On Non-Integer Bits-per-Symbol Modulation in DMT Modems

DANIEL FRANKLIN and IAN BURNETT
Telecommunications and IT Research Institute
University of Wollongong
Northfields Ave, Wollongong, NSW, 2522
AUSTRALIA
d.franklin@ieee.org http://ieee.uow.edu.au/~daniel

September 16, 2002

Abstract

This paper presents a method for encoding data for transmission in a discrete multi-tone modem using a non-integral number of bits per symbol. This allows a more precise assignment of bits per subchannel than is possible with existing techniques, and is achieved by spreading data over a series of constellations whose size is not a power of two. Advantages and disadvantages of the proposed technique are discussed.

1 Introduction

Discrete Multi-Tone (DMT) modems are increasingly being used in many telecommunications systems, from ADSL modems to the 802.11a and HIPERLAN/2 wireless standards [1], [2]. They use an efficient form of Orthogonal Frequency-Division Multiplexing (OFDM), which performs modulation and demodulation by Inverse and Forward Fast Fourier Transforms (FFTs) respectively [3], [4]. The block diagram of a conventional DMT modulator is shown in Figure 1.

Conventional DMT modems use different symbol alphabets in different subchannels, each with 2_i^N, $N \in \mathbb{Z}$ symbols. The objective of the optimisation algorithm used in most existing DMT modems is to achieve the same Bit Error Rate (BER) across all subchannels [5]. In practice, because of the constraints on the size of the constellation, adjustments of the number of bits per symbol in each subchannel can only be performed in a very coarse manner. When a single bit is added or removed from a particular subchannel, this causes a doubling or halving, respectively, in the size (and, if transmission power is constrained, in the density) of the constellation. Thus, moving a single bit from one subchannel to another in an attempt to even out the BER may cause a dramatic increase in the BER seen in the recipient subchannel. Therefore, it is desirable to be able to more finely control the dimensions of the constellation.

Techniques exist for encoding data for QAM modems using constellations which are not restricted to 2^N points. The simplest technique is to use a multidimensional generalised cross constellation, as described in [6]. This approach partitions the constellation into a number of cosets - for example, in the simplest case, consider a constellation with $2^N + 2^{N-1} = 3 \times 2^{N-1}$ points. In any given QAM symbol period, one data bit is used to select the coset, while another group of either N or $N-1$ bits chooses a point within the coset. In this way an average of $N + 0.5$ bits is transmitted in each symbol period. To spread the coset choice evenly amongst the two cosets, a convolutional code is typically applied to selection bitstream. The concept may be generalised to choose from amongst 2^M cosets.

The disadvantage of this technique, however, is that it does not allow arbitrary constellation size. According to [7], there are modest advantages in hexagonal codes. Further, shell mapping using spherical codes allows for many other constellation sizes in addition to 2^N [8]. We attempt to address this disadvantage with a novel technique based on non-binary representations of information.

2 Non-Integer Bits/Symbol Modulation

The Non-Integer Bits/Symbol (NIBS) scheme outlined here is based on the extension of the conventional DMT modulation algorithm. Standard DMT modulation takes a group of N bits (say, 1024) and breaks it into smaller groups of M_i bits (say, 2-16 bits at a time). Each group of bits is then used to choose a symbol from an alpha-

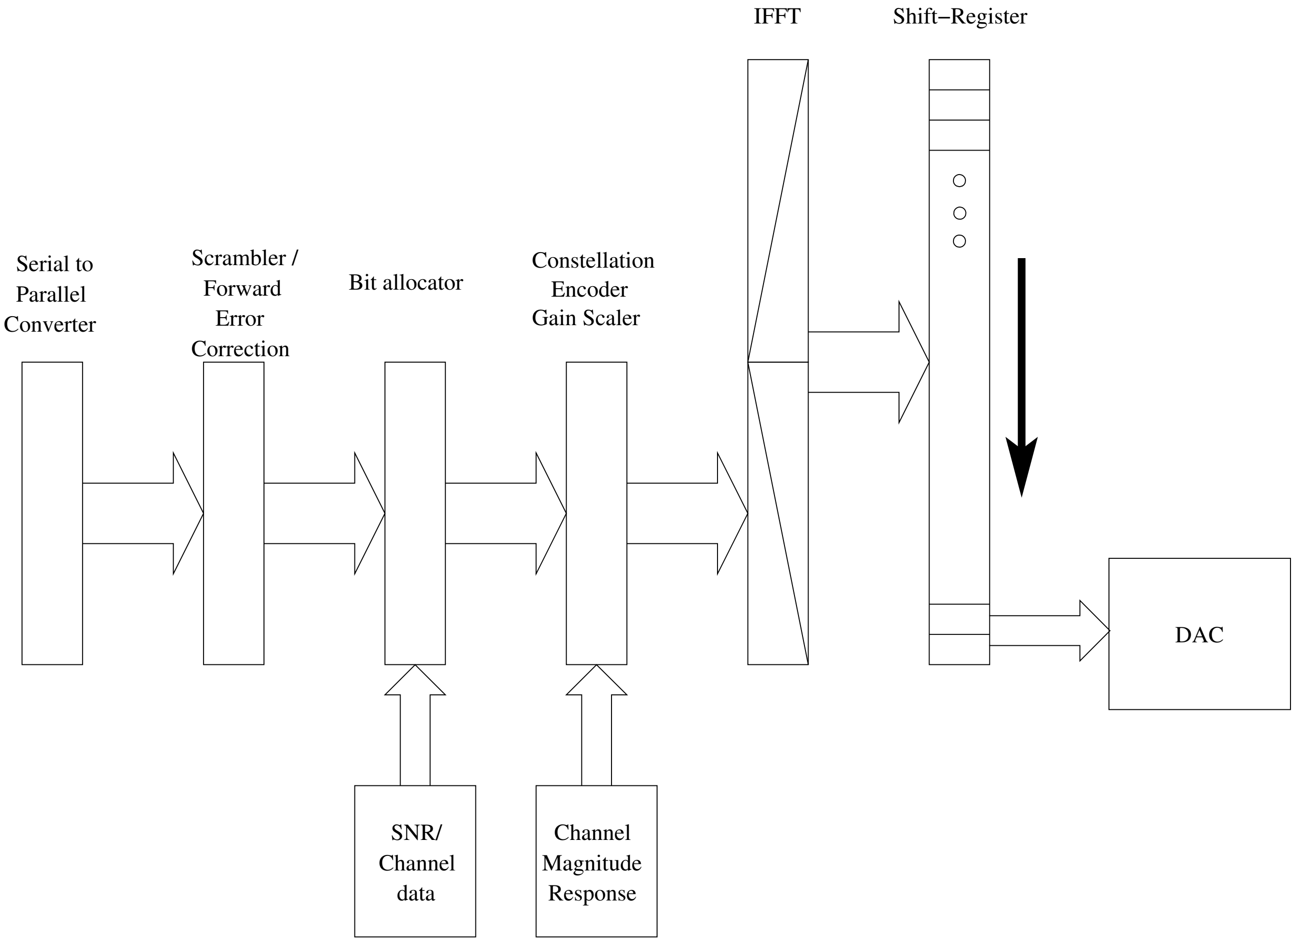

Figure 1: DMT Modulator - Block Diagram

bet of 2^{M_i} possible symbols. A total of K groups of bits can thus be formed, to give K symbols. In general, DMT modems use a set of constellations of complex numbers as the symbol alphabets, one constellation for each different possible value of M. The group of K complex numbers then populates the bottom half of a complex vector, the top half of which is constructed so as to provide Hermitian symmetry. This vector is used as the input to an inverse Fast Fourier Transform, which efficiently performs bulk modulation of the K carriers. The output of the IFFT is a set of time-domain samples which may be directly converted to an analog signal.

Therefore, we may think of the DMT encoding process as a way of representing an N bit number as a sequence of numerical digits, with each digit having a different base. With a standard DMT system, the choice of bases is limited to 2^M. The NIBS scheme extends this to a system which allows constellations whose alphabet size is *any* integer in a given range. To match the range of alphabet sizes in a conventional DMT system, we may say the integer must be between 4 and 65536. The effective number of bits which may be contained in each symbol selected from an alphabet of size i is $\log_2 i$, which in general is not an integer. From a practical point of view, this increases the number of possible constellations from 15 to 65532 for the case of the ADSL DMT implementation. For practical reasons, the actual increase may be smaller than this, since it may not be possible to construct or store all of these potential constellations.

3 Advantages of NIBS

The main advantage offered by the NIBS scheme is the ability to get closer to the ideal case of a uniform BER across all subchannels in a DMT modem. The coarseness of the adjustments required in standard DMT systems limits the effectiveness of the water-pouring algorithms which are used to allocate bits to subchannels. Although NIBS is restricted to a finite number of constellations, the flexibility it allows is considerably greater than that available with conventional DMT modulation.

4 Problems and Proposed Solutions

The NIBS encoder introduces a number of complications in the encoding process. The first and most important is that is is no longer possible to simply break up the 1024 bit number into arbitrary groups of of bits. The approach used to encode the number is the same algorithm as used for changing the digits of a number from one base to another. However, for this to work, it is necessary to have the least-significant digit (LSD) encoded with the largest of our chosen bases (alphabet sizes), and use smaller and smaller bases as we progress to the most-significant digit (MSD). This would at first seem to destroy the primary benefit to using such DMT, since it limits the choice of allocation of constellations to subchannels. However, this is not really the case, since the actual order of storing the 'digits' (i.e. in the input vector for the IFFT) is unimportant. Thus, for the purpose of breaking up the 1024 bit number, the base used for each digit in the mixed-base representation can be ranked from largest to smallest (going from LSD to MSD), and the digits can be re-arranged arbitrarily to suit the SNR per subchannel conditions of the channel. As long as the mapping is known, and can be communicated to the receiver, the transmitted number can be reconstructed in full.

A second problem is that it is not possible to perfectly represent a set of (say) 1024 bits using anything other than integer groups of bits without some 'wasted' capacity. For example, suppose we wish to transmit a 16-bit number using constellations containing 2, 3, 4, 4, 4, 5, 5 and 7 points respectively. The largest number which can be represented in this way would be written in mixed-base form as 12333446, or 67199_{10}. Thus, our set of digits is sufficient to store 16 bits but not 17 bits - and there are a number of possible symbol configurations which are not possible (in the range 65536 to 67199). However, the system still enjoys the benefits of being able to more finely tune the dimensions of the constellation to the SNR conditions of the channel.

The following example illustrates the procedure - suppose we wish to represent the number 47296_{10} using digits with the bases 2, 3, 4, 4, 4, 5, 5 and 7 (call it base k_i):

$$\begin{aligned}
47300 = & \; 0 \times (1 \times 7 \times 5 \times 5 \times 4 \times 4 \times 4 \times 3 \times 2)+ \\
& + 1 \times (1 \times 7 \times 5 \times 5 \times 4 \times 4 \times 4 \times 3)+ \\
& + 1 \times (1 \times 7 \times 5 \times 5 \times 4 \times 4 \times 4)+ \\
& + 0 \times (1 \times 7 \times 5 \times 5 \times 4 \times 4)+ \\
& + 3 \times (1 \times 7 \times 5 \times 5 \times 4)+ \\
& + 2 \times (1 \times 7 \times 5 \times 5)+ \\
& + 1 \times (1 \times 7 \times 5)+ \\
& + 1 \times (1 \times 7)+ \\
& + 2 \times (1) \\
& = 11032112_{k_i}
\end{aligned}$$

The order of these symbols is unimportant, thus for systems with good SNR at low frequencies and poor SNR at higher frequencies, we may assign the symbols to subchannels such that the 7-symbol-alphabet digit is given the best subchannel, and the 2-symbol-alphabet (i.e. 1-bit) digit is given the worst.

While the two problems described above relate to theoretical aspects of the modulation scheme, the third problem is related to the practicality of implementation. If we are to allow up to 65532 ($2^{16} - 4$) possible constellations, then it is necessary to store these numbers at both the transmitter end and receiver end. This would require enough storage to hold approximately 2^{31} constellation points, or 2^{32} single-precision floating-point numbers. If a single-precision floating-point number is used to store these values, this will require 2^{36} bytes of storage capacity - obviously an impractically large amount of data. Also, this ignores the difficulties of actually constructing all of these constellations.

Fortunately, many regular constellations can be constructed algorithmically - therefore, once both transmitter and receiver have decided upon the size of each constellation, it should be possible for them to design constellations of this size without the need to store all possible constellations. These can then be stored in a relatively small lookup table as is done with conventional DMT modems. Alternatively, by restricting the maximum size of the constellation to a smaller dimension (for example, 1024 symbols), a smaller capacity is required (the 16-bit constellation is only be rarely used in practice) which should make storage of the full set of available constellations possible. Also, it may be possible to construct a compromise system which only includes some constellations which are easy to construct algorithmically, such as staggered-concentric-circle APSK, square/pseudo-square as used in conventional DMT (see Figure 2(a)), quincunx (see Figure 2(b)), hexagonal lattice (see Figures 3(a) and 3(b)) and so on. One particularly interesting aproach, taken from

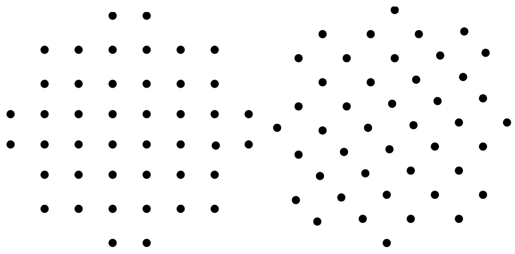

(a) Rectangular

(b) Quincunx

Figure 2: 44-point rectangular lattices

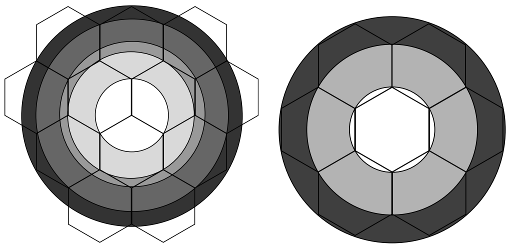

(a) Hexagonal, showing first four shells (6, 12 and 24 points)

(b) Hexagonal + $R\angle\frac{\pi}{6}$, showing first 5 shells (3, 9, 12, 18 and 24 points)

Figure 3: Hexagonal lattices

the field of vector quantisation (as applied to speech coding) is described in [9].

The final problem associated with NIBS encoding is the computational complexity of such a system - namely, the large number of multiplications required in the base conversion process. Fortunately, this task is well suited to vector processing, and hence implementation on a fixed-point DSP chip should be feasible.

5 Conclusion

DMT modulation based on non-integer bits-per-symbol constellations should permit a more accurate distribution of data amongst noisy subchnnels. Although there are a number of problems in the practical realisation of such a scheme, these can be mitigated or eliminated by careful design of the coder.

References

[1] T1E1.4 Working Group, *ANSI T1413-1999 Network and Customer Installation Interfaces - Asymmetric Digital Subscriber Line (ADSL) Metallic Interface*, American National Standards Institute, 1999.

[2] A. Doufexi, S. Armour, M. Butler, A. Nix, D. Bull, J. McGeehan, and P. Karlsson, “A comparison of the HIPERLAN/2 and IEEE 802.11a wireless LAN standards,” *IEEE Communications Magazine*, vol. 40, no. 5, May 2002.

[3] S. B. Weinstein and P. M. Ebert, “Data Transmission by Frequency Division Multiplexing Using the Discrete Fourier Transform,” *IEEE Transactions on Communications Technology*, vol. COM-19, pp. 628–634, Oct. 1971.

[4] B. Hirosaki, “An Orthogonally Multiplexed QAM System Using the Discrete Fourier Transform,” *IEEE Transactions on Communications*, vol. COM-29, pp. 982–989, July 1991.

[5] John A. Bingham, “Multicarrier Modulation for Data Transmission: An Idea Whose Time Has Come,” *IEEE Communications Magazine*, May 1990.

[6] Jr. G. David Forney and Lee-Fang Wei, “Multidimensional Constellations - Part I: Introduction, Figures of Merit, and Generalized Cross Constellations,” *IEEE Journal on Selected Areas in Communications*, vol. 7, no. 6, Aug. 1989.

[7] G. R. Lang, Jr G. D. Forney, S. Qureshi, F. M. Longstaff, and C. H. Lee, “Signal Structures with Data Encoding/Decoding for QCM Modulations,” *U.S. Pat. 4 538 284*, Aug. 1985.

[8] N. J. A. Sloane, “Tables of Sphere Packings and Spherical Codes,” *IEEE Transactions on Information Theory*, vol. 27, pp. 327–338, 1981.

[9] M. A. Ireton and C. S. Xydeas, “On improving vector excitation coders through the use of spherical lattice codebooks (SLCs),” in *ICASSP-89*, 1989, vol. 1, pp. 57–60.