P-37: High Quality Video over Powerline Networks

G. Berbecel, M. J. Omidi, H. Alavi, A. Chini
Valence Semiconductor Inc. 11 Allstate Pkwy. , Markham, Ontario, Canada

Abstract

The article describes a system that can be utilized to transmit compressed video through a Powerline network (based on HomePlug standard) to a digital video display. The issues that are specific to this type of medium, and the ways they are addressed for getting a high quality video are discussed.

1. Introduction

The need for Home Networking has brought to the surface three types of networks: *Wireless Networks*, *Wired Networks* and *No New Wires Networks*, each of them with its own specificity. The No New Wires Networks consist of three distinctive types of networks: *phone line* networks, *cable line* networks *and powerline* networks, among which the last ones are the most promising because of the large deployment in every house, no matter if it is a new or an old one. By having a plug in every room of the house, the powerline network becomes ideal for applications that do require some flexibility in placing the end point devices, but do not need to be truly mobile, and a possible one is to connect a video source and a display. Being a real-time process, the requirements of the compressed video stream come easily into conflict with the variability in time of the powerline channel. Also, because of the large range of the wiring quality, a solution that would work in the great majority of the situations has to be very robust and capable of a great deal of adaptability. The video stream is compressed up to a level that is compatible with the channel characteristics, and the compression rate is adaptively changed if the channel parameters are modified by some other factors that are beyond the user's control. Across the powerline, the signal is coded using Orthogonal Frequency Division Multiplexing (OFDM), in accordance with the HomePlug standard. The Media Access Control (MAC) parameters associated with HomePlug are adjusted to offer the maximum throughput for a given video source-display connection, that will need the entire bandwidth offered by the standard in order to offer a good video quality.

2. HomePlug capabilities and limitations

A complete list of the capabilities and limitations of HomePlug is beyond the scope of this paper and some can be found in the HomePlug 1.0 Specification. Instead, a short list of those features that are relevant for the video application under study will be presented.

- HomePlug 1.0 is a networking standard, that allows more than one source and more than one destination, but in the case of multiple connections happening at the same time over a physical network, the sum of bit rates across all the connections could not exceed the maximum specified by the standard.
- Medium sharing is accomplished by the Carrier Sense Multiple Access/Collision Avoidance (CSMA/CA) mechanism.
- Every connection can have assigned one out of four priority levels, and in the case of two or more devices having the same priority, they will contend for accessing the medium.
- The maximum bit rate on a physical network is about 14 Mbits/second (at the physical layer) and about 8 Mbits/second (at the MAC layer).
- It is a packetized protocol, with the ability of handshake for confirmation of packet reception and retrial in case of packet error.
- It is not a guaranteed connection, in the sense that after a given number of failures in packet re-transmission, the packet transmission is aborted.
- Although at the electrical level a bit duration is always the same, because of various values of some parameters (coding rate, number of carriers) there is a discrete set of 139 bit rates that are possible, and at each specific moment only one can be selected for each link. Some of the possible bit rates can be achieved in more than one way. The range of those discrete bit rates at the PHY level is from 1.13 Mbits/sec to 13.57 Mbits/sec.
- The channel capacity is measured/estimated and made available to the application layer periodically. The powerline channel undergoes long term and short-term variations. The long-term variations are handled by channel estimation process and the short-term variations are handled by buffering.
- When communication is established between two nodes, the bit rate can be any of the previously mentioned discrete ones depending on the quality of the wiring and the noise created by some loads on the same circuit that could be connected or disconnected during the data transmission.

3. Requirements of a good quality video connection

At the source and at the display level the Standard Definition video signal (SD) is 13.5 Mpixel/second which amounts to 216 Mbits/second assuming a 4:2:2 sampling grid and 8 bits per color component. At the compressed level, a bit rate higher than 6 Mbits/second is required to provide good video quality, so a compression of about 35:1 needs to be performed. The MPEG 2 video compression method is selected for this application, as it is the most common method today. The real-time aspect of video is its most demanding aspect. In order to reconstruct and display a full video field, enough compressed data has to be available every 1/60th of a second at the display level. Video compression algorithms (including MPEG 2) take advantage of the high redundancy that exists in the video signal and achieve the compression by eliminating it partially. From the viewer's perspective it would be ideal to have a constant video quality. However, the level of redundancy depends on the video scene, and the compression ratio will be variable. The variability of the compressed data rate is difficult to handle in a real-time system because it will increase the amount of required memory for the system to operate. There are, however, two options for achieving the compression: *constant* bit rate and *variable* bit rate. The constant bit rate is preferable since the amount of required memory (buffering) is smaller, and the overall system is more cost effective. Constant bit rate is achievable in most systems with guaranteed channel bit rate. With a constant bit rate, those scenes that have less redundancy will be compressed more compared to those that have higher redundancy. As a consequence, the final visual result for those scenes with high frequencies content will be worse. Constant bit rate is not an appropriate option for a channel with variable capacity, as it is shown in Figure 1.

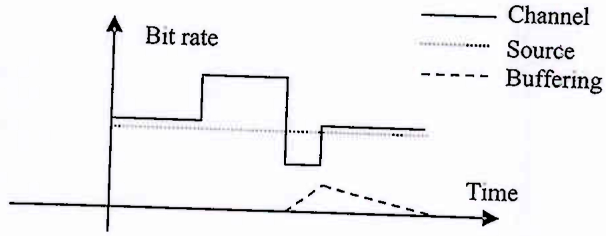

Figure 1. Constant bit rate transmission over a variable capacity channel

When the channel bit rate is higher than the source bit rate the channel is not utilized efficiently. When the channel bit rate is smaller than the source bit rate enough memory is required to allow buffering until the channel recovers. Even in the simple case of a constant bit rate, special care is needed to avoid overflow or underflow at the decoder. Figure 2 shows how the MPEG 2 decoder's buffer is filled in and emptied.

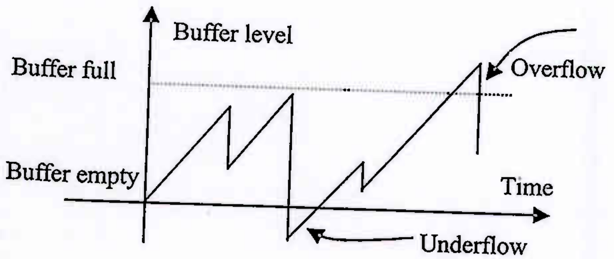

Figure 2. Illustration of overflow and underflow of the memory buffer with a constant bit rate source

Both buffer overflow and underflow have equally bad effects resulting in data corruption. The behavior of the decoder in this situation depends on the implementation through careful design; buffer overflow or underflow shall be avoided. Variable bit rate wouldn't suffer from the visual artifacts the constant bit rate does, but the requirements on memory buffering will be even tougher, and therefore more difficult to be met by a specific design.

4. Top level view of the proposed system

Figure 3 shows the overall view of a system proposed for video transmission over powerline channel. There are three subsystems involved: MPEG 2 encoder, powerline channel and MPEG 2 decoder, each of them being independent and asynchronous to the others. The asynchronous behavior needs to be understood in the context of each one having its own jitter. The overall system needs to be designed such that the worse case scenario in terms of jitter is still accommodated.

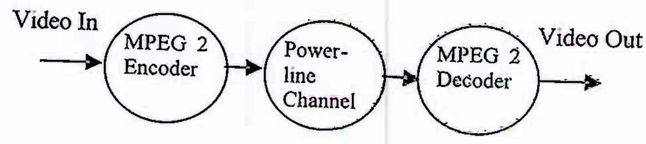

Figure 3. The top-level view of the proposed system

While the MPEG 2 encoder and decoder are well understood and well described by the MPEG 2 standard, the complexity of the HomePlug requires experimental measurements.

4.1 Experimental results of measuring the HomePlug packet jitter

At the application layer there is a packet jitter caused by variable overhead associated with each packet and also caused by variable number of re-transmissions of a packet upon a non-successful transmission. Graphically the packet jitter is described in Figure 4.

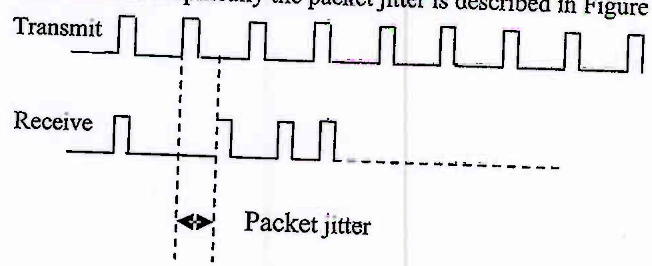

Figure 4. Graphical description of the packet jitter

The size of the buffer required compensating for this jitter is determined by the maximum accumulation of packet jitter. In a loose interpretation, the Transmit and Receive sections can be seen as two independent processes, and the required buffer will serve for synchronizing the two asynchronous processes. The experiment consists of establishing a point-to-point contention free connection between a HomePlug source and a HomePlug destination. A time stamp is attached to every packet at the transmission and reception and then the time from source to destination is calculated for each successful reception. It was repeated three times with 2000 packets, 1500 bytes each. The measured packet jitter was found to be bellow 10 microseconds, with a mean square root of about 1.5 microseconds. Note that, in a multi node connection jitter may well exceed these values.

4.2 Considerations about MPEG 2 encoder configuration

In order to transmit MPEG2 data over powerline, the essential parameter is the encoder's bit rate, which needs to be controlled. The 30 bit field specifying the bit rate is located in the video sequence's header as a combination of two sub-fields: the lower 18 bits resides in the field bit_rate_value and the higher 12 bits in the field bit_rate_extension. A change in the bit rate is possible only at the price of inserting a new sequence header in the video sequence, or just by having small video sequences (containing a single group of pictures) with the penalty of decreasing overall efficiency of the compression. For example, considering a group of pictures (GOP) of length equal to 6, and assuming a bit rate of 8 Mbits/second, the size of a GOP at the compressed level is 100 kBytes. This determines the required buffering when a request for decreasing the bit rate occurs. This delay (about 100 milliseconds corresponding to 6 video fields that run at 60 Hz) before adjusting the bit rate is larger compared to the few milliseconds latency in channel response to a request of measuring its bit rate, and it is the main limiting factor for the design. Some encoders will not allow the bit rate change without performing a full sequence of <Stop Encoding> <Start Encoding> and in that case the time to go *on* and *off* needs to be taken into consideration This can be done by adding the equivalent buffer size to the total size buffer. Without this precaution, the buffer underflow condition will occur because of the gap in proper functioning of the MPEG 2 encoder. If there is a possibility of selecting the interlaced or progressive video for MPEG encoding, the interlaced version should be chosen, because with appropriate video deinterlacing at the display, the overall video quality can be close to the progressive video but with almost half of the bit rate.

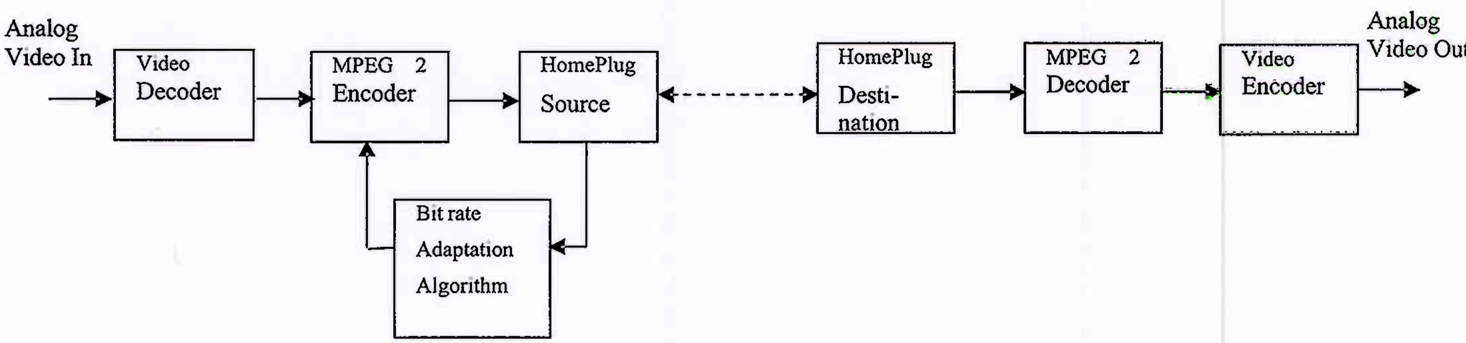

Figure 5. System block diagram

5. System block diagram

Figure 5 illustrates the overall system block diagram, where the analog input video is first decoded and converted into digital, compressed MPEG 2 in accordance with the bit rate specified by the rate adaptation algorithm and packetized by the HomePlug channel encoder. After going across the powerline, the opposite processing stages are applied in reverse order, producing the output analog video signal. In section 3 it was shown that neither constant nor the variable bit rate compression factor is appropriate for the MPEG 2 encoder, so the solution is a piecewise constant bit rate. Once the bit rate is changed, it will remain constant until the rate adaptation algorithm calculates a new value. The parameters that are processed by the rate adaptation algorithm are: the result of channel estimation and amount of data in the transmit buffer. Measurements performed by the HomePlug Technical Working Group [2] have shown the probability of getting a connection using HomePlug for each bit rate in the discrete range of possible bit rates. For North America a connection at 5 Mbits/second can be established with the probability of 80%. The quality of a connection depends on the wiring topology and also depends on various consumer loads connected on the same wire. Any other consumer load that is mounted on the same wire will be a source of noise. We have measured the dynamics of a powerline in two scenarios: one was to monitor the bit rate of the line during one hour without intentionally adding or removing loads on the line as in Figure 6. The other measurement was done for a shorter period of time but with an intense switching of various loads on the same line (and the results are shown in Figure 7)

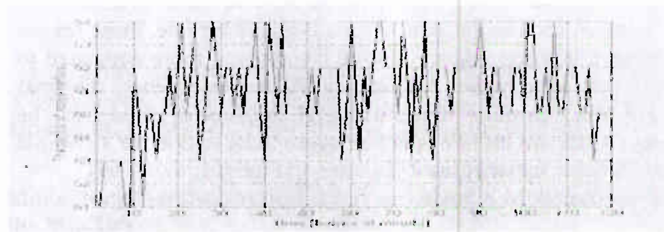

Figure 6. The powerline channel dynamics

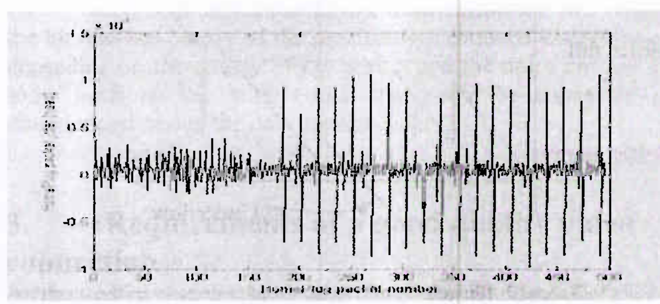

Figure 7. HomePlug packet jitter

The rate adaptation algorithm is illustrated graphically in Figure 8.

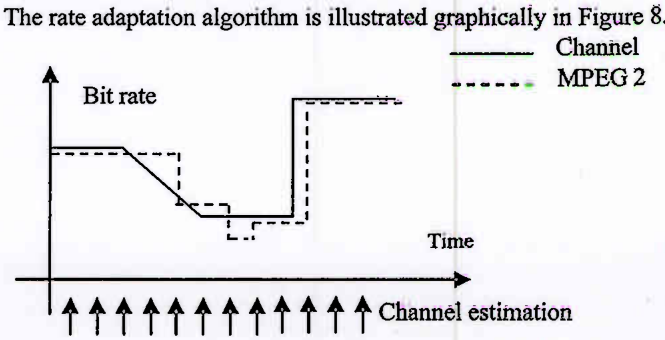

Figure 8. The rate adaptation algorithm

When the channel bit rate drops, the data is stored into a First- In-First-Out memory (FIFO) and that situation is detected by estimating the channel. Upon detection of the channel bit rate drop, the source bit rate is decreased accordingly to match the new bit rate. Since the FIFO is already close to full, initially the bit rate of the source is decreased more than the channel bit rate, to allow the FIFO to work closer to half full. This is shown as an undershoot in the source bit rate whenever there is a bit rate drop. The situation is not symmetrical and there is no overshoot when the bit rate increases, because there is no data to be flushed out from the FIFO. By implementing such an algorithm, reliable video can be transmitted across the powerline with bit rates between 0.7 Mbits/second and 8 Mbits/second, and variation in the level of noise will not influence the continuity of video transmission as long as the connection still exists. The effects of bit rate variation are visible, and more elaborated algorithms can be implemented to filter out some of the variations, which translates into changing the MPEG 2 compression factor less frequent and in bigger steps.

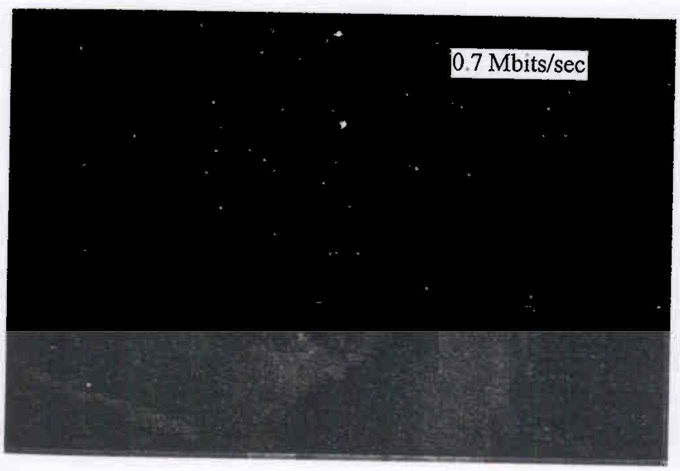

Figure 9. Video quality at four different bit rates

For illustration purposes the above four pictures in Figure 9 show the level of quality (in a monochrome representation) that is created by the system described in this paragraph.

6. Conclusions

This paper demonstrates how HomePlug 1.0 can be employed to carry a single channel of quality video, and opens the possibility to have a true network involving 3 or 4 simultaneous channels of standard definition video or one high definition video channel in the upcoming HomePlug A/V standard.

7. References

[1] HomePlug 1.0 Specifications, June 30th 2001.

[2] Final report on HomePlug Technologies Performance

[3] A comparative Performance Study of Wireless and Power Line Network-Yu-Ju Lin, Haniph A. Latchman, Richard E. Newman, Srinivas Katar –to be published in Communication Magazine

[4] MPEG 2-J. Watkinson

[5] ISO/IEC 13818-2 Generic coding of moving pictures and associated audio information, Part2: Video