

DATA TRANSMISSION

This invention relates to data transmission and in particular to data transmission in which a multi-level representation of a digital signal is transmitted.

5 Digital data represents a signal that has been quantised in both time and amplitude. The digital data thus approximates the actual value of an analogue signal. If an analogue signal is digitised, the range of the analogue signal is divided into a number of levels e.g. 16, the analogue signal is sampled at set intervals and the appropriate level at that instant is determined. Since only 16 levels are used,
10 the level which is closest to the actual level is chosen. A signal that is reconstructed from this digital data resembles but is not exactly the same as the original analogue signal.

 In binary digital data, a signal is represented by 0 or 1, a 0 being a 0 V pulse and a 1 being a 5 V pulse for example. If a sample of an input signal has an
15 amplitude greater than half of the maximum range, the signal sample is represented by a 1. A 0 will result from samples having an amplitude less than half. Thus a series of ones and zeros is produced.

 To reproduce the original signal, a receiver needs to know the threshold between the two levels. This is usually done by the transmitter signalling to the
20 receiver the maximum range of the signal, the number of levels and the spacing of the levels e.g. if they are linearly spaced etc. The receiver then determines the threshold(s) and decodes an incoming signal. In practice, in a multi-level system there are distortions in the signal due to the response of the network, namely overshoot and ringing. The instantaneous level of any received sample is not only
25 dependent on the transmitted sample, but on the recent previously transmitted samples and possibly the subsequently transmitted sample or samples.

 According to the present invention, there is provided a data transmission system comprising:

 a transmitter having means to encode an input signal to form coded data,
30 each element of said coded data having one of at least two discrete signal magnitude levels, the encoding means including in the coded data a periodic training sequence of data; and

a receiver to receive the coded data and, on the basis of the received training sequence, to adapt a threshold or thresholds to allow the discrete levels to be distinguishable from each other.

Such a system allows the receiver continually to adapt the threshold(s) in dependence on the dynamic conditions of the transmission link between the transmitter and receiver.

Preferably the training sequence comprises at least two elements, the receiver including means to monitor the effect of at least one of the elements of the training sequence on another of said elements of the training sequence and adapt the threshold(s) accordingly.

Thus the effects of the transmission link on groups of elements is accounted for.

The invention also relates to the transmitter and a receiver.

The invention will now be described by way of example only with reference to the accompanying drawings in which:

Figure 1 shows a data transmission system according to the invention;

Figure 2 shows an example of the line format of coded video data transmitted by the data transmission system of Figure 1;

Figure 3 shows a transmitter according to the invention;

Figure 4 shows a receiver according to the invention;

Figure 5 shows an example of the contents of a buffer after eleven lines of the training sequence have been received;

Figure 6 shows an example of a set of thresholds, generated from the contents of the buffer as shown in Figure 5; and

Figure 7 shows an example of a look-up-table generated from the thresholds of Figure 6.

As shown in Figure 1 the digital data transmission system comprises a transmitter 20, a receiver 30 and a communications link 40. Data is transmitted from the transmitter 20 to the receiver 30 via the communications link 40 which may take any suitable form. For example the communications link 40 may be part of a Public Switched Telephone Network (PSTN), a dedicated line such as provided by an Integrated Services Digital Network (ISDN), a radio link, coaxial cable, optical fibre etc.

For illustrative purposes, the data transmission system to be described refers to the transmission of data representing a video image. However the invention is applicable to any system that transmits multi-level digital data, in particular for transmission over an analogue transmission link e.g. cable modems,
5 higher bit rate teletext services.

The data transmission system to be described is suitable for distributing digital television signals to customers over an analogue hybrid fibre-coax network. In order to use existing analogue network infrastructure, the digital signals must be capable of being transmitted over the existing network in the same way as an
10 ordinary TV channel.

The digital signal must therefore occupy a similar bandwidth to the ordinary TV channel (6-7 MHz). It must also 'look' like a TV signal in terms of amplitude and have regular 'line' sync pulses at 15.625 kHz, as some dc restoration in the network relies on this. Signal to noise ratio is around 50dB and
15 there are non-linearities such as differential gain errors and sync pulse clipping to overcome.

In order to obtain reasonable capacity improvements over analogue, sufficient digital capacity to carry four multiplexed MPEG video streams of acceptable quality were desired.

20 The invention uses multi-level coding, whereby serial digital data is split into symbols of n bits. Each symbol is then coded as one of 2^n discrete levels within the active video region.

By choosing the symbol rate to be 6.75 MHz the minimum pulse width encountered is $1/6.75 \times 10^6$ i.e. 148 ns. This should pass through the bandwidth of
25 the system without too much degradation. Given the system signal to noise ratio, it would be reasonable to expect to be able to recover eight distinct levels, giving three bits per symbol (i.e. $n=3$).

Experiments have shown that it is not really necessary to maintain any frame timing as all clamps and dc restorers in the system use only the line syncs,
30 and black levels. This means the frame timing can be omitted and a continuous stream of active 'lines' used. This increases the data throughput and simplifies the transmitter and receiver design.

By choosing a master clock rate of 27 MHz, readily available TV sampling clock recovery chips can be used to provide a reasonably low jitter line locked clock at the receiver. The signal can be oversampled at this rate to determine the best sampling position. Readily available video analogue-to-digital converters
5 (ADCs) can be used as well as reasonable speed logic.

A diagram of the waveform to be transmitted is shown in Figure 2. It consists of a synchronisation (sync.) pulse 2 having a standard width and amplitude, repeating at 15.625 kHz. This is preceded and succeeded respectively by a front porch 4 and a back porch 6 to allow readily available TV sync pulse
10 separators and analogue to digital converters to be used.

After the back porch 6 is a start pulse (S) which is used by the receiver to determine the best sampling position. There then follow eight symbols, five of which form a training sequence (T) which step through a set sequence over a number of lines. The first of these symbols (M) is a marker to allow the start of
15 the training sequence to be determined. The exact nature and function of the training sequence will be described later.

There then follow a number of valid data symbols D. Each is nominally 148ns wide and is represented by one of eight distinct levels, nominally 0.1V apart. The valid data can be split up into blocks to allow the addition of a block
20 based forward error corrector (FEC). One overhead in systems which use block based FEC is the need to add framing bits to define the block boundaries, and the hardware at the receiver to search for and lock onto the framing. This is not necessary in this scheme as the data is already divided into 'lines' which can be further subdivided into blocks.

25 The choice of error corrector block size and total number of symbols per line depends on the required bit rate and the correcting power of the FEC.

The proposed system uses a BCH (Bose-Chaudhuri Hocquenghem) forward error corrector and divides the line into 17 blocks of 63 bits. Each 63 bit block contains 21 3-bit symbols, consisting of 19 symbols of data (57 bits) and 2
30 symbols of check bits (6 bits), giving a payload bitrate of $57 \times 17 \times 15625 = 15.140625$ Mbit/sec.

For the sake of having a 'round' bitrate, the last block has three symbols which are not filled with data, giving $((57 \times 17) - 9) \times 15625 = 15.00$ MHz. It is felt

that this is a reasonable rate into which four MPEG encoded TV channels can be multiplexed, giving the required quality per channel.

The FEC is capable of correcting one bit in error in each block. In order to reduce the likelihood of multiple-bit errors the symbols are Gray coded so that adjacent levels represent bit patterns with only one bit difference.

An example of the transmitter 20 is shown in Figure 3. The transmitter 20 can either be a slave to the MPEG multiplexer clock, or a master clock provider. The phase locked loop (PLL) and clock generator 201 generates the 6.75MHz symbol clock locked to the 15MHz data bit clock.

10 The incoming binary digital data is split into 3-bit symbols by a serial-in-parallel-out (SIPO) shift register 202 and stored in a first in, first out (FIFO) buffer 204. The FIFO 204 buffers the symbols between the continuous input data rate and the 'bursty' line and block structure. The symbols are read from the FIFO 204 and BCH FEC check bits are added by FEC encoder 206. At the start of each line, 15 the sync pulse, black level (i.e. front porch and back porch), start pulse (S) and training sequence are added by a unit 207 under control of a control block 208. The data is then Gray coded and converted to an 8-bit representation 210 before being presented to the digital-to-analogue converter DAC 212. A certain amount of pre-compensation can be added at this stage to help reduce overshoot in the 20 network. This effectively reduces the rise-time of the edges within the signal. The analogue output of the DAC 212 can then be transmitted across the network in the same way as a normal TV channel. An analogue post filter can be added if necessary to band limit the signal to suit the network.

A block diagram of the receiver 30 is shown in Figure 4. A sync. 25 separator 301 extracts sync. and black level pulses from the incoming signal. A PLL and voltage-controlled crystal oscillator (VCXO) 302 generates a line locked 27 MHz clock.

An 8-bit ADC 303 digitises the incoming signal into an 8 bit signal. The ADC 303 has an on-chip clamp and automatic gain control (AGC) which uses the 30 sync. and black level pulses. The effect of the AGC is to set the digital output at the base of the sync. pulse 2 to 0 and the black level (i.e. the level of the front and back porch 4 and 6) to 63. A nominal maximum level video input of say 0.7V is then given a level of 213. As the AGC gain is calculated using the sync. pulse

height, any clipping of the sync. pulse by the transmission network will alter this nominal maximum height. The ADC is clocked at 27 MHz.

Logic 304 which follows the ADC 303 contains a state machine which searches for the start pulses (S) after each line sync pulse 2. The start pulse S will
5 be rounded after having passed through the network and the best sampling position is taken as the sample closest to its peak. An inverted version of the clock can be used if this gives a better sample position.

The training sequence will now be described. The training sequence is 1024 lines long and occupies five symbols at the start of each line between the
10 start pulse S and the valid data symbols D. The first symbol (M) in the sequence indicates the start of the training sequence; it is high (level 7) on the first line of the sequence and low (level 0) on all the other lines. The next three symbols T_1 , T_2 , T_3 , symbols count through a sequence of all the possible combinations of eight levels for the three symbols (8^3 combinations), line by line, with the last symbol T_4
15 low (level 0) or high (level 7) giving 2×8^3 i.e. 1024 combinations, occupying 1024 lines (approx. 65 mSec.)

At the receiver, the level of the fourth symbol T_3 in the training sequence of each line is sampled and stored in a FIFO 306. Thus after 1024 lines the FIFO 306 contains examples of all levels of the fourth symbol with all combinations of
20 the two preceding symbols T_1 , T_2 and the succeeding symbol T_4 being low or high. Figure 5 shows an example of the contents of FIFO 306 after 11 lines of data. A microprocessor 308 calculates a set of seven decision thresholds for each of the combinations of preceding and succeeding levels and generates a look-up-table (LUT) which is stored in SRAM 312. For example, samples 1-8 in the FIFO 306
25 represent the levels of the fourth symbol of the training data, T_3 , when both of the preceding symbols T_1 and T_2 and the succeeding symbol T_4 are at level 0. The microprocessor 308 thus calculates the seven decision thresholds to be applied when the preceding two symbols are zero and the succeeding symbol is zero. This is typically achieved by setting each threshold for T_3 halfway between the two
30 received training levels i.e. $\text{threshold} = L_1 + [(L_2 - L_1)/2]$ where L_1 and L_2 are the received levels for successive T_3 symbols. Figure 6 shows an example of the thresholds for this example case, as stored in the microprocessor's RAM, 310.

The microprocessor then uses this set of thresholds to calculate a LUT, as shown in Figure 7, and stores it in SRAM, 312.

The LUT is then used to perform thresholding of the valid data D in real time. The 8-bit input data is applied to the LUT 312, via input a. The previous two samples of the input data are input to inputs b and c respectively. The level of the subsequent sample of the input data is input via input d. This input d is a simple high/low indication derived from the sample before the one presented to input a of the comparator 312. For valid data, the succeeding sample may have any value between the maximum and the minimum (213 and 0 respectively in this embodiment). A notional threshold is set midway between the maximum and the minimum. If the value of the subsequent sample is above this threshold, the value of the subsequent sample is considered to be high; if the value is below, it is considered to be low. Inputs b and c can be taken as latched outputs from the LUT 312 as then they have been quantised and are only three bits each, reducing the required LUT size.

In practice the LUT comprises two banks of SRAM. Once the microprocessor has calculated a LUT and written it to the SRAM, it 'pages' that LUT into the real-time data path. It then performs the whole cycle again, capturing a new FIFO full of training data and recalculating a set of thresholds. These can be averaged with the previous set to reduce the effects of random noise and a new LUT calculated. This is then paged-in in place of the previous LUT. The process repeats like this as fast as the processor can perform the tasks. Thus the system adapts to the response of the communications link 40 and tracks any long term changes in the response.

As a further enhancement, the microprocessor can use the samples stored in the FIFO 306 to measure the pulse response of the link. The pulse response may indicate that the level of any sample is more dependent on the subsequent sample rather than the sample two-previously, which may happen if the bandwidth of the communications link is low. The training sequence can then be sampled at its third symbol rather than its fourth and the input to the LUTs can be changed to input more bits of the subsequent sample rather than the two-previous sample. The processor would then have samples of all the combinations of previous and subsequent samples and could generate LUTs in a similar way.

The thresholded 3-bit symbols are then inverse Gray coded and passed through the BCH FEC detector/corrector 314 which corrects any single bit errors in each 64 bit block. The data then passes through a rate conversion FIFO 316 and is re-clocked out by control 318 at a continuous 15.0 MHz. This is passed to an

5 MPEG demultiplexer/decoder for decoding in a conventional manner.

**CLAIMS****THIS SHEET HAS BEEN
REPLACED**

1. Data transmission system comprising:
 - a transmitter having means to encode an input signal to form coded data,
 - 5 each element of said coded data having one of at least two discrete signal magnitude levels, the encoding means including in the coded data a periodic training sequence of data; and
 - a receiver to receive the coded data and, on the basis of the received training sequence, to adapt a threshold or thresholds to allow the discrete levels to
 - 10 be distinguishable from each other.
2. A data transmission system as claimed in claim 1, wherein the training sequence comprises at least two elements, the receiver including means to monitor the effect of at least one of the elements of the training sequence on another of
- 15 said elements of the training sequence and adapt the threshold(s) accordingly.
3. A data transmission system as claimed in claim 1 or 2, including means to generate a look-up-table according to the levels of the elements of the received training sequence and the associated adapted threshold(s).
- 20
4. A data receiver comprising:
 - an input for receiving input data, each element of said input data representing one of at least two discrete signal magnitude levels, said input data including a known training sequence, and
 - 25 means to adapt a threshold or thresholds to allow the discrete levels to be distinguishable one from another.
5. A data receiver according to claim 4, wherein the training sequence comprises at least two elements, the adaptive threshold means including means to
- 30 monitor the effect of at least one of the elements of the training sequence on another of said elements of the training sequence and adapt the threshold(s) accordingly.



6. A data receiver according to claim 4 or 5, including means to generate a look-up-table according to the levels of the elements of the received training sequence and the associated adapted threshold(s).
- 5 7. A data transmitter having means to encode an input signal to form coded data, each element of said coded data having one of at least two discrete signal magnitude levels, the encoding means including in the coded data a periodic training sequence of data.
- 10 8. A data transmitter according to claim 7, wherein at least one element of the periodic training sequence represents, in turn, each of the discrete signal magnitude levels.
9. A data transmitter according to claim 8, wherein the periodic training
15 sequence comprises at least two elements each of which, in turn, represent each one of the discrete signal magnitude levels.

This Sheet Has Been Replaced

