

Implementation and performance analysis of a DMT modem for broadband power line communications on the TMS320C6201

Manuel J. Salmerón-Moreno¹, José A. Cortés², Luis Díez², José T. Entrambasaguas²

¹CETECOM Telecommunications Division
Centro de Tecnología de las Comunicaciones, S.A.
Parque Tecnológico de Andalucía, Málaga (Spain)

²Departamento de Ingeniería de Comunicaciones
University of Málaga (Spain)

Email: msalmeron@cetecom.es, jaca@ic.uma.es, diez@ic.uma.es, jtem@ic.uma.es

Abstract

This paper describes the software implementation of a simple DMT (Discrete Multitone) modem for broadband indoor power line communications based on the Texas Instruments TMS320C6201 digital signal processor at 200 MHz. The structure of the DMT modulator and demodulator is described and a communication initialization procedure suitable for this application is proposed. Key signal processing algorithms, such as FFT, correlation and impulsive response length estimation are discussed and efficient implementation alternatives are proposed. The capacity of the DSP to execute the required tasks is evaluated both in the simulator provided with the Code Composer Studio and in the hardware platform. In this analysis, both the maximum sampling frequency and the memory usage, as a function of the number of carriers, are considered. It is concluded that the maximum transmission rate of the implementation is obtained with 256 carriers which, under ideal channel conditions, leads to a binary data rate of approximately 1 Mbit/s.

1. Introduction

Nowadays there is a growing interest in the use of power networks as a high-speed transmission medium. Two main applications are under consideration: to provide access to wide area networks, the so-called “last mile”, and to serve as a local area network transmission medium for indoor applications [1]. This work concentrates on the latter scenario.

Since the indoor power grid was not originally conceived as a data transmission medium, it presents serious impairments [2]. When bandwidths spanning several MHz are considered, the frequency response of these channels presents large amplitude variations and notches with a priori unknown positions. Noise in power line channels, mainly introduced by the electrical appliances connected in near the receiver end, is neither gaussian nor white. Finally, both its frequency response and noise power spectral density change with time.

In order to overcome these impairments, the DMT modulation [2] is widely considered one of the most appropriate transmission schemes for high-speed power line communications (PLC). This modulation divides the total available bandwidth among a set of orthogonal carriers. The constellation employed in each carrier can be chosen independently, according to its signal to noise conditions. Furthermore, when a high number of carriers is employed, it allows a very efficient frequency domain equalization (FEQ) method that employs just one tap per carrier.

This paper describes the software implementation of a DMT modem prototype on the TMS320C6201 DSP for indoor power lines applications. Section 2 gives a short overview of

the DMT modulator and demodulator and explains the modem initialization procedure. Key algorithms used in the modem are discussed in section 3, while section 4 gives low level details about the implementation on the chosen DSP. Section 5 analyzes the modem performance in terms of the maximum supported sampling frequency and the corresponding binary data rate. Finally, main conclusions drawn from the work are summarized.

2. DMT modem description

2.1 DMT modulator and demodulator

The transmitter and receiver block diagram of a DMT modem with N carriers [3] is shown in Figure 1. The transmitter produces DMT symbols by computing the $2N$ -point IFFT of the N complex symbols transmitted in each carrier plus their hermitic counterpart, needed to obtain a real signal. In order to avoid intersymbol (ISI) and intercarrier interference (ICI), each DMT symbol is preceded by a sort of guard time, called cyclic prefix (CP), consisting of the last part of the DMT symbol. As long as the CP remains longer than the channel impulse response, the complex symbols can be easily recovered undistorted at the receiver side. The implemented modem employs frequency division duplexing. Hence, a filter is finally applied to reduce the out of band emissions.

At the receiver side, the CP is removed before performing the $2N$ -point FFT used to extract the symbols transmitted on each carrier. Provided that neither ISI nor ICI occur, the effect of the channel is easily corrected using a one-tap FEQ. An adaptive FEQ based on the LMS algorithm is employed to track the channel frequency response time variations. Equalized symbols are fed to a detector, whose output is demapped to obtain an estimate of the original bit stream.

The design of a DMT modem should be completed by adding a synchronization mechanism in the receiver that estimates and compensates for the difference between the transmitter and receiver sampling frequencies. The compensation scheme that achieves the highest performance consists of a time domain interpolation filter [4]. The reception filter included in the demodulator models this block.

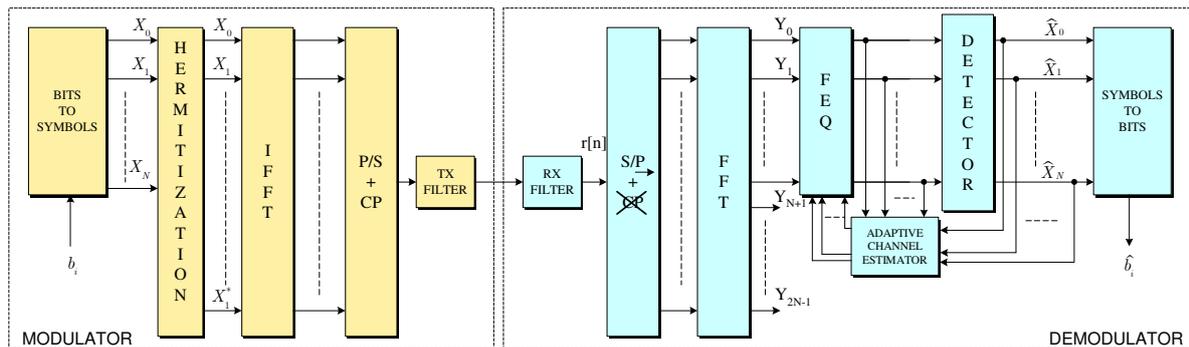


Figure 1: Block diagram of the DMT modulator and demodulator.

2.2 Modem initialization procedure

The modem initialization procedure, schematically shown in Figure 2, is the mechanism used by the two modems involved in the communication process to estimate channel characteristics and decide upon several communication parameters. The designed procedure has been divided in the following phases:

- *Activation*: the modem wishing to start the communications transmits a known DMT symbol repeatedly (activation signal). The other end detects this signal, initially using a power threshold and followed by a correlator whose output should present periodic peaks if the received signal is the activation signal.
- *Frequency error estimation*: a signal similar to the activation signal but having less active carriers is transmitted. It could be used by the receiver to estimate the difference between the sampling frequencies of both modems.
- *FEQ & CP length estimation*: the activation signal and the correlator are used again. The output of the correlator is a periodic repetition of the channel impulsive response, which is used to compute the initial value of the FEQ coefficients and whose length is employed to estimate an appropriate length for the CP.
- *CP decision*: both modems exchange the CP length determined on the previous phase and settle on the use of the greater one. This information is exchanged by using a robust procedure based on a very simple repetition-based coding and BPSK constellations on all carriers.
- *Per carrier BER estimation*: the BER (Bit Error Rate) of each carrier is estimated by transmitting a pseudorandom DMT signal. Those carriers whose BER is below a threshold are discarded while the rest will use a BPSK constellation after the initialization procedure.
- *Carrier load exchange*: the results of the previous phase are exchanged using the same coding and modulation scheme that in the CP decision phase.

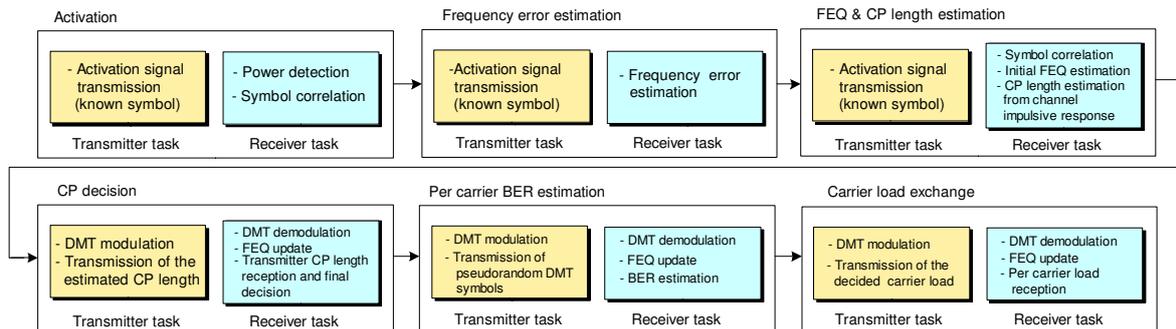


Figure 2: Diagram of the tasks executed during the initialization procedure.

3. Basic DSP algorithms

3.1 FFT/IFFT for real/complex hermitian signals

Since the modem only computes the FFT and IFFT of real signals or complex ones with hermitic symmetry, it is possible to take advantage of a more efficient algorithm [5] that replaces the $2N$ -point FFT/IFFT with an N -point FFT/IFFT and some additional operations. To perform these N -point FFTs/IFFTs the optimized implementations for the 'C62 included in the Texas Instruments DSP Library were used. This library provides the FFT radix-4 algorithm that was used when N was an even power of 2, and the FFT radix-2 algorithm that, because of its reduced efficiency, was only used for values of N that are an odd power of 2.

For a $2N$ -sample real input sequence, $g[n]$, the first step of this algorithm computes the N -point FFT of $x[n]$, built from $g[n]$ as shown in Eq. (1).

$$x[n] = g[2n] + jg[2n+1] \quad (1)$$

Then, the $2N$ -point FFT of $g[n]$, $G[k]$, is obtained applying Eq. (2) to the output of the previous N -point FFT, $X[k]$.

$$G[k] = \frac{1}{2} [X[k] + X^*[N-k]] + e^{-j\frac{\pi}{N}k} \cdot \frac{1}{2j} [X[k] - X^*[N-k]] \quad (2)$$

The $2N$ -point IFFT algorithm for complex signals with hermitic symmetry is analogous but follows the steps in reverse order.

3.2 FFT/IFFT scaling

The implemented FFT algorithm uses the Q.15 fixed point representation format for $g[n]$ and $G[k]$. As a consequence, both $|g[n]|$ and $|G[k]|$, will be equal or less than 1. On the other hand, the modulus of the $2N$ -point FFT of $g[n]$ will be generally less than $2N$ ($|G[k]| \leq 2N$). Thus, a previous scaling of the input of the FFT is needed to guarantee that the output will not overflow. This scaling should be kept as small as possible to avoid an excessive increase in the quantization error.

At the receiver FFT, it can be asserted that $|G[k]| \leq 1$ only if the following conditions meet. Firstly, that the modulus of the symbols transmitted on each carrier is less than unity. Secondly, that the equivalent channel between transmitter and receiver (including the Automatic Gain Control) does not amplify the signal and, finally, that after adding the channel noise, the modulus of the received symbols still meets the first condition. Since these conditions were met, no scaling would be needed in principle. However, for the FFT algorithm discussed in the previous subsection, the possibility of overflow in $X[k]$ must be taken into account. It can be determined that $\max\{|X[k]|\} \leq 2 \cdot \max\{|G[k]|\}$, so $g[n]$ must be divided by 2 (a simple shift on the DSP) before its FFT is computed.

The $2N$ -point IFFT of complex signals with hermitic symmetry was implemented using an N -point FFT according to Eq. (3).

$$x[n] = \text{IFFT}_N\{X[k]\} = \frac{1}{N} \text{FFT}_N^*\{X^*[k]\} = \frac{1}{N_1} \text{FFT}_N^*\left\{\frac{1}{N_2} X^*[k]\right\}; N = N_1 \cdot N_2 \quad (3)$$

Since the algorithm includes a division by N , no overflow would be possible if it is done before the N -point FFT. However, in order to reduce the quantization error, the division is performed in two steps: a division by N_2 previous to the computation of the FFT and a division of the output by $N_1 = N/N_2$. The value of N_2 is selected so that Eq. (4) holds.

$$\frac{\max_{k=0, \dots, N} \{|G[k]|^2\}}{N_2^2} < \frac{1}{(2N)^2} \quad (4)$$

The values selected for N_2 were restricted to powers of two and, hence, performed with binary shifts.

3.3 FFT based correlation

The correlation of the received signal with a DMT symbol is a basic processing in some phases of the initialization procedure. In spite of its great efficiency, the overlap-save algorithm [5] still has an excessively high computational load for this task.

To prevent the initialization procedure from becoming a bottleneck, an approximation was used. The overlap-save needs $2N$ -point FFTs because it computes N samples of a linear convolution by means of a $2N$ -point circular convolution. However, if the input signal repeats

itself periodically every N samples, which is true for the activation signal if noise is ignored, a simpler N -point circular convolution gives the same result. Consequently, the correlator was implemented using N -point FFTs and IFFTs. The noise component of the output remains different from the one of a true correlator, but, being of no interest, it does not constitute a problem.

3.4 Impulsive response length estimation

During the third phase of the initialization procedure, the received activation signal is correlated with one of its symbols. The result is a repetition of the channel impulse response, $h[n]$, every N samples. The CP length should be equal to the length of $h[n]$ to avoid ISI and ICI. However, since power line channels have a very long impulse response, meeting this condition may seriously reduce the symbol rate. Hence, the CP is selected to match the length of the most significant L_h samples of the impulse response, i.e. the shortest segment with a certain percentage of the energy of $h[n]$.

A straightforward algorithm to find the length and initial sample position of this segment, would require checking N^2 combinations. However, since the energy of the segment is monotonically increasing with L_h , a recursive binary search algorithm can be applied so that, at each step, the interval of possible lengths considered when looking for the optimal L_h is halved. Additionally, an efficient search of the maximum energy segment for the value of L_h under consideration can be achieved using Eq. (5), where E_{p,L_h} is the energy of the L_h -sample segment starting in sample p .

$$E_{p+1,L_h} = E_{p,L_h} - |x[p]|^2 + |x[(p+L_h)_N]|^2 = \sum_{n=p}^{p+L_h-1} |x[(n)_N]|^2 - |x[p]|^2 + |x[(p+L_h)_N]|^2 \quad (5)$$

As a result, in the worst case, the number of combinations of p and L_h tested is $N \cdot (1 + \log_2 N)$.

4. Implementation on the TMS320C6201

The modem was implemented on a Hunt Engineering HPEC9 PCI board [6] using a DSP 'C6201 at 200 MHz module. Being only a prototype intended for the evaluation of the DSP performance for this application, all the tests and performance measures were made using a closed loop, so transmitter and receiver were connected inside the DSP.

4.1 Real-time execution structure

The modem is organized around two main DSP/BIOS tasks, transmitter and receiver, that are executed concurrently by the DSP and have the same priority. Communication between them follows the block diagram shown in Figure 3.

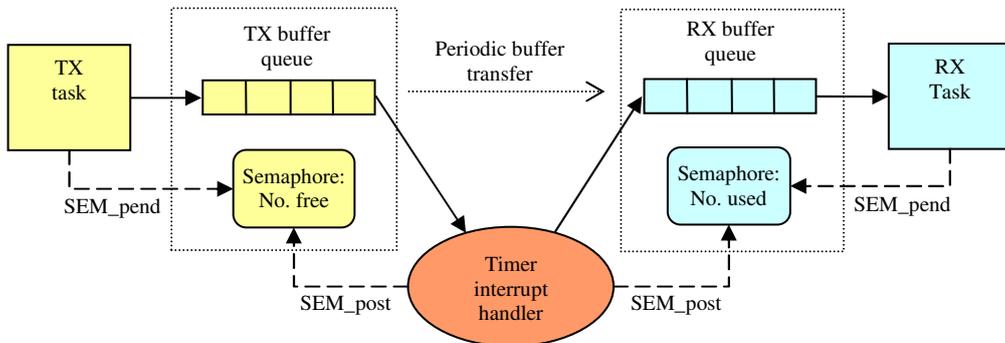


Figure 3: Communication between transmitter and receiver tasks.

The purpose of this structure is to put both, transmitter and receiver, under the same real-time execution conditions that would exist if they were connected to D/A and A/D converters. Each task has an associated buffer queue; the transmitter fills its queue with its output samples and, by means of a semaphore, blocks the task if the queue is filled. In a similar way, the receiver takes its input samples from another queue, having its own semaphore that is blocked when the queue is empty. As both tasks have the same priority, whenever one of them blocks, the other one has the opportunity to execute. Transfers from the transmitter to the receiver queue are driven by a periodic interrupt triggered by one of the internal DSP timers.

4.2 DSP memory usage

As the amount of memory required to store the modem code was less than 64 KB, the IPRAM was configured as normal memory and not as a cache. The IPRAM was only used for the main modem code (tasks and DSP/BIOS), while the rest (start up and Hunt API), not having a significant impact on the performance, was placed on external memory. The total amount of IPRAM used was 44 KB, leaving the 31.25 % free.

The IDRAM was used for data: stacks, precomputed tables, temporary arrays used for signal processing and buffer queues. The default stack size for DSP/BIOS tasks is 1 KB. However, taking advantage of the Kernel/Object View tool in Code Composer after a modem test execution, it was determined that no more than 256 bytes were really used, so data memory was saved shrinking the stacks accordingly.

Almost all the arrays internally used by both tasks, as well as buffers, were sized proportionally to the number of samples of the DMT symbols, or equivalently, the number of carriers. This is the reason why the amount of memory needed for data increases with the number of carriers as shown in Figure 4. For 512 carriers it even exceeds the size of the IDRAM, so the buffers had to be placed in external memory (SBSRAM), while for 1024 carriers it was also necessary to also move the arrays containing the DMT symbols of the activation signal and the synchronization signal.

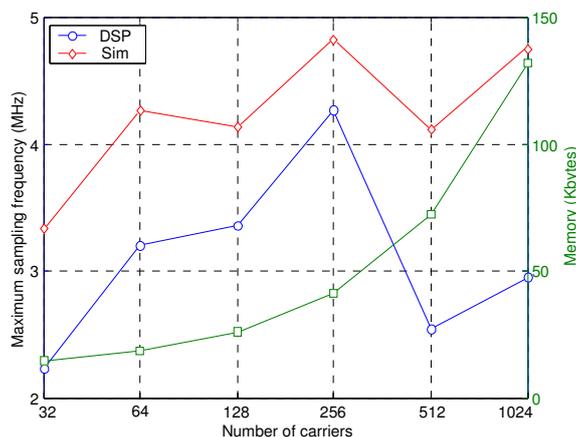


Figure 4: Memory usage and maximum sampling frequency estimated in the hardware platform and in the simulator.

5. Performance analysis

The 'C62 simulator that comes with the Code Composer Studio development environment was used at the first stage to evaluate the maximum achievable sampling frequency and

binary transmission rate. The execution time per DMT symbol for the transmitter and the receiver was measured for each phase of the initialization procedure as well as for the data transmission stage. The resulting numbers of cycles per sample obtained in the worst case, i.e. DMT symbols without CP, are shown in Figure 5 and Figure 6. By summing the values needed to execute the transmitter and receiver process during each stage of the initialization process and the final data transmission, the limiting phase can be identified and the maximum sampling frequency can be estimated. In this case, the limiting phase turns out to be the BER estimation one, leading to the maximum sampling frequencies shown in Figure 4.

Figure 4, Figure 5 and Figure 6 show two clear effects. Firstly, greater sampling frequencies can be achieved when the number of carriers is an even power of two. This is due to the superior performance of the FFT radix-4, employed in the former case, with respect to the radix-2 algorithm used in the latter. Secondly, the maximum sampling rate decreases when the number of carriers, N , is reduced, especially when N is equal or lower than equal 128. This is due to the inferior efficiency achieved by the software pipeline in loops with a small number of iterations. Of the three components of a software pipeline, prolog, epilog and kernel, the latter is usually the most efficient one since it may employ the 8 execution units of the 'C62. Therefore, as the number of iterations is increased, relatively more time is spent in the kernel and loop performance improves.

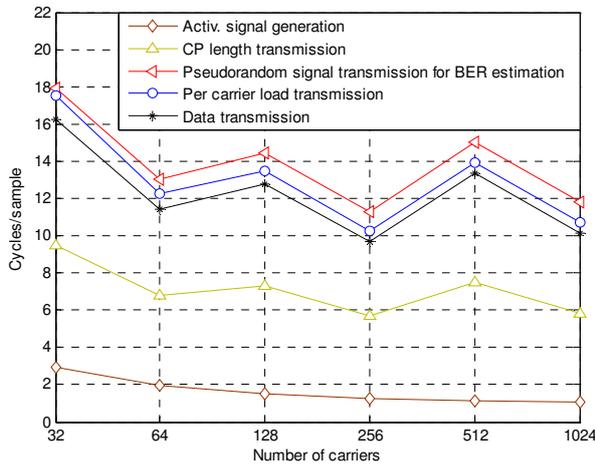


Figure 5: Cycles per sample for the transmitter tasks.

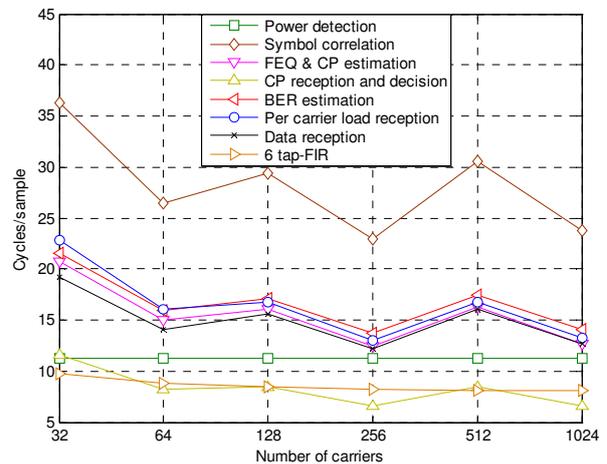


Figure 6: Cycles per sample for the receiver tasks.

In addition to the simulator results, the maximum sampling frequency was measured in the hardware platform [6] employing the structure shown in Figure 3. To this end, an iterative procedure was used. Tests executions were carried out with different timer periods. When the test failed, because the transmitter queue was empty or the receiver queue was full at the time of an interrupt, in the next iteration the period was enlarged, corresponding to a lower sampling frequency. When succeeded, it was shortened, so the tested sampling frequency was increased. Results obtained with this procedure are shown in Figure 4. It can be observed that the sampling frequencies are clearly lower than those derived using the simulator, especially when N is very small (32 carriers) or large (1024 carriers). The differences are mainly due to the transfer times to and from the buffers, which are not considered by the simulator. These times become very important above 256 carriers because buffers were placed in the external memory. On the other extreme, since task switching is more frequent when N small, because queues have less capacity and, hence, become fill or empty quicker, time spent in the task switching increases. Consequently, relatively longer queues and buffers are important to reduce the task switching frequency, even though they increase latency.

Taking into account all the factors that affect the modem performance in the proposed implementation, it can be concluded that 256 carriers is the optimum value to achieve the highest sampling rate, 4.3 MHz. This implies a maximum one-way binary data rate of approximately 1 Mbit/s under ideal channel conditions.

6. Conclusions

This paper has presented the software implementation and performance evaluation of a DMT modem prototype for broadband indoor power line communications on the TMS320C6201. Efficient implementation of the most important signal processing algorithms employed in the modulator and demodulator as well as the modem initialization procedure, such as the FFT and its scaling, correlation and impulsive length estimation have been presented and discussed. The memory usage as a function of the number of carriers has been determined. Performance analysis, in terms of the maximum sampling frequency as a function of the number of carriers, has been evaluated using the 'C62 simulator as well as in the final hardware platform [6]. It has been identified that 256 is the optimum number of carriers, since it allows the use of the most efficient FFT algorithm (radix-4 vs. radix-2), takes better advantage of the 8 execution units when executing loops, it does not need external memory (slower) and it is not so affected by task switching times as when the number of carriers is lower. The maximum estimated binary transmission rate of the developed modem is approximately 1 Mbit/s.

7. References

- [1] ETSI TS 101 867 V1.1.1 (2000-11), "Powerline Telecommunications (PLT): Coexistence of Access and In-House Powerline Systems", ETSI 2000.
- [2] F.J. Cañete, J.A. Cortés, L. Diez, and J.T. Entrambasaguas. "Modelling and evaluation of the indoor power line channel," IEEE Communications Magazine, Vol. 41, pp. 41-47, April 2003.
- [3] J. S. Chow, J. C. Tu, J. M. Cioffi. "A Discrete Multitone Transceiver System for HDSL Applications," IEEE Journal on Selected Areas in Communications, Vol. 9, Issue 6, pp. 895-908, Aug. 1991T.
- [4] Pollet and M. Peeters, "Synchronization with DMT modulation", IEEE Communications Magazine, Vol. 37, Issue 4, pp 80-86, April 1999.
- [5] J.G. Proakis, D.G. Manolakis. "Digital Signal Processing. Principles, Algorithms, and Applications". 3rd Edition, Prentice-Hall International, 1996.
- [6] P. Warnes. "HEPC9 Full Length PCI, Heart based HERON module carrier user manual," Hunt Engineering 2002.