2005 IEEE 6th Workshop on Signal Processing Advances in Wireless Communications

Joint Optimization of Transmit Pulse Shaping, Guard Interval length, and receiver side narrow-band interference mitigation in the HomePlugAV OFDM system

Kaywan H. Afkhamie[1], Haniph Latchman[2], Larry Yonge[1], Tim Davidson[3], Richard Newman[2]

[1]Intellon Corporation, [2]University of Florida, [3]McMaster University

ABSTRACT

The primary subject of this paper is the selection of a pulse-shaping waveform for in-home power line communications. As system performance is also determined by other parameters affecting the length and shaping of the OFDM symbol, the problem formulation is expanded to also include the simultaneous selection of guard interval length and Hanning window length, creating a joint optimization problem.

Given the constraints of allowing no transmit notch filters, and adequate receive side mitigation of narrow band jammers, we jointly optimize the selection of guard interval length, transmit pulse-shaping and receive side windowing for throughput performance on the average power line channel. Throughput performance is inferred from SNR data and associated guard interval overhead. We also present results of performance assessment and parameter selection, for the average power line channel, based on a collection of 120 measured power line channel impulse responses.

1. INTRODUCTION

One of the features that made Orthogonal Frequency Division Multiplexing (OFDM) such a popular choice among communication system designers, is that its multi-carrier nature gives it the ability to easily adapt in frequency. This frequency domain adaptability is exploited both in a dynamic sense, where given a frequency selective channel, an OFDM system can concentrate its efforts on those frequency ranges that yield the highest SNRs, and it can also be exploited in a static manner: for example, when a contiguous band of frequencies is not available, OFDM allows the communication system to efficiently use the frequencies where communication is possible.

This is precisely the case in the power line environment. Consideration of FCC regulations (in North America) and the inherent conductive properties of residential electrical wiring quickly determine that the frequency band of choice for in-home power line communications is the unlicensed band from 1.8 to 30MHz. Below 1.8MHz the power line signal could (and would) radiate and interfere with AM radio[1], and above 30MHz the limits for radiated emissions are much lower than in the range up to 30MHz (the resulting reduction in signal power would make communication impractical). Other unlicensed bands become available at higher frequencies, but those signals would undergo prohibitively high attenuation on the power line, leaving us with 1.8 to 30MHz again. Unfortunately, power line communications networks are not the only service in this band. Most notably, in North America there are ten distinct frequency bands (3.8 MHz collectively) reserved for amateur radio operators. To avoid interference, the radiated emissions from power line communications must be 30dB lower than the -50dBm/Hz that FCC regulations would otherwise permit. Figure 1 summarizes the restrictions imposed on the power spectral density of a power line communication system.

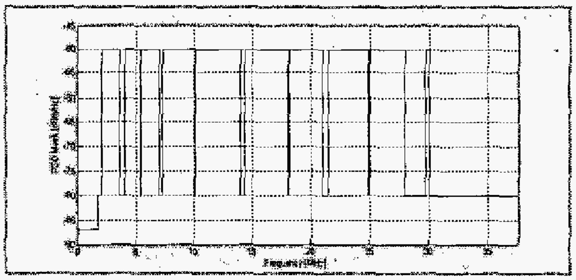

Figure 1. PSD Mask for PLC systems

Note that, Amateur Radio bands are not the only services that might impact the allowable transmit power of PLC systems. In the USA there are also potential considerations for RF bands for remote control toys, and several bands for aeronautical communications. In other countries there may be a whole different set of services altogether.

[1] Electrical wiring in the home acts as a big antenna thus some of the signal energy that is conducted on the power line is radiated; similarly the wiring makes power line communication systems susceptible to RF ingress from narrow band radio transmitters.

0-7803-8867-4/05/$20.00©2005 IEEE

We have already argued that due to all these requirements for spectral responsibility, OFDM is a good choice for power line communications. The remaining question is what combination of OFDM parameters will give the best performance on the average power line channel. In the remainder of this section we introduce some measurements of power line impulse responses and noise that were used to evaluate different OFDM systems. We then discuss some of the thoughts behind the parameter selection that led to the HomePlug AV OFDM PHY [1,2,3], and we formulate the problem of optimum pulse shaping which is the main topic of this paper. In Section 2, we describe the methodology we used to select the jointly optimized parameters in conjunction with the chosen pulse shape, and in Section 3 we discuss some simulation results and the performance of the final system.

1.1. The power line channel

Similarly to the wireless channel where a signal arrives at the receiver over multiple paths, due to reflections from floors and walls, the power line signal also undergoes reflections (due to un-terminated outlets, etc). The channel impulse response therefore has a delay spread that must be mitigated at the receiver [4,5,6]. As is well known, mitigation of delay spread is another feature of OFDM which accomplishes this by simply pre-pending the regular symbol by a cyclic prefix that should be at least as long as the delay spread of the channel. As the multipath nature of the power line channel (as well as attenuation, and noise characteristics) varies considerably in different houses, and different outlet pairs, we created a database of channel models by directly measuring channel impulse responses in 6 different homes and 20 different paths per home (a full mesh of 5 outlets). The channel impulse responses, together with the local noise profiles of each outlet, are used in the simulations section to evaluate the performance of differently configured OFDM systems.

Two samples of channel impulse responses are shown in Figure 2 showing one very good and one very bad channel (together with their respective frequency domain magnitude responses), and Figure 3 shows the distribution of delay spreads of all 120 channels. For purposes of this investigation we define delay spread as the minimum portion of a channel impulse response that contains 95% of the energy.

From Figure 3 we can see that most channels have a delay spread of 4μs or less. This implies that the communication system should be designed with a cyclic prefix of the same order and a symbol length that is preferably much longer than 4μs, so that the overhead due to cyclic prefix is not too costly.

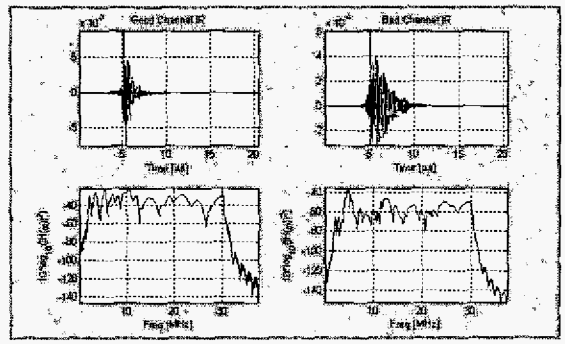

Figure 2. Sample Power Line Channels

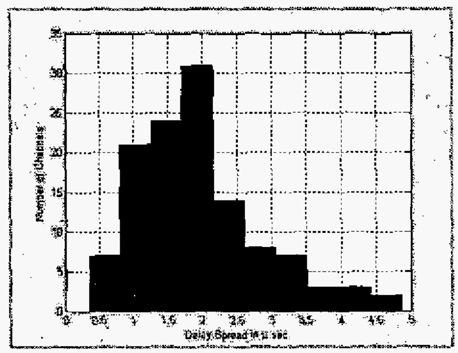

Figure 3. Delay Spread Histogram

At each end node, 40ms of local noise profile were measured with a digital capture card (calibrated to be able to accurately reproduce relative noise power levels in simulations later). Figure 4 depicts two of the noise profiles that include both background noise, and local impulse noise.

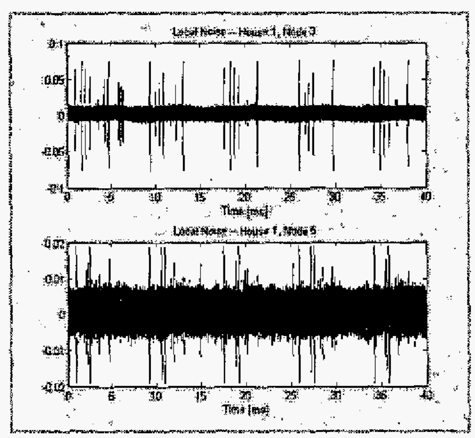

Figure 4. Measured Noise Profiles

1.2. HomePlug AV Parameters

The predecessor technology to HomePlug AV, HomePlug 1.0, uses the frequency band up to about 20.7MHz, and to allow a sufficient band gap to accommodate anti-aliasing filters the sampling frequency is 50MHz. In HomePlug AV the sampling frequency is

75MHz. The corresponding Nyquist frequency of 37.5MHz also gives ample room for band-pass filters, and moreover, the carrier frequencies of HomePlug 1.0 can be easily maintained by just increasing the basic FFT length from 256 (in the case of HomePlug 1.0) to 384. With a 384point FFT, the OFDM symbol length would remain 5.12 µs, and thus the addition of a several µs long cyclic prefix, would make channel use relatively inefficient. In HomePlug AV, the FFT length is therefore increased further to 3072 (8 times 384). The increased FFT length, now leads to a 40.96µs symbol length. A further benefit of the increased FFT length is that it reduces the bandwidth occupied by each single carrier, which helps to accommodate the notches necessary due to amateur radio bands.

Table 1. HomePlug 1.0 and HomePlug AV parameters

	HomePlug 1.0	HomePlugAV
FFT length	256	3072
Nyquist	25MHz	37.5 MHz
Symb Length	5.12µs	40.96µs

The associated cost of the longer FFT size, are (i) that a larger FFT block must be implemented in hardware, and (ii) the reduced carrier spacing increases the systems susceptibility to clock frequency error, and inter-carrier interference.

1.3. Optimization of OFDM symbol pulse-shaping

Essentially there are two options through which the system can meet the PSD requirements illustrated in Figure 1. One option is to transmit regular OFDM symbols, and rely on notch filtering (with either FIR or IIR implementations) to achieve the 30dB notching requirements. The reason that this option is considered undesirable for power line communication systems, is the diverse and uncertain nature of the regulatory environment. The fixed frequency notch filters are not well suited for diverse regulations in differing geographic environments, and it is cumbersome to implement filters with tunable frequency filtering capabilities. The second option is to simply turn off carriers that are located in and around the frequency bands that must be vacated for other uses. For this option, evidently, it is highly desirable to have carrier shapes that roll off very quickly in frequency, so that a minimum number of carriers are turned off for a given notch.

Pulse-shaping waveforms are therefore applied to each OFDM symbol to soften the time-domain waveform, and hence contain the effective frequency domain width of each carrier. Figure 5 illustrates the way HomePlugAV creates a time-domain waveform, including the addition of a cyclic prefix, application of pulse shaping and overlapping of symbols. The transmitter's IFFT function produces 3072 discrete points that represent the length T=40.96µs waveform in Figure 5. A prefix of length t_{prefix} is pre-pended to the symbol, and subsequently both sides of the time-domain OFDM symbol are tapered off, by simply multiplying a length RI pulse-shape to the OFDM symbol. Finally, the tapered end of one OFDM symbol, will be overlapped with the tapered beginning of the next OFDM symbol. This reduces the overall overhead incurred without adversely affecting the frequency domain properties of the signal. GI denotes the effective guard interval length, and T_s denotes the effective symbol length (after prefix and overlap).

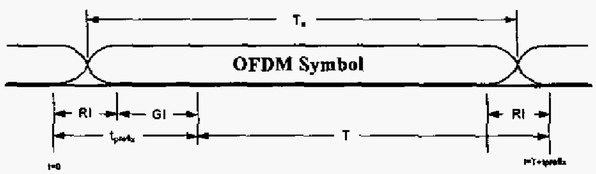

Figure 5. OFDM Pulse Shaping

The goal of selecting the OFDM pulse shape is then to maximize the bandwidth utilization of the system, while minimizing the length of the pulse-shaping waveform (since this reduces the effect of adding a cyclic prefix), and minimizing any adverse effect of pulse shaping on receive side signal processing (e.g., narrow band jammer mitigation).

2. METHOD

As described above pulse shaping waveforms influence the performance of a power line communications system in multiple ways. Namely, they affect the frequency domain carrier shapes, the amount of system overhead, the effectiveness of the cyclic prefix, and the effectiveness of receive side windowing for narrow band jammer rejection. It is thus very difficult to analytically evaluate the benefits of one pulse shape compared to an alternative. In this paper, we therefore resort to evaluating the performance of differently configured OFDM power line communications systems over a set of measured channel and noise models. This section describes how we selected the different configurations of OFDM systems to evaluate.

2.1. Candidate pulse-shaping waveforms

Three different pulse shapes were evaluated: one based on a raised-cosine window, one based on a Gaussian window, and a third pulse shape defined as a piece-wise linear taper. The window lengths for each of the types of pulse shapes were tuned manually, to allow the system to use as many carriers as possible, while having minimum overhead due to the roll-off section of the OFDM symbol. In the case of the Raised Cosine pulse shape, a window length of 432 samples or 5.76µs is used, which allows 899 carriers to be used and still meet the masking requirements of Figure 1. The Gaussian window is also of length 432, but it allows 917 carriers to be used (one more carrier at each boundary point), and finally the piece wise linear taper is defined with a length of 372 samples, and it allows 917 used carriers. The candidate pulse shapes are depicted in Figure 6 (the piece-wise linear taper is easily identified, and the Gaussian waveform is the one that does not decay to zero), and Figure 7 shows the respective PSDs of the different systems. The raised cosine window demonstrates the deepest notches and best out-of-band roll-off but the other two tapers make more efficient use of the bandwidth (18 more carriers), while still meeting the spectral mask requirements.

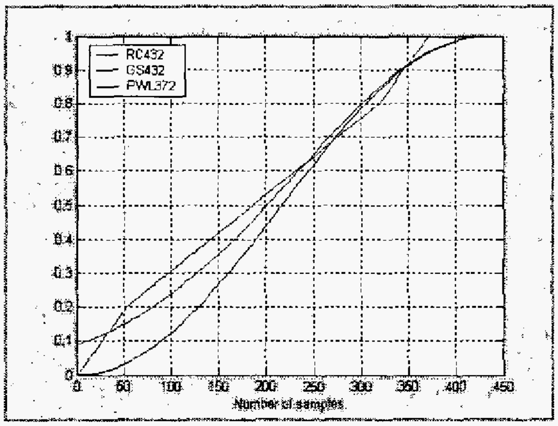

Figure 6. Candidate Pulse Shapes

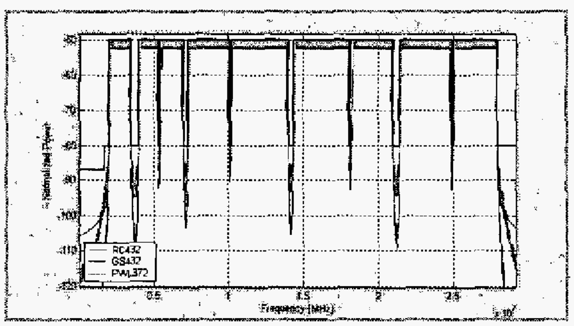

Figure 7. PSDs of candidate OFDM systems

2.2. *GI* range

The next step in defining system configurations is to determine the approximate range of cyclic prefix lengths to use with each of the candidate pulse shapes. We ran clean packets (with each of the three pulse shapes) through a receiver, noting that the receiver will also apply a Hanning window to each OFDM symbol in order to mitigate narrow band interference. It is found quickly that the Hanning window must have a length of at least 90 samples for adequate jammer suppression. With receivers configured to apply a 90 sample Hanning window we measure carrier SNRs on the noise free received packets. For each pulse shape three different length for *GI* were examined, -2μs, 0μs, and +2μs. Table 2 summarizes results of this exercise.

Table 2. GI range for each pulse shape (Receiver Hanning Window Length = 90 samples)

	GI = -2μs	0μs	+2μs
RC432	Min/Avg SNR = 25/31dB	36/38	38/38
GS432	18/26	25/34	38/38
PWL372	15/22	25/30	38/38

In Table 2, each of the table entries shows both the SNR observed on the worst carriers, and the average SNR observed through all 899, or 917 carriers, respectively. In the case of the RC432 pulse shape, *GI*=-2μs, and *GI*=0μs both yield acceptable receiver performance. The 2μs *GI* has the highest SNRs, however this case has very high overhead costs as well. A 2μs *GI* is deemed inefficient with a pulse shape that already has higher frequency domain guard bands (18 fewer carriers for RC432). In the cases of GS432 and PWL372, the lowest *GI* leads to high distortion. Ideally, the receiver should see in the neighborhood of 25-30dB SNRs when noise free packets are fed into it.

3. SIMULATION RESULTS

In Section 2, the system configurations have been identified except for the exact length of *GI* to use in each configuration. In this section we describe how the system performances are evaluated in a simulation environment using our database of a 120 measured channel impulse responses. The diagram of Figure 8 describes the basic simulation scenario.

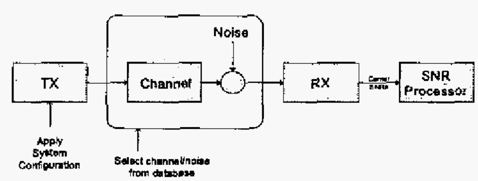

Figure 8. Simulation Environment

A transmitter model is configured with one of the system configurations under test. The transmitter generates QPSK modulated packets that are transmitted through one of the stored channel responses. The noise waveform corresponding to the receive node of the selected channel impulse response is added. A receiver model then implements the signal processing necessary (and common to all configurations) to generate SNR data for each packet, and for each carrier. An SNR processing block, determines average carrier SNRs on each of the paths, converts carrier SNRs to a number of bits supported on each carrier, and finally sums the bits over all carriers and divides by the present configurations symbol time T_s to determine an overall data rate for the channel. Once all channels have been run in the simulation, an overall average data rate is determined by simply averaging the rates over the 120 possible channels.

Note, that since the exact cyclic prefix (or *GI*) lengths are still undetermined they are obtained here by comparing simulations with different values of *GI*. Table 3 summarizes some of the results of these system level simulations. For each pulse shape results are shown at the *GI* length that yields maximum system performance.

Table 3. Simulation Results

Pulse Shape	Num Carr.	*GI* len.	Hanning Window Length	Avg Mbps (120ch)
RC432	899	-1.0μs	1.2μs	98.448
GS432	917	+0.4μs	1.2μs	99.163
PWL372	917	+0.6μs	1.2μs	100.48
PWL372'	917	+0.6μs	1.2μs	101.60

The raised cosine pulse shape shows it's best performance (somewhat surprisingly) at a negative value of *GI*. A negative *GI* implies that the input to the FFT on the receiver's end will have included some of the samples that have been pulse shaped, even if the channel has delay spread of 0μs. The fact that the raised cosine performs well, with such a low *GI* value, indicates that it has a very smooth transition into the unshaped part of the symbol, thus introducing only very little distortion. Nevertheless, the RC432 shape has the disadvantage of supporting only 899 carriers, which it cannot make up, even with the aggressive selection of guard interval length.

The Gaussian pulse shape GS432, requires a positive *GI* length, but even here the length of the guard interval is relatively small compared to the average channel delay spread, indicating that the distortion due to symbol shaping does not totally eliminate the ability of the shaped portion of the symbol to absorb inter-symbol interference.

The best average channel data rate is demonstrated by the piece-wise linear pulse shape. Having had the lowest overhead to begin with (best bandwidth utilization, at shortest pulse shape) it can afford to spend a little more time on guard interval, allowing better overall ISI mitigation. The average data rate of 100.48 Mbps (over 120 channels) represents a 2% improvement over the data rates with RC432. This system is further fine-tuned by optimizing the starting position at which samples are input into the FFT processor (i.e., which 3072 samples out of 3072+789-372=3489 to perform the FFT on). This increases data rates by a further 1.12 Mbps, for an overall improvement of 3.2% over the raised cosine shape. In absolute terms, the data rate has increased by about 3.15 Mbps, which also equates to at least a ½ dB system performance improvement.

4. SUMMARY

This paper addresses primarily the selection of a pulse-shaping waveform for in-home power line communication systems. Pulse shaping becomes particularly important in the HomePlug AV system, because there is a requirement to achieve deep spectral notches without any additional filtering circuitry. Due to the many system impacts that the selection of the pulse shape has, the problem formulation is expanded to also include the simultaneous selection of guard interval length and Hanning window length, creating a joint optimization problem. To pick a configuration that optimizes the overall physical layer system throughput, simulations are employed and run over a collection of 120 power line channel models. The best performing system is the configuration using the piece-wise linear taper that has the lowest intrinsic overhead. It has been verified that the resulting system still has adequate ability to reject narrow band jammers, and despite its short guard interval duration performs well in long delay spread channels.

6. REFERENCES

[1] Kaywan H. Afkhamie, Srinivas Katar, Larry Yonge, and Richard E. Newman, *"An Overview of the upcoming HomePlug AV Standard,"* Proceedings of the 9th International Symposium on Power Line Communications and its Applications, pp. 400-404, Vancouver, British Columbia, April 2005.

[2] Brent Mashburn, Kaywan H. Afkhamie, Larry Yonge, Srinivas Katar, Richard E. Newman, George Peponides, "HomePlug AV Technology -- Part I: Physical Layer and Field Test Performance," Journal on Selected Areas of Communicastions, Submitted for Publication.

[3] HomePlug Powerline Alliance, http://www.homeplug.org

[4] Y. Lin, H. Latchman, S. Katar, and M.K. Lee*, "A Comparative *Performance Study of Wireless and Power Line Networks,"* IEEE Communications Magazine, Vol. 41, No. 4, pp. 54-63, May 2003.

[5] Barnes, J.S., *"A Physical Multipath Model For Powerline Channels at High Frequencies"*, Procceedings of the International Symposium on Powerline Communication and Its Applications, 1998, 76-89.

[6] Minkyu Lee, Haniph A. Latchman, Richard E. Newman, Srinivas Katar, and Larry Yonge, *"Field performance comparison of IEEE 802.11b and HomePlug 1.0"*, 27th Annual IEEE Conference on Local Computer Networks, Pages: 598-599, 2002.