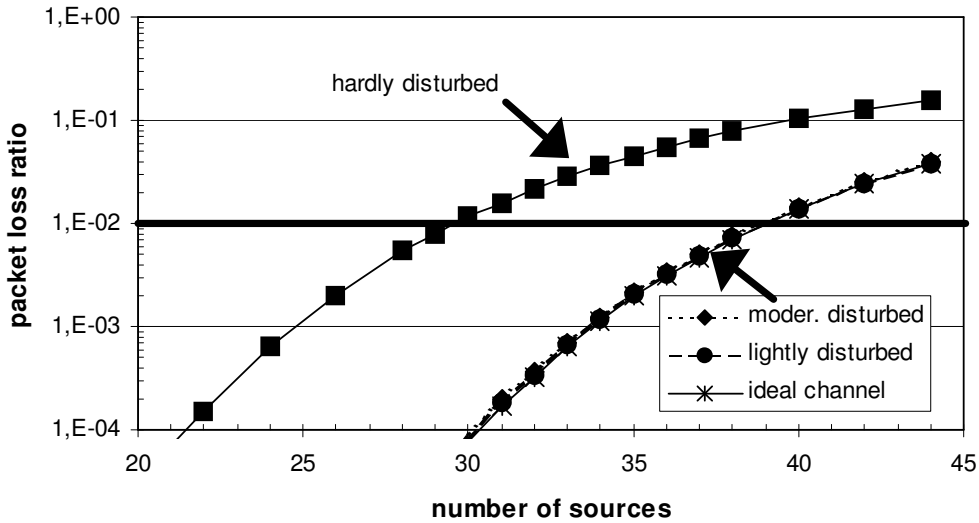

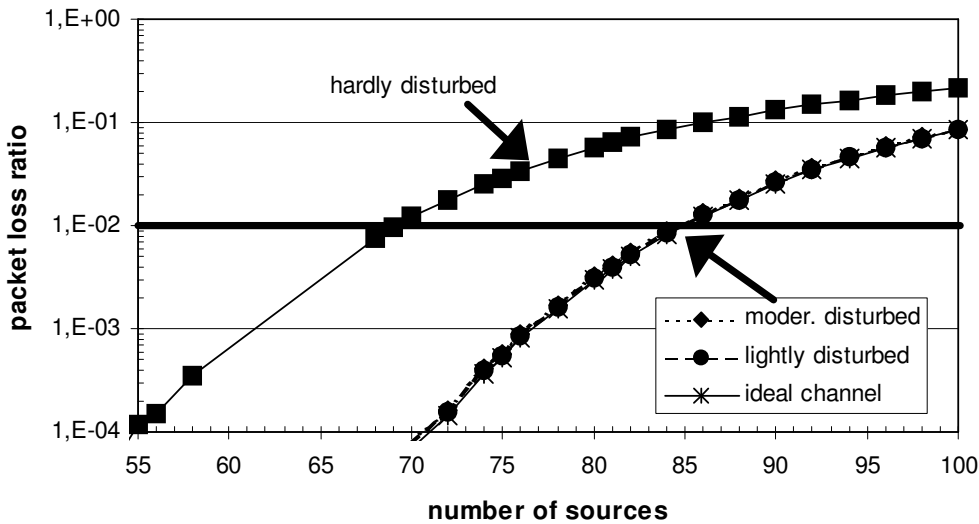*Fig. 5 - Packet loss probability versus number of calls for R= 720 Kbps (F=20) and m=4*

Fig. 6 - Packet loss probability versus number of calls for R=1.44 Mbps (F=40) and m=5.

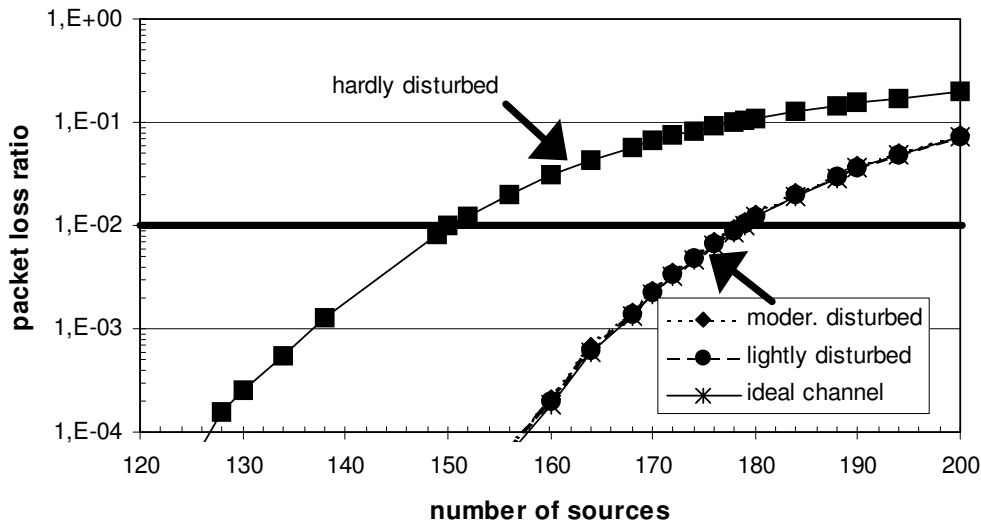*Fig. 7 - Packet loss probability versus number of calls for R=2.88 Mbps (F=80) and m=8*

A. Simulation results for homogeneous sources

Simulation results are shown in Fig. 5, Fig. 6, and Fig. 7, for the case of channel rate equal to 720 Kbps, 1.44 Mbps, and 2.88 Mbps, respectively. The packet loss probability versus the number of offered voice calls is plotted for the three considered disturbance noise scenarios. For each channel rate considered, the number m of mini-slots composing a reservation slot has been set to a sufficiently high value in order to achieve optimal performance (i.e. the performance obtained are extremely close to that achieved by using an ideal – not implementable - reservation mechanism where collision is resolved in a deterministic manner). It can be noticed that the value m necessary to achieve optimal reservation efficiency increases with the channel rate. The reason is evident by considering that the reservation load is given by the rate at which traffic sources move from idle to contending state, which increases linearly with the number of

accommodated voice calls. The adoption of sub-optimal values for *m* leads to performance impairments, as a greater amount of time is spent in resolving contention for access to the reservation mini-slots. This problem can become dramatic when a too small value for *m* is considered, as the reservation channel behavior may start suffering from instability problems (the channel is trashed by repeating collisions in the reservation minislots), that cause a sharp increase in the loss curves and in the confidence intervals measured in the simulation runs.

From Fig. 5, Fig. 6, and Fig. 7, a number of important considerations can be drawn. First, the use of reservation on a talkspurt basis allows to significantly increase the number of accommodated calls in ideal channel conditions. Assuming a Quality of Service requirement of 10^{-2} packet error ratio, and setting the link capacity to, e.g., 1.44 Mbps (Fig. 6), the number of admitted voice calls is 84. This number is given to the statistical multiplexing capability of the proposed HRP scheme, and it has to be compared with the maximum value of 45 (1440 Kbps channel rate divided by 32 Kbps source peak rate – neglecting overhead) achievable by fixed TDMA systems. Tab. 2 summarizes the capacity of the HRP protocol for the various channel disturbance models. The capacity is defined as the maximum number of voice calls that can be admitted, in order to achieve a target packet loss lower or equal than 1%. For reference purposes, the capacity of a traditional TDMA system is also reported in the table, but remark that traditional TDMA performance is impaired when hardly disturbed channels are considered, as about 15% of the transmissions are corrupted in an unrecoverable manner.

	R=720 Kbps (F=20)	R=1.44 Mbps (F=40)	R=2.88 Kbps (F=80)
Hardly disturbed	29	69	149

Moderately disturbed	38	84	178
Lightly disturbed	38	84	178
No disturbed	38	84	178
Traditional TDMA	**22**	**45**	**90**

Tab. 2 - Number of admitted voice calls with a packet loss not exceeding 1%

Second, The figures show that the system efficiency is virtually not affected by lightly to moderately disturbed powerline channels: the performance results in light and moderately disturbed channels lay, in practice, on the curve obtained with no channel error.

Finally, and most interesting, the proposed scheme allows to reach the target 10^{-2} packet error ratio even on hardly disturbed channels, despite the fact that the probability of a corrupted packet transmission is of the order of 15% (see Tab. 1). As shown by the curves, this is achieved by simply reducing the number of accommodated voice calls (x-axis), until the packet loss ratio drops below the target 10^{-2} value. We conclude that packet loss ratio and system utilization are tightly related: by reducing the number of accommodated calls, we increase the spare capacity available for retransmitting corrupted packets, and therefore we improve the packet loss performance.

B. Simulation results for prioritized sources

We now evaluate the HRP MAC performance when the priority mechanisms described in section 3.E are employed. The aim is to show that a our scheme allows to achieve service differentiation, and allows to protect the performance of a given service class from overload occurring on another service class.

At first, we have evaluated the capacity improvements that can be obtained when the offered traffic is generated by non homogeneous sources, which require different packet

loss or different delays. If no form of priority is used, the most-demanding packet loss requirement or the most stringent delay requirement must be provided to all sources, with a consequent waste of capacity. On the contrary, by using the multiplexing priority, a more efficient channel utilization can be achieved. Fig. 8 shows an example of this capability. The figure plots the packet loss performance obtained by two different traffic classes with the same *MD*: class 1 (high priority) is represented by a constant number of 10 stations requiring a packet loss lower than 1%, while class 2 (low priority) is represented by a varying number of stations which can tolerate higher packet loss. In order to investigate the prevalent effect of the congestion on the packet loss performance, we considered a moderately disturbed noise scenario. The x-axis reports the total number of stations. From the figure, we see that, in absence of priority, in order to guarantee class 1 requirements, we can accept no more than 38 stations. The introduction of service differentiation allows to increase such a number to 40 (5% of capacity improvement).

Fig. 9 refers to a complementary situation: two traffic classes share the channel with different *MD* requirement. Class 1 and class 2 require respectively a lifetime of 40 and 400 slots. The number of less demanding sources is constant and equal to 10. Again, moderately noise conditions are assumed. Fig. 9 shows that class 2 stations experience less packet loss probability, even if no multiplexing priority is considered. This effect is due to the larger delay tolerated for class 2 sources, which allows a better scheduling efficiency. When high multiplexing priority is given to class 2 sources, their packet loss probability significantly decreases, while the corresponding increase in packet loss probability for class 1 stations is negligible. Moreover, we have verified that the average service delay for the two classes do not change. This behavior is particularly

effective in more realistic data-voice traffic integration environment. In fact, voice traffic has generally stringent time constraints, but not severe packet loss requirement, while data traffic requires lower packet loss probability and accepts a greater delay. In such a case, voice sources would be classified as class 1 and data sources as class 2.

In presence of very degraded channel conditions, multiplexing priority is no more able to differentiate packet loss performance, because most of the packet loss occurs for retransmission requests expiration. The increasing service delay due to the multiplexing priority operation can make useless the advantage represented by the low discard probability of the reservation requests. In fact, when the delay is high, the served requests have a small lifetime and transmission failures cannot be recovered before expiration.

In order to overcome this problem, we have introduced a second form of priority, that we called transmission priority (section 3.E). The effectiveness of the complete HRP scheme is illustrated in Fig. 10. Results have been obtained considering in conjunction both forms of described priority.
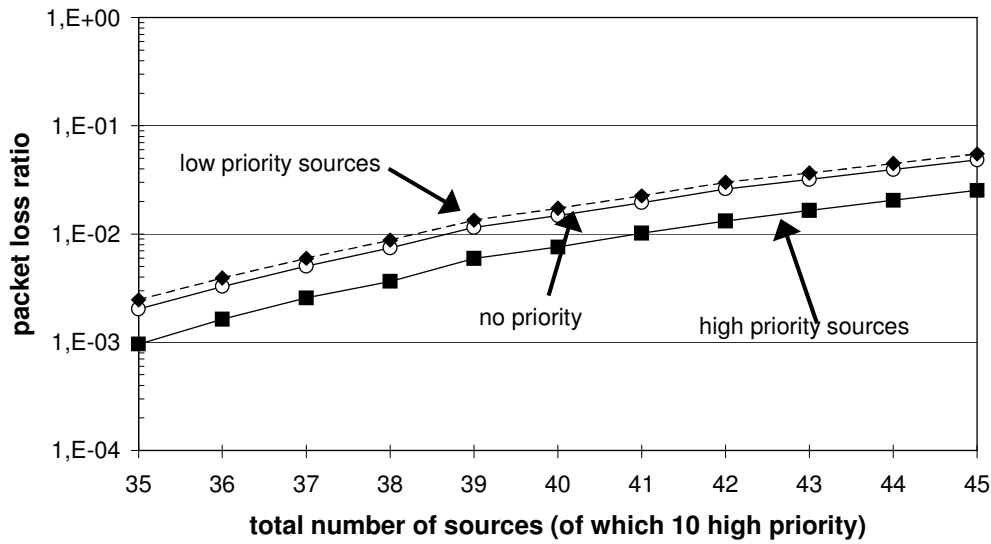
Fig. 8 – performance of high multiplexing priority sources versus variable low multiplexing

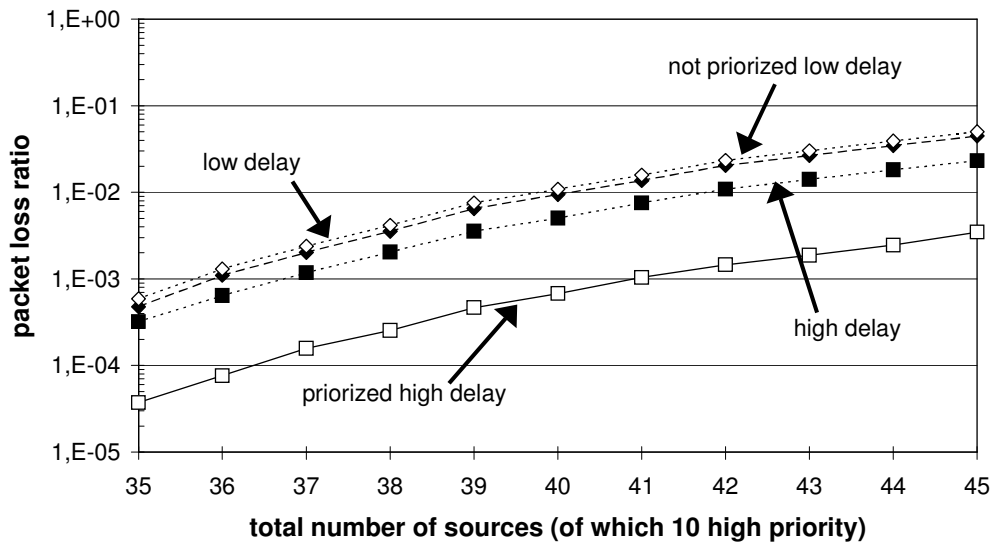*priority load Moderately disturbed channel*

Fig. 9 – performance of high delay sources versus variable low delay load with and without

multiplexing prioritization –Moderately disturbed channel

This figure reports a simulation scenario where 10 *HP* stations compete for channel access with a number of *LP* stations. The highly disturbed noise model is used. The x-axis reports the total number of stations. The y-axis reports the packet loss experienced by *LP* and *HP* stations, the average packet loss and the packet loss obtained in absence of service differentiation. From the figure, some considerations can be drawn. First, *HP* station performance are not affected by the number of low priority stations. In fact, as we explained in section III.E, they are multiplexed and served regardless of the presence of *LP* stations. Second, *LP* stations performance are not too far from the case of homogeneous stations. Third, service differentiation slightly degrades the average packet loss conditions.

Finally, Fig. 11 shows the performance assuming a constant total number of stations (N=29). The x-axis reports the number of *HP* stations, while the y-axis reports again the average packet loss and the packet loss for the two service classes. A comparison with Fig. 5, shows once more that *HP* station share the channel without suffer for the presence of *LP* stations. Conversely, *LP* stations performance degrade as the number of *HP* stations increases. .
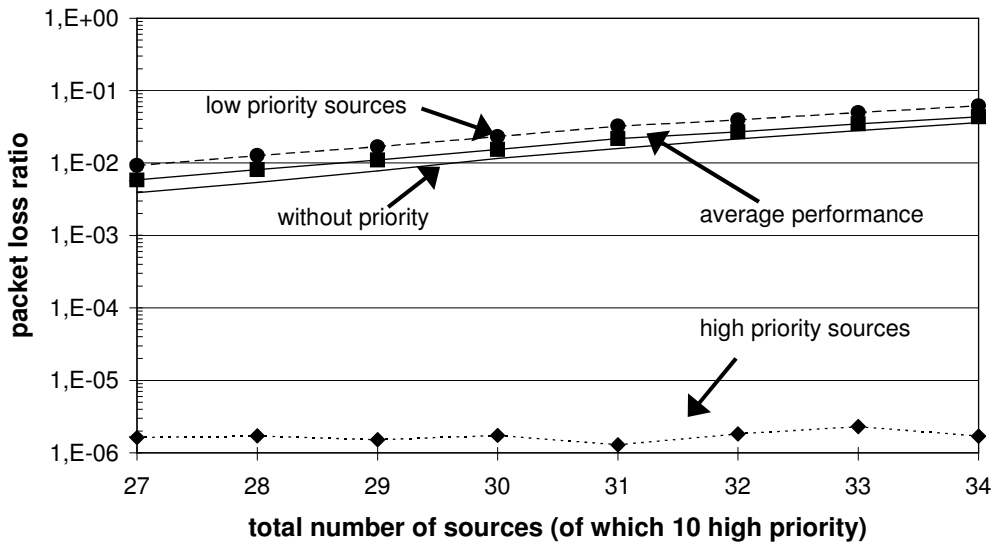
Fig. 10 – performance of high priority sources versus variable low priority load - Hardly
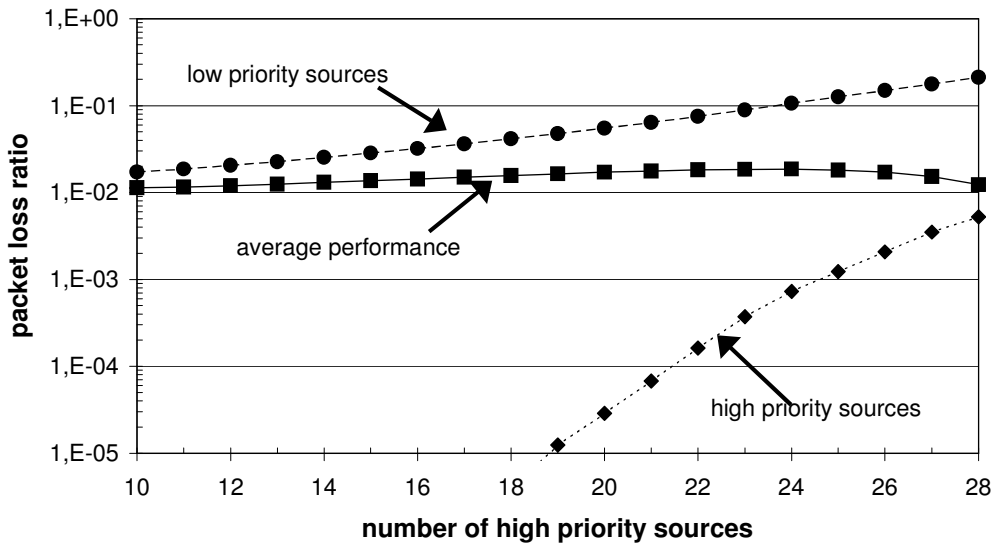
disturbed channel

Fig. 11 – performance of high priority sources versus variable low priority load - Hardly

disturbed channel

V. Conclusions

In this paper we have presented an Hybrid Reservation-Polling MAC protocol for Powerline Communication systems. An Access Unit having packets to transmit first uses a reservation procedure, based on a random access protocol, to insert its nominative in a dynamically managed polling list. To take advantage of statistical multiplexing, the reservation procedure is carried out on a per traffic burst base rather than on a per-session base. Then, a centralized controller manages access to the channel via a command/response operation, i.e. by sending transmission grants to the stations with more urgent needs.

The main feature of the proposed protocol is the capability to achieve high performance on hardly disturbed PLC channels using fast retransmission of corrupted packets.

Performance evaluation has been carried out by means of simulation. Different noise scenarios, channel rates, and station priority models have been considered. Simulation results are obtained for the more critical case of delay-sensitive traffic sources (voice calls). Results show that the system efficiency is virtually not affected by lightly to moderately disturbed PLC channel. Moreover, and most interesting, at the price of a slight reduction of the system utilization, a target 1% packet loss ratio is achieved even in the presence of frequent error burst due to hard impulsive noise, a scenario in which traditional TDMA operation is not possible because of a too high packet error ratio.

Formattati: Elenchi puntati e numerati

References

[1]   W. Liu, H.P. Widmer, J. Aldis, T. Kaltenschnee, ``Nature of Power Line Medium and Design Aspects for Broadband PLC systyem", Proceedings of IZS 2000 - 2000 International Zurich Seminar on Broadband Communications, pp. 185-189.

[2]   H. Hrasnica, A. Haidine, ``Modeling MAC Layer for Powerline Communications Networks", SPIE's (The International Society for Optical Engineering), symposium on Information Technologies; conference "Internet, Performance and Control of Network Systems"; November 5-8, 2000 – Boston MA, USA

[3]   C.N. Krishnan, P.V. Ramakrishna, T.V. Prasad, S. Karthikeyan, "Power-Line As Access Medium-A Survey", International Millennium Conference on Affordable Telecom and IT Solutions for Developing Countries, COMMSPHERE 2000, February 28-March 2, 2000, Indian Institute of Technology Madras, Chennay, India.

[4]   M. Zimmerman, K. Dostert, ``An Analysis of the Broadband Noise Scenario in Powerline Networks", Proceedings of ISPLC 2000 - 4th International Symposium on Powerline Communications and its Applications, 5-7 April 2000, Limerick, Ireland

[5]   H. Hrasnica, A. Haidine, R. Lehnert, ``Reservation MAC Protocols for Powerline Communications", Proceedings of ISPLC 2001 – 5th International Symposium on Powerline Communications and its Applications, April 3-6, 2001, Malmö, Sweden

[6]   G. Bianchi, F. Borgonovo, L. Fratta, L. Musumeci, M. Zorzi, ``C-PRMA: a Centralized Packet Reservation Multiple Access for Local Wireless Communications", IEEE Transactions on Vehicular Technology, May 1997.

[7]   P. T. Brady, ``A model for generating on-off speech patterns in two-way conversations", Bell Syst. Tech. J., vol. 48, no. 7, September 1969, pp. 2445-2472