

Efficient Framing and ARQ for High-Speed PLC systems

Srinivas Katar, Larry Yonge

Intellon Corporation
Ocala, FL, USA
srinivas.katar@intellon.com

Richard Newman, Haniph Latchman

University of Florida.
Gainesville, FL, USA
nemo@cise.ufl.edu

Abstract

In powerline communications (PLC) more so than in many other media, there is a strong interaction between the Medium Access Control (MAC) layer and the Physical (PHY) layer. This paper investigates one of the manifestations of this interaction, that of MAC framing and error recovery with PHY transmission characteristics. MAC Framing for high-speed PLC systems should be capable of concatenating multiple MAC Service Data Units (MSDUs, e.g., Ethernet or IP packets) in order to amortize the high and fixed overheads associated with each MAC Protocol Data Unit (MPDU). By itself, this is easy. However, transmission errors may damage portions of an MPDU, which can further rob a system of efficiency. This paper presents a 2-level MAC framing mechanism that enables efficient retransmission of corrupted data. Both analytical and simulation results are used to compare the performance of the proposed MAC framing scheme with other existing viable approaches. The results show that the proposed scheme is significantly superior to existing approaches. A slightly modified version of this 2-level framing mechanism provided in this paper was used in the new generation 200Mbps class HomePlug Standard, HomePlug AV.

Keywords

PLC, BPL, MAC framing, concatenation, QoS, HomePlug

1. INTRODUCTION

Powerline communication (PLC) systems have come a long way from their modest beginnings in low speed control centric systems to the current generation high-speed solutions to Local Area Networking in Small Office Home Office (SOHO) environments. PLC systems based on HomePlug 1.0 standard [1, 2, 3] are widely available both in North America and Europe. The current generation PLC-based home networking technologies are data centric. While these are good enough for applications like web browsing and file transfer, they cannot support multimedia applications like audio/video streaming and multi-player gaming. These applications require much higher bandwidth than provided by current generation PLC technologies.

Supporting multimedia applications over powerline requires a physical layer (PHY) that is designed to operate close to Shannon channel capacity along with a Medium Access Control (MAC) layer that can guarantee QoS while maintaining high efficiency. One important aspect of MAC design is the MAC framing, namely the process of convert-

ing the payloads received from higher layers (MAC Service Data Units) into a format that can be handed over to the PHY for transmission to the destination (MAC Protocol Data Units). MAC Framing for high-speed PLC systems should be capable of concatenating multiple MSDUs (like Ethernet or IP packets). Concatenation is essential due to the high and fixed overheads (like the PHY layer header, interframe spacings, etc.) required for each MPDU transmission. PLC systems should also deal with recovering PHY layer errors by using proper Automatic Retransmission (ARQ) at the MAC layer. In this paper, we present a 2-level MAC framing mechanism that enables efficient retransmission of corrupted data. Both analytical and simulation results are used to compare the performance of the proposed MAC framing scheme with other viable approaches. The results show that the proposed scheme is significantly superior.

The organization of this paper is as follows. Section 2 provides background on powerline channels and the need for efficient ARQ. Section 3 gives an overview of MAC Framing and provides an intuitive definition of Framing Efficiency, while Section 4 discusses various approaches for MAC Framing. Simulation and Analytical results on the performance of each of these approaches are presented in Section 5. Section 6 provides the conclusions.

2. Powerline Channels and ARQ

Power lines were originally devised for transmission of power at 50-60 Hz and at most 400 Hz. At high frequencies power line is very hostile for signal propagation [4, 5]. Power line networks are usually made of a random interconnection of conductor types and cross sections. Therefore a wide variety of characteristic impedances will be encountered in the network. Further, the network terminal impedance will tend to vary both with communication signal frequency and with time as the consumer premises load pattern varies. This impedance mismatch results in a multipath effect causing deep notches at certain frequencies. In a typical home environment the attenuation on power line is between 20dB to 60dB and is a strong function of load. Further the attenuation characteristics are unique for each transmitter/receiver pair. Hence PLC systems (like HomePlug 1.0) use a mechanism that adapts uniquely to the channel between each transmitter and receiver.

The major sources of noise on power line are from electrical appliances, which utilize the 50/60 Hz electric supply and generate noise components that extend well into the

high frequency spectrum. Typically these generate impulse noises. Electric appliances can be divided into three categories based on the nature of noise they produce in the high frequency bands. They are,

1. Single impulsive noise
This kind of noise is produced when electric switches are turned ON or OFF.
2. Periodic impulsive noise
The most common periodic impulse noise sources are triac-controlled light dimmers. These devices introduce noise as they connect the lamp to the AC line part way through each AC cycle. These impulses occur at twice the AC line frequency as this process is repeated every $\frac{1}{2}$ AC cycle.
3. Continuous Impulsive noise:
This kind of noise is produced by a variety of series-wound AC motors. This type of motor is found in devices such as vacuum cleaners, drills, electric shavers, and many common kitchen appliances. Commutator arcing from these motors produces impulses at repetition rates in the several kilohertz range. Continuous impulsive noise is the most severe of all the noise sources.

High bandwidth digital devices communicating on power line devices need to use forward error correction (FEC) coding along with appropriate modulation techniques to overcome these impairments. FEC codes like Turbo Codes and LDPC Codes perform very close to capacity and hence are a good choice for PLC systems. Apart from these, the MAC layer should use error detection and ARQ to overcome unrecoverable errors at the PHY layer.

Channel adaptation and ARQ go hand-in-hand for ensuring reliable delivery. In general, the channel adaptation mechanism is used to reduce the PHY layer errors to a range where they can be effectively corrected by the MAC layer. For example, the channel adaptation can be chosen such that the FEC block errors are less than 1/1000 and the MAC retransmissions can be used to recover the uncorrectable errors. Such a technique is not viable under impulse noise environments. Impulse noise channels will display temporal effects in channel capacity. In between two impulse noise events, the channel capacity can be very high, while at the impulse noise the capacity is very low. To obtain 1/1000 FEC block error rate, the channel has to be adapted close to the lower end (i.e., capacity near impulse noise events) of its capacity and hence is undesirable. A better approach is to adapt aggressively and use MAC ARQ mechanism to correct errors.

MAC Framing mechanisms that are capable of efficiently handling ARQ are extremely important for new generation PLC systems since impulse noise are very common on Powerline channels.

3. MAC FRAMING AND FRAMING EFFICIENCY

MAC framing process depends not only on the PHY, but also on MAC features like channel access mechanism, net-

work architecture, Layer-2 encryption, etc. For example if one of the MAC features is to provide privacy, the MAC framing process will include fields for passing encryption related information. For a token passing MAC, the MAC framing process will contain fields to indicate busy token, etc. When the underlying PHY is unreliable, MAC framing will also have to add fields for error detection and retransmission. Thus, the MAC framing process depends heavily on system details and the approach used at the MAC layer to address them. Once the MAC framing process generates an MPDU, it is handed over to the physical layer. The physical layer adds the PHY header and transmits it to the destination.

In this paper, we are interested in studying the performance of the MAC Framing process with respect to recovery of errors. To decouple analysis from other system level details, we assume that the only overheads incurred are those required for error recovery at the MAC layer. We define the MAC Framing Efficiency to be the average number of MSDU bits delivered per MPDU divided by the total bits per MPDU. MAC Framing Efficiency depends on PHY errors, MAC framing approach, MSDU size, and MPDU size.

4. MAC FRAMING STRATEGIES

Concatenation is essential for high speed PLC systems due to the large and fixed overheads per MPDU. These occur due to PHY overheads and interframe spacings. In Home-Plug 1.0 [2], each MPDU transmission will incur a minimum overhead of 351.02 microseconds. If a system is ported to 100 Mbps without any concatenation, the maximum efficiency (assuming 1518 byte Ethernet packet as the MSDU) that can be obtained per transmission is ~25%, even in the absence of errors. Thus, concatenation is essential.

Three different packet concatenation mechanisms are investigated here. All three assume that some of the contents of an MPDU may be deliverable even if the MPDU is partially damaged. Without this stipulation, the efficiencies are so low in the presence of errors that these approaches are not considered further. In each case, we assume that the MPDU body is sent in one or more fixed size FEC blocks of a PPDU.

Given that multiple MSDUs may be concatenated to form an MPDU, we consider several basic stages. First, the MPDU may be acknowledged as a whole - either all of the constituent MSDUs are received correctly or none are, thereby requiring retransmission of the entire MPDU contents. This approach has a very low overhead, but is too inefficient when there are errors. The second main stage is to allow each MSDU to occupy its own, individually acknowledged, MAC frame. In this case, only the MAC frames that could not be recovered from a damaged MPDU must be retransmitted. We consider two variations of this approach; Simple Concatenation and Explicit Demarcation

based Concatenation. Finally we consider two-level concatenation and framing. Here, the individual FEC blocks hold their own header information, which allows retransmission of uncorrectable FEC blocks as well as resynchronization of MAC frames.

4.1 Simple Concatenation

This is a very straightforward extension of the case without concatenation. Each MSDU is prepended with a Sequence Number (SN) and a Length Field (Len). A CRC based error detection field is attached at the end of the MSDU. Sequence numbers are needed to ensure in-order delivery and the Length field enables demarcation of each packet at the receiver. Thus, multiple MSDUs can be concatenated at the transmitter and de-concatenated at the receiver. Figure 2a shows the MAC Frame format for simple concatenation. This approach has two main disadvantages. Firstly, each MPDU has to be padded to fit into the PPDU. Secondly, a single FEC Block error will result in loss of information contained in all subsequent FEC Blocks.

4.2 Concatenation with Explicit Demarcation

One of the main disadvantages of simple concatenation is that all information following the location of the first FEC block error is lost. This loss is mainly due to the loss of MAC Frame demarcation. So, it is quite likely that several MAC Frames might have been received without errors at the receiver but the receiver has to discard them due to loss in demarcation information. One simple approach is to add a Header Check Sequence (HCS) for detecting errors on the header fields in the MAC Frame. Thus, whenever the receiver detects a MAC Frame with errors, it can scan through the received bits to check for the next header. Once a proper header is confirmed by means of the HCS, it has the necessary demarcation information for receiving the remaining MAC Frames within the MPDU. MAC Frame header should also include an Identifier (ID) field. ID can be set to zero for the first MAC Frame in an MPDU and incremented for each subsequent MAC Frame. This enables the receiver to use a simple bit map for indicating the reception status of each MAC Frame within a MPDU. Figure 2b shows the MAC Frame format for Concatenation with Explicit demarcation.




Figure 1. MAC Frame format for Simple Concatenation and Concatenation with Explicit Demarcation

4.3 2-Level Framing

PHY errors occur on an FEC block basis. Hence, an optimal framing strategy is one in which only the corrupted FEC blocks need retransmission. 2-Level framing is based on this principle. Here, framing is done in two steps.

In the first step each MSDU is prepended with a Length (Len) field. This entity is referred as a MAC Frame. Multiple MAC Frames that are awaiting delivery are concatenated together to form a MAC Frame Stream. MAC Frame Stream thus generated is treated as a stream of bytes and is then divided into segments of fixed size that match the FEC Block size. Each segment is prepended with a Sequence Number (SN) and a MAC Frame Boundary Offset (MFBO) field. The SN provides the relative location of the Segment within the stream while the MFBO field indicates the location of the first MAC Frame boundary (if any) within the Segment. A CRC based error detection field is attached at the end. The resulting entity is referred to as a PHY Block. Each PHY Block is mapped on to FEC Block at the PHY layer (refer Figure 2). It might be necessary to add a pad to the MAC frame stream when there are not enough bytes to transmit a Segment. Note that a pad is not necessary if there are enough pending MAC Frames. All the results presented in this paper assume that the padding overhead for 2-Level Concatenation is negligible.

5. Efficiency Analysis

We have derived equations for the efficiency of the Simple Concatenation and the 2-Level Concatenation schemes. For Explicit Demarcation based concatenation, simulation results were used to obtain the efficiency. For the results presented in this section, the overheads shown in Table 1 are used. We also assume that all MSDUs are of the same size. Results are presented for the case of 520-byte FEC Blocks with 25 FEC Blocks per MPDU. The FEC block errors in an MPDU are assumed to be independent of each other.

5.1 Simple Concatenation

In the following analysis of Simple Concatenation, we assumed that a single FEC block error would result in loss of all MAC Frames that are contained (either partially or fully) within it. Due to the loss in demarcation information of MAC Frames, all information subsequent to the first FEC Block error is also lost. Thus, if k^{th} FEC block is the first FEC block that was corrupted the number of MAC Frames that will be successfully delivered is given by $\left\lfloor \frac{k L_{\text{fec}}}{L_{\text{mf}}} \right\rfloor$, where L_{fec} is the length of the FEC Block and

L_{mf} is the length of the MAC Frame. L_{mf} is the sum of the size of the MSDU, L_{msdu} , and the overhead added per MSDU, i.e., $L_{\text{OH,SC}}$. If p is the probability of an unrecoverable FEC Block error and N is the number of FEC Blocks, the efficiency for Simple Concatenations η_{SC} is given by

$$\eta_{\text{SC}} = \left[\sum_{k=1}^{N-1} (1-p)^k p \left\lfloor \frac{k L_{\text{fec}}}{L_{\text{mf}}} \right\rfloor \right] + (1-p)^N \left\lfloor \frac{N L_{\text{fec}}}{L_{\text{mf}}} \right\rfloor \frac{L_{\text{msdu}}}{N L_{\text{fec}}}$$

Figure 3a shows the efficiency of Simple Concatenation as a function of FEC Block error rate.

Table 1. MAC Framing parameters

	SN	Len	HCS	ID	CRC	MFBO
Simple Conc.	2	2	-	-	4	-
Explicit Demarc.	2	2	2	1	4	-
2-level	2	2	-	-	4	2

5.2 Explicit Demarcation based Concatenation

For Explicit Demarcation based concatenation, we assume that a single FEC block error will result in loss of all MAC Frames that are contained (either partially or fully) within it. The number of MAC Frames successfully delivered will not only depend on the number of FEC Block errors, but also on how and where they occur. For example, when two side-by-side FEC blocks have errors, then they tend to produce fewer MAC Frame errors than if they are far apart. Hence, it is not possible to derive a simple analytical result for this case. Figure 3b shows the efficiency of Explicit Concatenation as a function of FEC Block error rate.

5.3 2-Level Concatenation

With 2-Level Concatenation, only FEC Blocks that were corrupted need to be retransmitted. If p is the probability of FEC Block error, then the average number of FEC Blocks delivered per MPDU is $(1-p)$ times the number of FEC Blocks, N . The efficiency with 2-level concatenation, η_{2L} , is given by

$$\eta_{2L} = (1-p) \left(\frac{L_{fec} - L_{OH, sb, 2L}}{L_{fec}} \right) \left(\frac{L_{msdu}}{N L_{fec}} \right),$$

where $L_{OH, sb, 2L}$ is the overhead per FEC block. Figure 3c shows the efficiency of Explicit Concatenation as a function of FEC Block error rate.

5.4 Comparison of Performance between various Framing Schemes

Using Simple Concatenation, a single FEC block error results in loss of all MAC Frames that it intersects. This approach is very sensitive to FEC block error rate, p , and number of FEC blocks per MPDU, N . For example at 10% FEC block error rate, the efficiency drops to 30%. The effect of MSDU size on efficiency is small compared to that of the FEC block loss probability. Larger MSDUs will also incur higher overheads due to padding (since a number of whole MSDUs are required to fit in each MPDU), thus resulting in lower performance.

Using Explicit Demarcation based Concatenation, a single FEC block error results in loss of all MAC frames that it

intersects. Simulation results show that the efficiency loss is linear with p and is strongly dependent on the MSDU size. At low p , smaller MSDU sizes show poor performance due to framing overhead. However, as p increases, larger MSDUs sizes suffer higher loss in efficiency since each single FEC block error results in larger number of MAC frame bits lost.

With 2-Level Concatenation, only FEC blocks that were corrupted need to be retransmitted. Results show that the efficiency is linear with respect to p , and independent of N . Figure 3d shows the relative performance of the three framing approaches for 1518 byte MSDU size with varying FEC block error rate. The results show that the performance of 2-level concatenations is superior to that of other approaches.

6. CONCLUSIONS

This paper has highlighted the necessity of concatenation when framing at the MAC level in high-speed PLC communication systems. Five approaches were considered, with three analyzed and simulated in some detail. Approaches varied according to their ability to retransmit at the MAC frame or the FEC block level, and in their ability to resynchronize at the MAC frame level. The results show that 2-Level Concatenation was far superior to the others in its efficiency. It was also seen to be independent of the number of FEC blocks per MPDU, unlike the other two approaches.

7. REFERENCES

- [1] HomePlug Powerline Alliance, www.homeplug.org
- [2] M.K. Lee et al., "HomePlug 1.0 Powerline Communication LANs –Protocol Description and Comparative Performance Results", Special Issue of the International Journal on Communication Systems on Powerline Communications, April 2003
- [3] Intellon Corporation, www.intellon.com.
- [4] Philipps, Holger, "Performance Measurements of Power line Channels at high frequencies." Proceedings of International Symposium on Power-line Communications and its Applications, U.K., 1998.
- [5] Brown, Paul A., "Some Key Factors Influencing Data Transmission Rates in the Power line Environment when Utilizing Carrier Frequencies above 1MHz." Proceedings of International Symposium on Power-line Communications and its Applications, Japan, 1998.

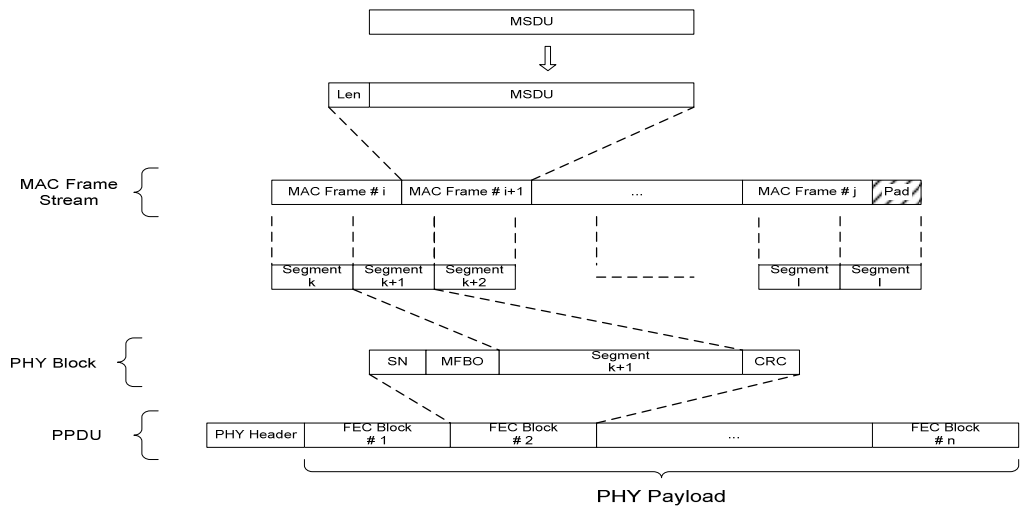


Figure 2. 2-Level Concatenation and Segmentation

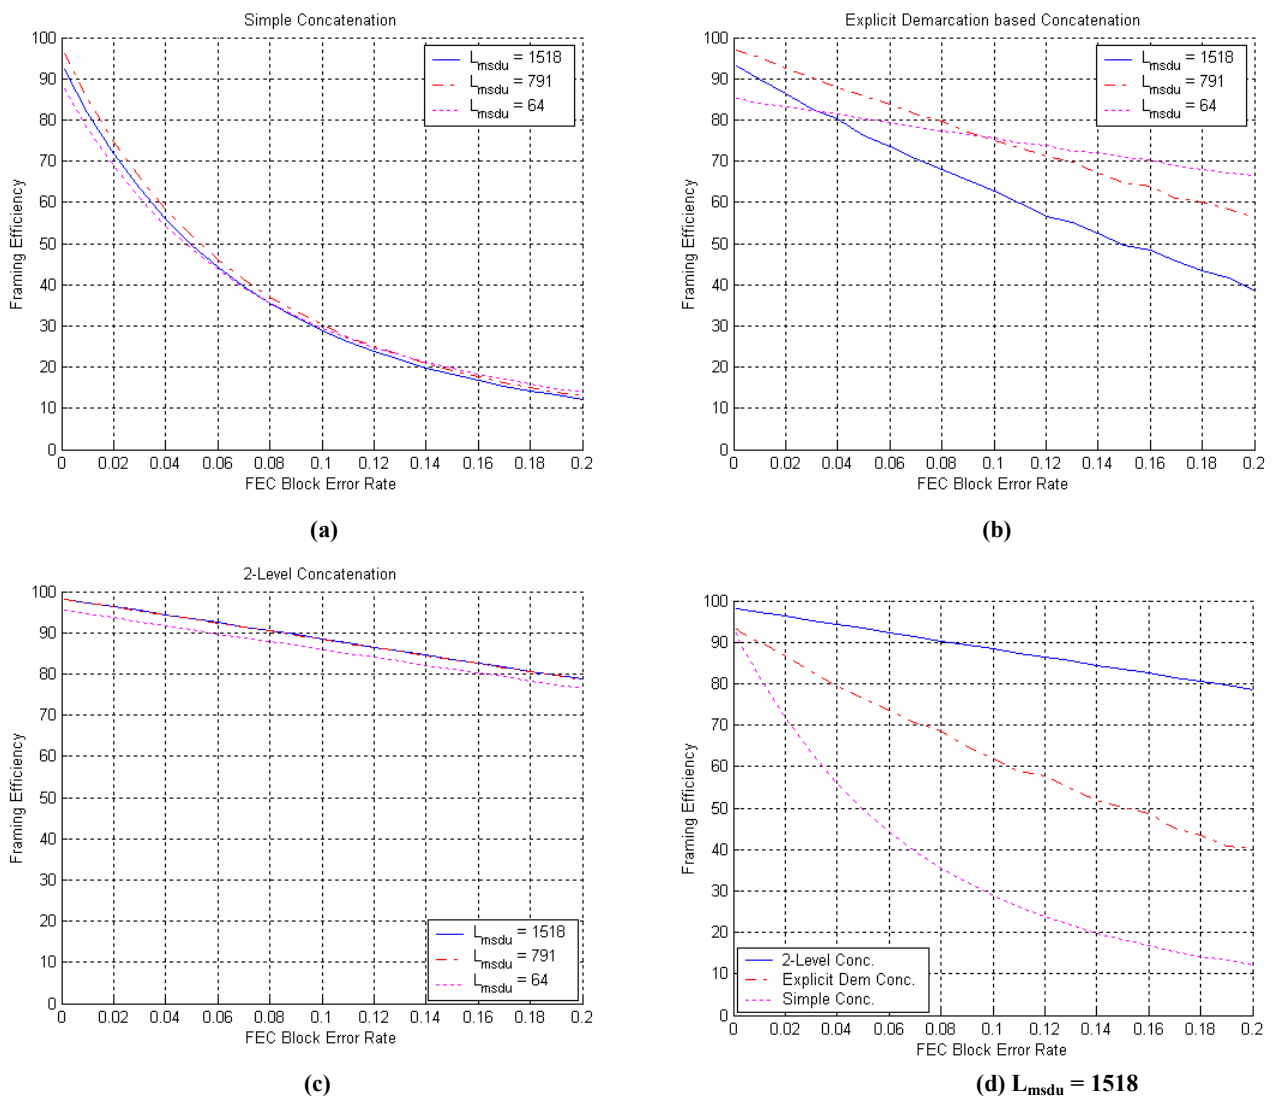


Figure 3. Performance of MAC Framing Strategies