



G00002



159159

Protocols for transmission of data, in particular over telephone lines

Field of the invention

The present invention relates to methods for transmitting data, in particular over telephone lines (typically, copper telephone lines) or similar lines. It
5 further relates to systems arranged to perform the methods.

Background of Invention

The use of fast Internet connections has grown rapidly over the last few years, and consequently the demand for broadband (high-speed) connections is
10 increasing.

One technology which is very well known in the market is Asymmetric Digital Subscriber Line (ADSL) technology. This employs the frequency spectrum indicated schematically in Fig. 1. "Upstream" communications (that is in the
15 direction from the home or office user premises, "customer premises equipment" or "CPE", to the "central office", or "CO") are transmitted on frequencies in the range 25kHz (i.e. above the maximum audible frequency of 4kHz) to 138kHz. "Downstream" communications are in a higher frequency band from 138kHz to an upper limit. According to the first two versions of
20 ADSL (ADSL and ADSL2) the downstream band goes up to 1.1MHz, whereas in ADSL2+ it goes up to 2.2MHz. Within each of the upstream and downstream bands, the range is divided into 4kHz intervals, "tones", so that the downstream band includes 256 tones in ADSL and ADSL2 (which is capable of transmitting 8MBps), and 512 tones in ADSL2+ (which is capable
25 of transmitting 28MBps). Each tone is encoded by quadrature amplitude modulation ("QAM"), and can encode between 0 and 15 bits. During a training phase, the line conditions (signal to noise ratio, SNR) of each of the tones is

estimated, and the number of bits which will be encoded in each tone during each frame is selected.

5 In a typical ADSL modem, the main sections are (i) a Digital Interface (which may use asynchronous transfer mode (ATM), (ii) a Frammer (also referred to here as a framing unit), (iii) an DMT Modulator, and (iv) the AFE (Analog Front End).

10 The framer multiplexes serial data into frames, generates FEC (forward error correction), and interleaves data. FEC and data interleaving corrects for burst errors. This allows DMT-based ADSL technology to be suitable for support of MPEG-2 and other digital video compression techniques. For the transmit signal, the Encoder encodes frames to produce the constellation data for the DMT Modulator. It assigns the maximum number of bits per tone (based on
15 measured SNR of each tone) and generates a QAM constellation where each point represents a digital value. Each constellation point is one of N complex numbers, $x + iy$, where x and y are the phase and amplitude components. The summation of bits in all carriers, multiplied by the frame rate (4kHz), represents the data rate. For the receive signal, the framer includes a decoder
20 which converts QAM symbols back into the data bitstream.

In the DMT Modulator, a frequency domain processor implements FFT/IFFT and associated processing. In the transmit path, the Inverse Fast Fourier Transform (IFFT) module accepts input as a vector of N QAM constellation
25 points and duplicates each carrier with its conjugate counterpart so the 2N output samples are real. The 2N time domain samples have the last 2N/16 samples appended as a cyclic extension (which may include a cyclic suffix, a windowing function and/or a cyclic prefix extension), and are then delivered to a DAC (digital to analogue converter). The set of time domain samples
30 represents a summation of all the modulated sub-channels, for the duration of

- one data frame. In the receive path, the first $2N/16$ samples from the ADC (cyclic prefix) are removed. A FFT module transforms the carriers back to phase and amplitude information (N complex QAM symbols). Correction for attenuation of the signal amplitude and phase shifts (i.e. overall distortion) is implemented. If the QAM constellation is thought of as points in a grid where rows and columns represent phase and amplitude information respectively, then the grid effectively rotates reference to the constellation points to correct for these distortions.
- ADSL2 also include a feature known as "tone ordering". Based on the SNR which has been established for the tones, they are classified based on the SNR such that a "path" is selected for each tone through the encoding device, and each of the tones is transmitted along to the framing unit through the corresponding selected transmission path. This is illustrated in Fig. 2(a), in which the framing unit 1 receives data along two paths 2, 3. Each path 2, 3 leads to a respective block 4, 5 which constructs respective portions of frames. The frame is shown in Fig. 2(b), including a portion 6 generated by block 4, and a portion 7 constructed by a block 5 (which may be an interleaver). The outputs of the blocks 4, 5 are stored respectively in a fast buffer 8 and interleaved buffer 9, until they are transmitted out of the framing unit 1. Since the interleaver 5 interleaves data over a period of time, data transmitted along path 3 will have a different (higher) latency than data transmitted along the path 3. Thus, these two paths are referred to as different "latency paths" (e.g. they may be referred to as LP1 and LP2). Note that if the SNR of the tones is changed, the association between latency paths and tones is changed also. The trellis encoding is performed on the output of the framing unit 1.
- Furthermore, the number of bits which are transmitted by each of the tones may be modified if the estimated SNRs of the tones are revised: increasing

the number of bits stored per frame in some tones and correspondingly reducing the number of bits stored per frame in other tones. This process is known as "bit swapping".

- 5 For further details of the ADSL2 standard the reader is referred to the document ITU-T Recommendation G.992.3 published by the International Telecommunication Union).

While ADSL provides Internet connections that are many times faster than a
10 56K modem, it still are not fast enough to support the integration of home services such as digital television and Video-on-Demand. However, another DSL technology known as very high bit-rate DSL (VDSL) is seen by many as the next step in providing a complete home-communications/entertainment package.

15

In contrast to ASDL, a conventional VDSL standard (here referred to as VDSL1) uses a number of bands, e.g. as shown in Fig. 3, which may go up to for example 12MHz. Frames are typically 16 times larger than those of ADSL, e.g. including 8k samples. VDSL has a number of further differences from
20 ADSL. For example, VDSL1 has different framing method from ADSL2 (for example, with no sync symbol). It does not include Trellis encoding. Its interleaving system is different. In the VDSL1 system, the tone ordering is applied to all the tones used for communication in a given direction. Up until now two sets of memories are required on chip. If this feature is incorporated
25 into future versions of VDSL, here referred to as VDSL2, with 4k tones or higher, each of the bit allocation table, gain tables, tone ordering tables each for 4k tones requires significant on-chip memory.

Summary of the Invention

The present invention aims to provide new and useful protocols for transmitting data through lines such as telephone lines. Typically these protocols have transmission rates of over 24Mbps, and often much higher.

- 5 A first aspect of the invention proposes in general terms that the tones of a signal are grouped, and that Trellis encoding together the data which is to be transmitted on the tones of each group.

The concept can be generalised to give a second aspect of the invention, which is that a portion of data is associated with each of the groups of tones,
10 and that the tones used together to transmit the corresponding portion of data (e.g. in a way determined by a measured SNR of each tone, or possibly by a selection made by the unit receiving the transmitted data).

A selection may be made from the tones of the group such that each of the tones of the group is allotted a respective role in transmitting the data.

- 15 Thus, a second aspect of the invention proposes in general terms that any one or more of tone ordering, bit swapping, selection of gains for each tone (e.g. gain scaling) and/or selection of paths for each tone is done within the members of a group. This idea is applicable both to tone ordering etc following the training stage, and also to the dynamic configuration changes
20 subsequently, for example bit swapping.

Implementations of each of the first and second aspects of the invention may offer two advantages (compared to treating all the tones of a given direction equivalently): (i) the computational cost of coding and decoding the data is reduced, and (ii) errors are not propagated between the groups, thereby
25 increasing the robustness of the system.

Furthermore, it can significantly reduce the memory requirements of the encoder and decoder. By contrast, the current tone re-ordering system implemented in ADSL system allows for tone re-ordering to happen across all the tones during the seamless rate adaptation. To allow that it is required to have two copies of the constellation gains, bit allocation tables and tone reordering tables for all the tones. For an VDSL2 system with potential 4k tones, this requires significant memory on chip. By means of applying tone ordering only to the tones of groups, it is proposed to reduce the above requirements and provide provision to achieve the possible means to do dynamically adapt the data rates.

The groups in the first and second aspects of the invention may be defined in various ways within the scope of the invention, and need not be the same.

A first possibility is for the groups to correspond to the bands (i.e. that all the consecutive tones which transmit information in a given direction are in the same group). A second related possibility is for there to be multiple groups within each of the bands, but such that all the tones of any given group are all within a single one of the bands.

A third possibility is for the groups to be selected irrespective of the bands, e.g. such that each group of tones are consecutive ones of the tones for transmitting data in a given direction. For example, the groups of upstream (downstream) tones may be consecutive, with some of the groups extending over more than one of the upstream (downstream) bands.

In any of these three possibilities, the groups may be predefined, i.e. as part of the protocol. However, this feature is not required by the invention, which may instead allow the groups to be defined based on the SNR ratios.

For example, a fourth possibility which is particularly suitable for the first aspect of the invention, is for the groups to be defined based on the SNR

ratios, but for the number of tones in each group to be limited. The trellis encoding is performed only on the tones within each group. In one form of this possibility, for example, the tones of a given group may be selected from any of the tones associated with a given direction of communication (e.g. based
5 on conventional tone ordering) but each group is limited to a maximum number of tones, and the tones of this group are trellis encoded together. This too would help to ensure that errors are not propagated between the groups.

These various possibilities can be performed in combination. For example, a further way to select groups would be based on SNR ratios (as in the fourth
10 possibility) but such that the tones of any given group are within only a single one of the bands (as in the first possibility).

A particularly preferred example is for tone ordering to be performed during the learning phase for the tones which are members of groups which are entirely within a single band. Furthermore, the latency paths selected for each
15 tone are preferably determined based on the SNRs only of that tone and the other tones which are members of the same group. Subsequently, bit swapping is only performed for to re-allocate the bits between the tones of these same groups.

The third to seventh aspects of the invention relate to techniques which
20 reduce the computational burden of the IFFT.

The third aspect proposes in general terms that the sizes of the bands for the transmit and/or receive directions are selected, so as to give efficient IFFT.

By contrast, the fourth to sixth aspects relate specifically to modes of operation which require fewer tones than the full number provided by the
25 protocol. Such modes may be useful when the volume of data which is to be transmitted is reduced. The invention proposes that this is done in such a way as to reduce the IFFT/FFT computational burden, to give a "low power"

transmission mode (in contrast to the “high power” transmission mode in which all the tones are fully used). This also has the advantage that reduced memory is required, especially in the decoder.

5 The fourth aspect proposes in general terms that the number of tones in the low power mode is reduced, such that only a selected number of the tones are present in the signal input to the IFFT and FFT. The IFFT/FFT module (which normally employs an algorithm designed to be capable of processing data specifying the full range of tones available for transmission in the corresponding direction) is switched to a mode of operation which uses the
10 information that certain of the tones are not-present to reduce the computation it performs.

The fifth aspect of the invention similarly proposes that the number of tones is reduced (e.g. so as to be as in ADSL), and an IFFT is used which only processes low frequency tones. The number of frequency components is then
15 increased by up-sampling. This aspect makes it possible for a single line driver and DAC (digital-to-analogue converter) to be employed in both of the two modes.

An alternative, which is the sixth aspect of the invention, is to effect the low power mode by switching from VDSL communication to ADSL communication
20 (or, more generally, from a first band plan to a second band plan which employs a maximum frequency for data transmission which is lower than the maximum frequency used for data transmission in the first band plan). During the transitional period between the VDSL (or, more generally, first band plan) and ADSL (or more generally second band plan) and vice versa data
25 may perhaps be lost, but in certain implementations this cost may still be worthwhile for the power saving of ADSL.

Note that optionally the second band plan may include transmitting signals which do not encode data in frequencies above the maximum frequency

which it uses for data transmission. For example, signals which do not encode data (e.g. signals which are predefined) may be transmitted in the second band plan at all the frequencies above the maximum frequency which is used to transmit data in the second band plan and up to the maximum frequency which is used to transmit data in the first band plan. An advantage of transmitting such signals is that if this is not done the very fact that signals of these high frequencies are not being transmitted may mean that there are changes in the noise environment. Optionally, these frequencies may optionally all be transmitted with the same gain, or more generally with a gain which is not chosen by an analysis of the SNRs of the individual tones in these frequencies. This reduces the memory which is required to store gains for each of the tones.

Note furthermore that the first and second aspects of the invention are advantageously combinable with the sixth aspect of the invention, since in the case that the groups are defined (e.g. predefined) to include groups of low frequency and groups of relatively high frequency, the transition from one protocol to the other can be performed by simply ceasing to transmit data on a plurality of the high frequency groups of tones.

The invention, in its various aspects, may be expressed as a method of transmitting and/or receiving data (i.e. the method performed by the apparatus at any one end of the line, or the method performed by the system as a whole including the apparatus at both ends of the line). Alternatively, the invention may be expressed as the apparatus which performs those methods, and having means for implementing each of the method steps.

25 Brief Description of The Figures

Preferred features of the invention will now be described, for the sake of illustration only, with reference to the following figures in which:

Fig. 1 shows the frequency usage of a conventional ADSL technique;

Fig. 2, which is composed of Figs. 2(a) and 2(b), illustrates tone ordering in a conventional ADLS technique;

Fig. 3 shows the frequency usage of a conventional VDSL technique;

Fig. 4 shows a grouping of tones in a first embodiment of the invention;

5 Fig. 5 shows a butterfly of an IFFT/FFT stage of a third embodiment of the invention;

Fig. 6, which is composed of Fig. 6(a) to 6(d), shows stages of IFFT processing in the third embodiment of the invention;

10 Fig. 7, which is composed of Figs. 7(a) to 7(c), shows stage of FFT processing in the third embodiment of the invention;

Fig. 8 shows schematically the structure of a fourth embodiment of the invention;

Fig. 9, which is composed of Fig. 9(a) and 9(b) shows schematically the data transmission rate in a fifth embodiment of the invention;

15 Fig. 10 shows a known trellis encoder which may be used in an embodiment of the invention;

Fig. 11 shows a known convolutional encoder which may be used in the encoder of Fig. 10;

Fig. 12 is a known trellis diagram;

20 Fig. 13, which is composed of Figs. 13(a) to 13(d), shows a specific example of grouping the tones in an embodiment of the invention;

Fig. 14 shows a known band plan;

Fig. 15 shows the result of pre-processing the band plan of Fig. 14;

Fig. 16 is a radix-2 butterfly;

25 Fig. 17, which is composed of Fig. 17(a) to 17(d), shows a first band plan used in a second embodiment of the invention, and the steps of processing it; and

Fig. 18, which is composed of Fig. 18(a) to 18(d), shows a second band plan used in the second embodiment of the invention, and the steps of
30 processing it;

Fig. 19 shows the implementation of the low power mode in a fourth embodiment of the invention;

Fig. 20 shows the implementation of the low power mode in a fifth embodiment of the invention; and

5 Fig. 21, which is composed of Figs. 21(a) to 21(f), shows possible band plans produced by either the fourth or fifth embodiments of the invention.

Detailed Description of the embodiments

10 Protocols which are embodiments of the present invention will now be described in detail.

Referring to Fig. 4, a schematic view is shown of bandwidth allocation in a protocol which is a first embodiment of the invention. Many features of the
15 protocol which are not described are generally according to the ADSL standard (which is described for example in the document ITU-T Recommendation G.992.3 published by the International Telecommunication Union). In particular the protocol preferably includes any one or more of the following: (i) trellis encoding, (ii) tone ordering (for selection of transmission
20 paths), and (iii) bit swapping.

Generally, speaking in this first embodiment the upstream and downstream allocation of the bandwidth may be as in VDSL (i.e. as shown in Fig. 3), but within the scope of the invention any other allocation is possible also. Indeed
25 an embodiment of the invention having a novel allocation of the bandwidth into bands is described below.

The first embodiment has the feature that in at least one, and more preferably both, of the upstream and downstream directions the tones are grouped into
30 groups 10, 11 (in Fig. 4 the downstream groups are labelled 10 and the

upstream groups 11). In various versions of the first embodiment within the scope of the invention this grouping is done in various ways. A first possibility, for example, is to predetermine the number of tones in each group 10, 11 (e.g. such that all groups are the same number N of tones, such as 4 tones),
 5 and then to allocate tones to groups according to a simple scheme, for example such that the groups 10 are respective sets of N consecutive tones.

The groups are used for encoding the data, in one or both of the following ways:

- 10 1) Trellis encoding is used only to combine the N tones within the groups 10, 11. The decoding of such data is much cheaper computationally than if trellis encoding is performed encoding together data which is to be placed in all the tones (as in ADSL)
- 2) Tone ordering (e.g. in connection with selecting the number of bits to
 15 be transmitted per tone, selection of the gains and/or selection of latency paths) and/or bit swapping, are performed only in relation to the tones of single groups 10, 11. For example, within a certain group the SNR of each of the tones may be established during a training phase, and based on this SNR value a selection may be made of which of a plurality of paths each of the
 20 tones in the groups should be transmitted along. The plurality of paths preferably include at least one path with interleaving, and at least one path without interleaving. The two paths may, for example, be generated by a system as shown in Fig. 2(a). However, note that the invention is not limited in this respect, and in particular there may be more than two possible paths.

25

Both of these two features have the consequence that the computational cost of decoding the data is much reduced. For example, if the trellis encoder employs four registers in its algorithm, then decoding of the Trellis data can be performed in about 20 stages in a Viterbi decoder (that is, the number of
 30 registers multiplied by a parameter which from statistical theory is known to be

about 5). If the number of tones in each group is lower, then the computational task performed by each stage is less.

We now describe a first particular example of grouping in the first embodiment: grouping used for Trellis encoding. In this example, the groups are defined such that no group spans two bands (i.e. each group is entirely made up of tones which are all in a single band). This ensures that errors are not propagated from one band to another: decoding of such data limits the error propagation across all bands and provides the benefits of the DMT technology. At the same time the decoding is much cheaper computationally than if trellis encoding is performed encoding together data which is to be placed in all the tones (as in ADSL).

For example, each group may contain maximum of 256 tones. A 4-dimensional 16-state Wei's encoder widely used in ADSL2 is shown in Fig. 10 where bits are encoded from input data labelled as u to form encoded data shown as v and w . For example, bits for encoding on tone pairs u_1 and u_2 are input to a convolution encoder shown in Fig. 11. For the encoding, two tone pairs are taken. The known Trellis diagram associated with the encoder is shown in the Figure 12. Other types of encoder may also be used. Block processing of Wei's 16-state 4-dimensional trellis code is supported to improve system performance. An algorithmic constellation encoder is used to construct constellations with a maximum number of bits defined in the standard. Data bits from the data frame buffer are extracted according to the bit allocation table b_i' , least significant bit first. Because of the 4-dimensional nature of the code, the extraction is based on pairs of consecutive b_i' , rather than on individual ones. Furthermore, due to the constellation expansion associated with coding, the bit allocation table, b_i' , specifies the number of coded bits per sub-carrier, which can be any integer from 2 to 15. Trellis

coding is performed on pairs of consecutive b'_i values, $(x = b'_{2 \times i}, y = b'_{2 \times i + 1})$, in the order $i = 0$ to $(NSC/2) - 1$. The value b'_0 is prepended to the reordered bit table b'_i to make an integer number of pairs and is set to 0. Within each group in order to force the final state to the zero state $(0, 0, 0, 0)$, the 2 LSBs u_1 and u_2 of the final two 4-dimensional symbols within the group of the DMT symbol are constrained to $u_1 = S_1 \oplus S_3$, and $u_2 = S_2$.

The last two 4-dimensional symbols in the DMT symbol are chosen to force the convolutional encoder state to the zero state. For each of these symbols, the 2 LSBs of u are predetermined, and only $(x + y - 3)$ bits are extracted from the data frame buffer and allocated to $t_3, t_4, \dots, t_z$. The above requirements imply a minimum size of the b'_i table of 4 non-zero entries. The minimum number of non-zero entries in the corresponding b'_i table could be higher. Other types of the encoder could also be used.

The groups of the Trellis bands could be defined. By limiting the span of the Trellis, and forcing the reset after a certain number of tones, the algorithm is made less computationally expensive while the Viterbi decoding is performed. This provision provides several advantages for the decoding process:

- a) It helps to improve the overall performance of the system, reducing the error propagation only within certain groups of tones and frequency bands due to Trellis. Inherently it makes the system very much independent of the complete frequency bands. It enhances the capabilities of the DMT system making it more robust and provides better noise correction capabilities.
- b) It helps to reduce the dynamics of the Path matrices which is required to store while the decoding is performed. While performing the trace-back during the Viterbi decoding, depending on the number of states, the branch matrices are calculated and at each input the accumulated path

matrices are calculated for a given state. In the current ADSL system, depending on the frequency tones, the path metrics will be accumulated for all the tones (namely 256 tones for ADSL2 and 512 tones for ADSL2+). In VDSL system, where the number of tones are very high (approx for 4096 tones or
 5 higher), and considering 2.5k tones in one direction, the cumulative path matrices requires large bit-widths without providing any benefit.

c) It helps to reduce the total memory requirements for the Viterbi decoder state memory. During the Viterbi decoding process, depending on the constraint length the survivor path, XY decode as well as the cumulated
 10 path matrices is required to be stored for decoding process depending on the constraint length. Hence this proposal allows for

The figure 13 shows the downstream and upstream spectrum for VDSL1
 15 services according to one form of this embodiment of the invention. In Fig. 13(a), the two upstream bands are shown as US1 and US2, while the two downstream bands are shown as DS1 and DS2. As shown respectively, in Fig. 13(b) and Fig. 13(c), it is proposed that each of the downstream and upstream frequency bands be further grouped in several bands $GD_1, \dots, GD_n$
 20 and $GU_1, \dots, GU_m$, here n and m are respectively the number of bands for downstream and upstream transmission. These bands may be limited by maximum frequency tones within each group.

Alternatively, as shown in Fig. 13(d) (which shows, for the sake of example
 25 only, the downstream direction), the frequency tones within each of the groups may also span across complete frequency spectrum, which could be more efficient for trellis and could be also dependent on the latency paths or tone ordering. In Fig. 13(d) areas with the same shading level represent tones within the same group. Thus, for example, the tones of Fig. 13(d) labelled GD_1
 30 are the tones which together constitute the band GD_1 .

It is also possible that the frequency bands DS1 be covered with a single trellis group and DS2 with second trellis group and similarly for the upstream bands.

5

We now describe a second particular example of grouping in the first embodiment of the invention: grouping for the tone ordering (and subsequent dynamic configuration). This is performed using many of the techniques known already for tone ordering, and reference is made in particular to the
10 ADSL standard G992.3, and in particular Fig. 8-7.

Specifically, during initialization, the receive PMD function shall calculate the numbers of bits and the relative gains to be used for every subcarrier, as well as the order in which subcarriers are assigned bits (i.e., the tone ordering).

15 The calculated bits and gains and the tone ordering shall be sent back to the transmit PMD function during a later stage of initialization (see 8.5.3.3).

The pairs of bits and relative gains are defined, in ascending order of frequency or subcarrier index i , as a bit allocation table b and gain table g (i.e., b_i and g_i , for $i = 1$ to $NSC - 1$, with b_1 bits to be allocated to subcarrier 1
20 and $b_{NSC - 1}$ bits to be allocated to subcarrier $NSC - 1$). If trellis coding is used, the receive PMD function shall include an even number of 1-bit subcarriers in the bit allocation table b .

The tone ordering table t is defined as the sequence in which subcarriers are assigned bits from the input bitstream (i.e., t_i for $i = 1$ to $NSC - 1$, with
25 constellation mapping beginning on subcarrier t_1 and ending on subcarrier $t_{NSC - 1}$). The tone ordering table t shall remain static for the duration of the session.

Following receipt of the tables b , g and t , the transmit PMD function shall calculate a reordered bit table b' and a reordered tone table t' from the original
30 tables b and t . Constellation mapping shall occur in sequence according to the

re-ordered tone table t' , with the number of bits per tone as defined by the original bit table b . Trellis coding shall occur according to the re-ordered bit table b' .

If trellis coding is not used, $b' = b$ and $t' = t$.

- 5 If trellis coding is used, the reordering of table t shall be performed by the transmit PMD function. The reordered tone table t' shall be generated according to the following rules:

- Indices of all subcarriers supporting 0 bits or 2 or more bits appear first in t' , in the same order as in table t .
- 10 • Indices of all subcarriers supporting 1 bit appear last in table t' , in the same order as in table t .

If the bit allocation does not include any 1-bit subcarriers, the reordered tone table t' is identical to the original tone table t .

- 15 The (even number of) 1-bit subcarriers shall be paired to form 2-dimensional constellation points as input to the trellis encoder. The pairing shall be determined by the order in which the 1-bit subcarriers appear in the original tone ordering table t .

- The table b' is generated by scanning the reordered tone table t' and reordering the entries of table b according to the following rules (with
- 20 NCONEBIT representing the number of 1-bit subcarriers in the bit allocation table b):

- The first $\text{NCONEBIT}/2$ entries of b' shall be 0, where NCONEBIT is the (by definition, even) number of subcarriers supporting 1 bit.
- The next entries of b' shall be 0, corresponding to the subcarriers that
- 25 support 0 bits.
- The next entries of b' shall be non-zero, corresponding to the subcarriers that support 2 or more bits. The entries shall be determined using the new tone table t' in conjunction with the original bit table b .
- The last $\text{NCONEBIT}/2$ entries of b' correspond to the paired 1-bit
- 30 constellations (i.e., 2 bits per entry).

The table b' is compatible with the G.992.1 trellis encoder.

The tables b' and t' shall be calculated from the original tables b and t as shown in the tone pairing and bit re-ordering processes below.

```

5  /* TONE RE-ORDERING PROCESS */
   t_index=1; /* tone order index t_index is index of array t */
   t'_index=1; /* tone paired index t'_index is index of array t' */
   while (t_index<NSC) {
       tone=t[t_index++]; bits=b[tone];
10  if (bits==0) { t'[t'_index++]=tone; }
       if (bits==1) { }
       if (bits≥2) { t'[t'_index++]=tone; }
       }
   while (t'_index<NSC) t'[t'_index++]=1;
15
   /* BIT RE-ORDERING PROCESS */
   NC1=0; /* NCONEBIT is the number of tones with 1 bit */
   NCL=0; /* NCUSED is the number of used tones (at least 1 bit) */
   for (i=1; i<NSC; i++) { if (b[i]>0) NCL++; if (b[i]==1) NC1++; }
20  b'_index=1; while (b'_index<(NSC-(NCUSED-NCONEBIT/2))) b'[b'_index]=0;
   t'_index=1; while (t'_index<NSC) {
       tone=t'[t'_index++]; bits=b[tone];
       if (bits==0) { }
       if (bits==1) { b'[b'_index++]=2; t'_index++; }
25  if (bits≥2) { b'[b'_index++]=bits; }
       }

```

If on-line reconfiguration changes the number or indices of 0-bit subcarriers or 1-bit subcarriers, then tables t' and b' shall be recalculated from the updated table b and the original table t .

The constellation encoder takes L bits per symbol from the PMS-TC layer. If trellis coding is used, the L bits shall be encoded into a number of bits L' matching the bit allocation table b and the reordered bit table b' , i.e., into a number of bits equal to $L' = \sum b'_i = \sum b_i$. See 8.6.2 of the standard. The value of L and L' relate as:

$$L' = \sum b'_i = \sum b_i = L + \left\lceil \frac{NCUSED - \frac{NCONEBIT}{2}}{2} \right\rceil + 4$$

with the $\lceil x \rceil$ notation representing rounding to the higher integer. The above relationship shows that using the 1-bit subcarrier pairing method, on average, one trellis overhead bit is added per set of four 1-bit subcarriers, i.e., one trellis overhead bit per 4-dimensional constellation. In case trellis coding is not used, the value of L shall match the bit allocation table, i.e., $L = \sum b_i$.

A complementary procedure should be performed in the receive PMD function. It is not necessary, however, to send the re-ordered bit table b' and the re-ordered tone table t' to the receive PMD function because they are generated in a deterministic way from the bit allocation table and tone ordering tables originally generated in the receive PMD function, and therefore the receive PMD function has all the information necessary to perform the constellation demapping and trellis decoding (if used).

All these features are incorporated in this second example of the first embodiment. Assuming that the groups of the Trellis coding are as shown in Figs. 13(b) and 13(c), or alternatively in Fig. 13(d), the tone ordering may be performed using only the tones of the same groups used for Trellis encoding, will all the tones of the group transmitted using the same latency path.

For the on-line reconfiguration, for efficient implementation, the idea is to limit the tone ordering process within a trellis band for single latency path. The bit-

swap procedure should be limited to one trellis group at one time. After the procedure is completed for one trellis group.

Optionally, as noted above, the Trellis bands may be limited to a single
5 transmission band. This helps to smoothly transit to low power modes. During the low power modes (as described in detail below) the transmission could be restricted to only few bands. Hence if it is allowed to smoothly transit to one mode of transmission to another, it allows for smooth transition of the mode. In low power mode, the same set of tone ordering table, bit gain tables could
10 be used with the exception that the bands not used for transmission during the low power modes may not be used. The trellis pairs also will not be affected and hence allows for seamless transition to low power mode. This mode of operation provides efficient implementation.

15 We now turn to a second embodiment of the invention, having features which are freely combinable with those of the first embodiment of the invention. In general terms, the proposal is to define the band plans for the downstream and upstream directions so as to reduce the power of the FFT/IFFT module and have efficient computations. With the band plans appropriately defined,
20 only a small number of computations will be required and saving significant power. The idea here is to propose a scheme such that the FFT/IFFT only processes the non-trivial data, hence saving a large amount of power. The power saving can be more than 50% depending on the number of tones that are used to carry the relevant data. Consequently, it lowers the heat
25 dissipation requirements. This is particularly important for remote DSL equipment, where heat is a challenging problem.

If the downstream band is configured to tones which are power of 2, i.e. efficient implementation could be drawn and significant savings could be
30 achieved. Due to the architecture of the FFT/IFFT, the samples taken for each

butterfly computation in stage 0 involves at most one non-trivial sample. This significantly reduces the computation since the samples with value zero do not contribute to the computation and can be bypassed.

- 5 For butterflies of different radices there are potential savings. As an illustration, consider the radix-2 butterfly shown in Fig. 16. In the pre-processing stage, for IFFT each butterfly processing tone N takes the other sample from tone $(N-4096)$. With the proposed band plan 997 (shown in Fig. 14, where the two shaded areas (tones 1 to 696, and 1182 to 1634) represent
- 10 the two bands used for transmission in one direction), each butterfly at the pre-processing stage, has at most one non-trivial input. Thus no complex addition is needed in these butterflies, only one complex multiplication is needed. There could be many other butterflies where both inputs are trivial inputs and could be skipped. If the number of non-zero tones is 1k in the
- 15 downstream direction, then only 1k complex multiplications are required. Other 3k butterfly processing stages could be skipped. Referring to Fig. 14, no more butterfly computation is needed in this stage. The result of the pre-processing is shown in Fig. 15. In the next subsequent stages, if the band separation is the number of such reduced radix-2 butterflies doubles from one
- 20 stage to another as the stride taken halved from one stage to another.

- The following explanation is based on 4K tones and radix-2 butterfly. There are 8 stages of FFT/IFFT processing and a pre-processing or post-processing stage. In the VDSL band plans, tones 513 to 575 are used for all downstream
- 25 transmission profile (TS 101 270-2 V2.0.8). Assume that instead the band plan we adjusted as one possible illustration below whereby the tones configures in direction are aligned to get more non-trivial computations. As in Figure 17(a) for the data distribution in the memory. The radix-2 butterfly takes in 2 samples in stride of 2k for the first stage, for example $\{X_1, X_{2048}\}$.
- 30 Since only X_{2048} is zero, the two outputs of the radix-2 butterfly are $\{X_1, X_1W_1\}$

where W_1 , is the associated twiddle factors. From Figure 17(a) and (b), it is clear that all the radix-2 butterflies in the pre-processing phase and also at Stage 0 have only one non-zero value each. It is possible that the pre-processing and Stage 0 twiddle factors be combined together into a single stage. These twiddle factors are a subset of the set of twiddle factors needed for the 4K tones IFFT, so no pre-computation and additional storage are required. The outputs are again the product of the non-zero input and a twiddle factor. This process saves one complete stage and significant reduction in computational cycles as well as power consumption. Those butterflies with all inputs of zero value do not have to be processed at all. The radix-2 butterflies in the first stage have either 1 non-zero input or all inputs of value zero. After the first stage of radix-2 processing, some samples still have zero value if the frequency band allocation is done appropriately. If the frequency band groups are chosen with steps of two to power n , 2^n , where n is an integer value most computational savings could be achieved while IFFT/FFT processing. The pre-processing stage of the IFFT computation could be combined with the processing of 1st stage of the IFFT processing hence saving computational cycles, memory bandwidth and power consumption.

20

A second example of a suitable proposed band plan according to this embodiment is shown in Fig. 18(a). In this band plan too the two bands in a given direction (tones 63 to 575, and tones 1473 to 1984) each have 2^n tones, and are spaced apart by 2^m tones (note that n is not equal to m). Pre-processing leads to the usage of tones shown by shading in Fig. 18(b); stage 0 processing to the usage of tones shown by shading in Fig. 18(c); and stage 1 processing to the usage of tones shown by shading in Fig. 18(d).

Note that while it is preferred in this embodiment for each of the bands to have exactly 2^n tones, and be spaced apart by exactly 2^m tones (n not

30

necessarily equal to m), as in Fig. 17(a) and 18(a), exact conformity with this rule is not required, and lower (but still appreciable) savings compared to the band plan of Fig. 14 are possible with small departures from the rule, provided that the number of tones per band is substantially equal to a power of two, and the number of tones spacing bands apart is also substantially equal to a power of two. For example, reasonable savings may be achieved if the number of tones is say $2^n - k$, where k is preferably less than 10% of 2^n and more preferably less than 5% of 2^n . Similarly, the spacing of the bands is preferably a number of tones which may be written $2^m - j$, where j is preferably less than 10% of 2^m and more preferably less than 5% of 2^m .

We now turn to embodiments of the invention which are protocols permitting a low power mode of operation, principally saving power consumption in the IFFT in the case of the encoder, or the FFT in the case of the decoder.

In a third embodiment of the invention, this is achieved by providing the possibility of a low power mode in which data is only transmitted on certain of the frequencies. The other tones are all zeros. The IFFT algorithm may be modified to a form which incorporates knowledge of what frequencies are unused and which has a reduced computational requirement compared to the form of the IFFT used for the full range of tones. As discussed below, this can result in a power saving of over 50%, depending on the number of tones which are used to carry data in the low power mode. This is particularly important for remote DSL equipment, where heat is a challenging problem.

FFTs and IFFTs may be regarded as a plurality of stages, each involving a "butterfly" calculation. In the embodiment, the IFFT (or FFT in the case of the decoder) is designed (and the non-zero tones selected) such that in the first stage ("stage 0") the samples taken for each butterfly computation involve at

most one non-trivial sample, and all the trivial samples (involving coefficients which are all zero, or only 1 non-zero coefficient) are bypassed.

The following explanation, with reference to Fig. 5, is based on the use of 4k tones and a radix-16 butterfly, which is implemented using radix-4. There are three stages of FFT/IFFT processing, and also a pre-processing or post-processing stage. Suppose that the full set of tones are labelled 0 to 4095, that in the high-power mode tones 513 to 575 are used for downstream data flow (as in VDSL according to TS 101 270-2 V2.0.8), and that in the low power mode only these 63 tones are used and the remaining downstream tones are zero. This is illustrated in Fig. 6(a) where the non-zero tones are shaded.

After the pre-processing, tones 513 to 575 and 3521 to 3583 are non-zero. Fig. 6(b) shows the data distribution.

The radix-16 butterfly of Fig. 1 takes in 16 samples in strides of 256, for example $\{X_1, X_{256}, X_{513}, \dots, X_{3585}\}$. Each radix 4 butterfly in the first phase of the radix-16 butterfly takes the samples in strides of 1024; thus, samples X_{513} , X_{1537} , X_{2561} and X_{3585} are taken for butterfly processing. Since only X_{513} is non-zero, the four outputs of the radix-4 butterfly are $\{X_{513}, X_{513}W_1, X_{513}W_2, X_{513}W_3\}$ where W_1 , W_2 and W_3 are the associated twiddle factors.

From Fig. 5 it is clear that the four radix-4 butterflies in the second phase have only one non-zero value each. The outputs are again the product of the non-zero inputs and twiddle factors. Hence, the radix-16 butterfly outputs are $\{X_{513}, X_{513}T_1, X_{513}T_{10}, X_{513}T_{11}, X_{513}T_{12}, X_{513}T_{13}, X_{513}T_{14}, X_{513}T_{15}\}$, where T_i ($i=1, \dots, 15$) are the combined twiddle factors. These twiddle factors are a subset of the set of twiddle factors needed for the 4k tones of the FFT/IFFT, so no pre-computation or additional storage are required. The radix-16 butterflies in the

first stage have either 1 non-zero input or all inputs of value zero. Those butterflies with inputs of zero value do not have to be processed at all. After the first stage of radix-16 processing, the coefficients are as shown in Fig. 6(c). Some samples still have zero value.

5

The second stage of the IFFT processing takes data in strides of 16, for example {Y1, Y17, Y33, Y49, Y65, Y81, Y97, Y113, Y129, Y145, Y161, Y177, Y193, Y209, Y225, Y241}. The first 4 samples and the last 4 samples are non-zero. Therefore, for each radix-4 butterfly in the first segment, two
 10 samples are non-zero and two are zeros. The number of complex additions is halved. There would be no saving in the second segment since the inputs to the 4 radix-4 butterflies are all non-zero.

There is no potential saving in the third stage. Table 1 shows the breakdown
 15 of the computation saving in each stage as a factor of the number of non-trivial tones. For the number of tones from 48 to 63, the pre-processing stage, stage 0 and stage 1 have a computation reduction of approximately 97%, 80% and 10% respectively (1 CMAC $\approx$ 5 CADD). Overall saving is about 40%. For a number of tones less than 16, even stage 1 can be reduced by up to
 20 40%.

There is no saving in the IFFT processing, if the number of tones exceeds 127. Hence, it is preferable to limit the number of tones to this number, and specifically to the range from 513 to 639.

25

Table 1 Computation Complexity Table for IFFT

Number of non-zero tones	Pre- processing	Stage 0	Stage 1	Stage 2

0-15	0-15BF	0-30 BF16'	256*(2 BF4 ₁ +4 BF4 ₂) (1 trivial BF4 ₁)	256 BF16
16-31	16-31 BF2	32-62 BF16'	256*(4 BF4 ₁ +4 BF4) (1 trivial BF4 ₁)	256BF16
32-47	32-47 BF2	64-94 BF16'	256*(2 BF4 ₁ +2 BF4 ₂ +4 BF4) (1 trivial BF4 ₁)	256 BF16
48-63	48-63 BF2	96-126 BF16'	256*(4 BF4 ₂ +4 BF4)	256 BF16
64-79	64-79 BF2	128-158 BF16'	256*(2 BF4 + 2 BF4 ₂ + 4 BF4)	256 BF16
80-95	80-95 BF2	160-190 BF16'	256 BF16	256 BF16
96-127	96-127 BF2	192-254 BF16'	256 BF16	256 BF16
128-143	128-143 BF2	256 BF16	256 BF16	256 BF16
4K	2K BF2	256 BF16	256 BF16	256 BF16
		BF16'=15 CMAC, 0 CADD	BF4 ₁ =3 CMAC, 0 CADD (room for optimisation)	

			BF ₄ ₂ =4 CADD, 3 CMAC	
			BF ₄ =8 CADD, 3 CMAC	

Here the term "BF₄" means a radix 4 butterfly, BF₂ mean a radix-2 butterfly, etc.. CADD stands for a complex adder operation, and CMAC for a complex multiplier-accumulator.

5

We now analyse the power saving for the FFT processing. Assume that in low-power mode, the 64 tones from 960 to 1023 are non-trivial, as shown by shading in Figure 7(a). These tones fit into the upstream spectrum for all transmission profiles. In stage 0, there are 64 radix-16 butterflies with a single non-zero input. These butterflies, similar to those discussed in the IFFT section, do not need any complex addition, only 15 complex multiplications are required. The non-zero values output from the first stage are shown in Fig. 7(b). The remaining butterflies in this stage are trivial. In stage 1, all 256 radix-16 butterflies have to be processed, but with substantial saving in the computation. Each of the 4 radix-4 butterflies in the first segment has only a single non-zero input. Hence, they do not need any addition, only 3 multiplications are required. There is no potential saving in the next segment since all the samples are non-trivial. There are no saving in stage 2 and the post-processing. Computation saving in stage 0 and stage 1 are 84% and 18% respectively. Overall saving is 35%. The tones recommended for the low-power mode is from 960 to 1023. Only 64 tones are proposed in this embodiment since the upstream has lower traffic.

20

Number of	Stage 0	Stage 1	Stage 2	Post-
-----------	---------	---------	---------	-------

non-zero tones				processing
0-16	0-16 BF16'	$256*(BF4_1+4 BF4_1)$	256 BF16	2K BF2
17-32	17-3216 BF16'	$256*(2 BF4_1+4 BF4_2)$	256 BF16	2K BF2
33-48	33-4816 BF16'	$256*(3 BF4_1+4 BF4_3)$	256 BF16	2K BF2
49-64	49-6416 BF16'	$256*(4 BF4_1+4 BF4_4)$	256 BF16	2K BF2
65-80	65-8016 BF16'	$256*(3 BF4_1 + BF4_2 + 4 BF4_4)$	256 BF16	2K BF2
81-96	81-9616 BF16'	$256*(2 BF4_1 + 2 BF4_2 + 4 BF4_4)$	256 BF16	2K BF2
97-112	97-11216 BF16'	$256*(BF4_1 + 3 BF4_2 + 4 BF4_4)$	256 BF16	2K BF2
113-128	113-12816 BF16'	$256*(4 BF4_2 + 4 BF4_4)$	256 BF16	2K BF2
129-144	129-14416 BF16'	$256*(BF4_3 + 3 BF4_2 + 4 BF4_4)$	256 BF16	2K BF2
145-160	145-16016 BF16'	$256*(2 BF4_3 + 2 BF4_2 + 4 BF4_4)$	256 BF16	2K BF2
161-176	161-17616 BF16'	$256*(3 BF4_3 + BF4_2 + 4 BF4_4)$	256 BF16	2K BF2

		BF4)		
177-192	177-19216 BF16'	256*(4 BF4 ₃ + 4 BF4)	256 BF16	2K BF2
193-208	193-20816 BF16'	256*(3 BF4 ₃ + 5 BF4)	256 BF16	2K BF2
209-224	209-22416 BF16'	256*(2 BF4 ₃ + 6 BF4)	256 BF16	2K BF2
225-240	225-24016 BF16'	256*(1 BF4 ₃ + 7 BF4)	256 BF16	2K BF2
241-246	241-24616 BF16'	256 BF16	256 BF16	2K BF2
4k	256BF16	256 BF16	256 BF16	2K BF2
	BF16'=15 CMAC, 0 CADD	BF4 ₁ =0 CADD, 3 CMAC		
		BF4 ₂ =4 CADD, 3 CMAC		
		BF4 ₃ =6 CADD, 3 CMAC		
		BF4=8 CADD, 3 CMAC		

Table 2 Computation Complexity Table for FFT

A fourth embodiment of the method is illustrated in Fig. 8. The upper part of
5 the diagram illustrates a known structure for encoding and transmitting data in
VDSL, and which performs the "high power" mode of operation of the

embodiment. The output of the IFFT unit 20 is transmitted to a filter 21 which produces data having a frequency of 17 or 35 MHz depending upon the sampling rate. The output of the filter 22 is passed to a DAC 23, from there to a line driver 24, and from there to a transformer 25 which increases the
5 voltage for transmission onto the line.

In the fourth embodiment, the low power mode of operation is performed by a second signal path which bypasses the IFFT unit 20 (thus considerably saving power) and instead employing an IFFT unit 30 such as used conventionally in
10 an ASDL system. That is one which is able to receive fewer tones than the unit 20, preferably only about 16th of the number of tones. This maybe 128 or 256 tones, for example. The output of the IFFT 21 is transmitted to an upsampler 31 which upsamples the frequency, e.g. by a factor of 16. The output of the upsampler 31 is transmitted to the filter 21. The power
15 consumption of the path of the lower part of Fig. 8 is considerably less than that of the IFFT unit 10, and thus there is a considerable power saving, but this does not disrupt the operation of the units 21 to 25 which receive (as far as they are concerned) equivalent input signals.

20 A further example of the structure of a device according to this embodiment of the invention is shown in Fig. 19. This structure is explained below, following the explanation of Fig. 20.

In a fifth embodiment of the invention, the protocol simply converts from a
25 high data transmission rate mode (such as VDSL) referred to here as L1, to a low data rate transmission mode (such as ADSL) referred to here as L2. This is illustrated in Fig. 9(a). At the transmission points some data may be lost (since the device receiving the encoded data will in general not be aware that the protocol is being changed), so that retransmission of this data may be

required. However, the cost of the retransmission may still be low in comparison to the power saving.

The fifth embodiment of the invention suggests switching the transmission the mode of transmission from VDSL to ADSL mode of transmission. Hence
 5 instead of power consumption of 2 Watts during the full traffic mode of VDSL, during the low power mode the power consumption will be much less than 1 Watt. Also it is suggested that the during the low traffic modes, the bands of transmission be limited to say 2k tones(8Mhz spectrum) or 1k tones(4Mhz
 10 spectrum) likewise. The data in the remaining tones is all zeros. The FFT/IFFT does not have to process the samples with value zero, and some stages in the processing can be skipped or reduced significantly.

15 A possible structure of the fifth embodiment of the invention is shown in Fig. 19. The command message to switch to low power mode is communicated by the ATU-C (**ADSL termination unit – central office**). Once the ATU-R (**ADSL termination unit – remote**) receives the message and may either accept the command or reject the same. If ATU-R in its acknowledgement
 20 sends a accept message, after a fixed duration coupled either with sync symbol, the ATU-C may change the state from VDSL to low band mode either of M0, M1, M2, M3, M4 or M5 mode of transmission. It will be a coordinated transition based on the command protocol, in which all of the IFFT, filter, DAC and line driver change their modes of operation. Since the ATU-C knows the
 25 end to end delay, it may either synchronize the transition to new state after a fixed duration of transmitted frame or could also use the synchronization frame as a reference. The ATU-C may also switch to different mode of transmission after a fixed synchronization frame delay. A command message is appropriately defined to provide with the information for the mode of
 30 transmission which could consist of mode of transmission(the bands used, or

limiting the FFT and frequency spectrum, the delay after which the switch could happen if required. The frames with a fixed reference could be used to specify the delay.

- 5 Hence, the advantage is a reduction of the power usage during low-power mode. The power saving can be more than 70% depending on the mode that are used to carry the relevant data. Consequently, it lowers the heat dissipation requirements. This is particularly important for remote DSL equipment, where heat is a challenging problem.

10

As mentioned above, Fig. 19 shows a example of a possible structure of the fourth embodiment of the invention. In contrast to Fig. 20, in Figure 19 it is suggested to use the transmission limited to 256 tones but to maintain the data rate on the interface. To ensure the same a simple interpolator (16x)
 15 could be used. The interpolator could assist to bring the frequency back to desired 17Mh spectrum. Other kind of circuitry could also be used. The interpolator could be also located after the filter or before the filter. Hence significant power savings could be achieved for the frequency domain processing. The above shows the scenario for transmit but the same is also
 20 valid for receive direction.

Fig. 21 shows band plans which can be produced in the various modes of operation of either the fourth or fifth embodiments of the invention. Fig. 21(a) shows the full frequency range (up to 17MHz) being used for transmitting data
 25 (i.e. the VDSL high power mode).

Fig. 21(b) shows a different mode of operation, in which only frequencies up to 1.1MHz are used for transmitting data. Those frequencies above 1.1MHz are shown with low values to indicate that no signal is transmitted on those
 30 frequencies, or (in other forms of the embodiments) that signals are

transmitted which do not carry data. For example, it may be advantageous to broadcast signals which do not carry data on frequencies 1.1MHz to 17MHz, to avoid changing the noise environment.

- 5 In yet further versions, the maximum transmission rate may be lower for the range of frequencies which are shown as having a low value. In other words, the range of frequencies (a "high frequency range") between 1.1MHz and 17MHz may be used to carry a certain data load in the first mode of operation, and a lower amount of data in the mode of Fig. 21(b).

10

Fig. 21(c) shows a third mode of operation in which frequencies up to 12MHz are used for transmitting data. Just as described above in relation to Fig. 21(a), the frequencies above 12MHz which are shown as "low" may be unused, used for non-data transmission, or used for data transmission at a lower data transmission rate than in Fig. 21(a).

15

Fig. 21(d) shows a fourth mode of operation in which frequencies up to 8MHz are used for transmitting data. Just as described above in relation to Fig. 21(a), the frequencies above 8MHz which are shown as "low" may be unused, used for non-data transmission, or used for data transmission at a lower data transmission rate than in Fig. 21(a).

20

Fig. 21(e) shows a fifth mode of operation in which frequencies up to 5.3MHz are used for transmitting data. Just as described above in relation to Fig. 21(a), the frequencies above 5.3MHz which are shown as "low" may be unused, used for non-data transmission, or used for data transmission at a lower data transmission rate than in Fig. 21(a).

25

Fig. 21(f) shows a sixth mode of operation in which frequencies up to 4.4MHz are used for transmitting data. Just as described above in relation to Fig.

30

21(a), the frequencies above 4.4MHz which are shown as "low" may be unused, used for non-data transmission, or used for data transmission at a lower data transmission rate than in Fig. 21(a).

- 5 Note that the various power saving features of the third to fifth embodiments may be freely combined, such that there is a choice of different power saving modes. This is illustrated in Fig. 9(b) in which the protocol uses a first power saving mode L_2 at a first time, and a second power saving mode L_4 (such as the one of the third or fourth embodiment of the invention) at other times.

10

Although only a few embodiments of the invention have been disclosed in this application, many variations are possible within the scope of the invention as will be clear to a skilled reader. For example, although all the features in the embodiments described above can be combined, the various aspects of the
15 invention can also be used separately. For example, the aspects of the invention relating to grouping of tones may be used when the modulation technique is other than one using an IFFT.

Appendix 1

1. What technical problem(s) is (are) to be solved by your invention?

The current tone re-ordering system implemented in ADSL system allows for tone re-ordering to happen across all the tones during the seamless rate adaptation. To allow that it is required to have two copies of the constellation gains, bit allocation tables and tone reordering tables for all the tones. For VDSL2 system with potential 4k tones, this requires significant memory on chip. Here it is proposed to reduce the above requirements and provide provision to achieve the possible means to do dynamically adapt the data rates.

This requires Trellis modulation has been incorporated to improve the performance of the DSL line system. In VDSL-DMT with multiple bands, if traditional trellis encoding is done to span over the complete band will be efficient since error will propagate from one band to another with traditional trellis implementation. Error correction capability of the Viterbi decoder hence is reduced.

2. How has the problem been solved to data? What are the disadvantages of the prior art?

In past as in ADSL system the tone ordering is . In VDSL1 system, the tone ordering is. New system development is on the way. Uptill now two sets of memories are required on chip. With VDSL2 system with 4k tones or higher, each of the bit allocation table, gain tables, tone ordering tables each for 4k tones requires significant on-chip memory.

During initialization, the receive PMD function shall calculate the numbers of bits and the relative gains to be used for every subcarrier, as well as the order in which subcarriers are assigned bits (i.e., the tone ordering). The calculated bits and gains and the tone ordering shall be sent back to the transmit PMD function during a later stage of initialization (see 8.5.3.3).

The pairs of bits and relative gains are defined, in ascending order of frequency or subcarrier index i , as a bit allocation table b and gain table g (i.e., b_i and g_i , for $i = 1$ to $NSC - 1$, with b_1 bits to be allocated to subcarrier 1 and b_{NSC-1} bits to be allocated to subcarrier $NSC - 1$). If trellis coding is used, the receive PMD function shall include an even number of 1-bit subcarriers in the bit allocation table b .

The tone ordering table t is defined as the sequence in which subcarriers are assigned bits from the input bitstream (i.e., t_i for $i = 1$ to $NSC - 1$, with constellation mapping beginning on subcarrier t_1 and ending on subcarrier t_{NSC-1}). The tone ordering table t shall remain static for the duration of the session.

Following receipt of the tables b , g and t , the transmit PMD function shall calculate a reordered bit table b' and a reordered tone table t' from the original tables b and t . Constellation mapping shall occur in sequence according to the re-ordered tone table t' , with the number of bits per tone as defined by the original bit table b . Trellis coding shall occur according to the re-ordered bit table b' . If trellis coding is not used, $b' = b$ and $t' = t$.

If trellis coding is used, the reordering of table t shall be performed by the transmit PMD function.

The reordered tone table t' shall be generated according to the following rules:

- Indices of all subcarriers supporting 0 bits or 2 or more bits appear first in t' , in the same order as in table t .
- Indices of all subcarriers supporting 1 bit appear last in table t' , in the same order as in table t .

If the bit allocation does not include any 1-bit subcarriers, the reordered tone table t' is identical to the original tone table t .

The (even number of) 1-bit subcarriers shall be paired to form 2-dimensional constellation points as input to the trellis encoder. The pairing shall be determined by the order in which the 1-bit subcarriers appear in the original tone ordering table t .

The table b' is generated by scanning the reordered tone table t' and reordering the entries of table b according to the following rules (with $NCONEBIT$ representing the number of 1-bit subcarriers in the bit allocation table b):

- The first $NCONEBIT/2$ entries of b' shall be 0, where $NCONEBIT$ is the (by definition, even) number of subcarriers supporting 1 bit.
- The next entries of b' shall be 0, corresponding to the subcarriers that support 0 bits.
- The next entries of b' shall be non-zero, corresponding to the subcarriers that support 2 or more bits. The entries shall be determined using the new tone table t' in conjunction with the original bit table b .
- The last $NCONEBIT/2$ entries of b' correspond to the paired 1-bit constellations (i.e., 2 bits per entry).

The table b' is compatible with the G.992.1 trellis encoder.

The tables b' and t' shall be calculated from the original tables b and t as shown in the tone pairing and bit re-ordering processes below.

```
/* TONE RE-ORDERING PROCESS */
t_index=1; /* tone order index t_index is index of array t */
t'_index=1; /* tone paired index t'_index is index of array t' */
while (t_index<NSC) {
```

```

tone=t[t_index++]; bits=b[tone];
if (bits==0) { t'[t'_index++]=tone; }
if (bits==1) { }
if (bits≥2) { t'[t'_index++]=tone; }
}
while (t'_index<NSC) t'[t'_index++]=1;

/* BIT RE-ORDERING PROCESS */
NCL=0; /* NCONEBIT is the number of tones with 1 bit */
NCL=0; /* NCUSED is the number of used tones (at least 1 bit) */
for (i=1; i<NSC; i++) { if (b[i]>0) NCL++; if (b[i]==1) NCL++; }
b'_index=1; while (b'_index<(NSC-(NCUSED-NCONEBIT/2)))
b'[b'_index]=0;
t'_index=1; while (t'_index<NSC) {
tone=t[t'_index++]; bits=b[tone];
if (bits==0) { }
if (bits==1) { b'[b'_index++]=2; t'_index++; }
if (bits≥2) { b'[b'_index++]=bits; }
}

```

Figure 8-7 presents an example to illustrate the tone reordering and bit reordering procedures, and the pairing of 1-bit subcarriers for trellis encoding.

Tone ordering table t (as determined by the receive PMD function, NSC=24)

7	14	21	4	11	18	1	8	15	22	5	12	19	2	9	16	23	6	13	20	3	10	17
---	----	----	---	----	----	---	---	----	----	---	----	----	---	---	----	----	---	----	----	---	----	----

Bit ordering table b (as determined by the receive PMD function, 37 bit/symbol)

0	1	2	3	2	1	2	1	0	2	0	2	1	1	3	3	3	2	1	0	2	3	2
---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---


Tone reordered table t' (moving 1-bit tones to the end of the table)

7	21	4	11	18	1	15	22	5	12	9	16	23	20	3	10	17	14	8	19	2	6	13
---	----	---	----	----	---	----	----	---	----	---	----	----	----	---	----	----	----	---	----	---	---	----

Bit reordered table b' (moving 0-bit tones to begin of the table)

0	0	0	0	0	0	0	2	2	3	2	3	3	2	2	3	2	2	2	3	1+1	1+1	1+1
---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	-----	-----	-----

Trellis pairs (encoding 25 data bits into 37 trellis bits) and bit mapping to tones



G 992.3_F08-7

Figure 8-7/G.992.3 – Example of frequency ordering and pairing of one-bit carriers

If on-line reconfiguration changes the number or indices of 0-bit subcarriers or 1-bit subcarriers, then tables t' and b' shall be recalculated from the updated table b and the original table t .

The constellation encoder takes L bits per symbol from the PMS-TC layer. If trellis coding is used, the L bits shall be encoded into a number of bits L' matching the bit allocation table b and the reordered bit table b' , i.e., into a number of bits equal to $L' = \sum b'_i = \sum b_i$. See 8.6.2. The value of L and L' relate as:

$$L' = \sum b'_i = \sum b_i = L + \left\lceil \frac{NCUSED - \frac{NCONEBIT}{2}}{2} \right\rceil + 4$$

with the $\lceil x \rceil$ notation representing rounding to the higher integer. The above relationship shows that using the 1-bit subcarrier pairing method, on average, one trellis overhead bit is added per set of four 1-bit subcarriers, i.e., one trellis overhead bit per 4-dimensional constellation. In case trellis coding is not used, the value of L shall match the bit allocation table, i.e., $L = \sum b_i$.

A complementary procedure should be performed in the receive PMD function. It is not necessary, however, to send the re-ordered bit table b' and the re-ordered tone table t' to the receive PMD function because they are generated in a deterministic way from the bit allocation table and tone ordering tables originally generated in the receive PMD function, and therefore the receive PMD function has all the information necessary to perform the constellation demapping and trellis decoding (if used).




Figure 4: VDSL Downstream and Upstream frequency bands and Bands of Trellis coding

For the on-line reconfiguration, for efficient implementation, the idea behind is to limit the tone ordering process within a trellis band for single latency path. The bit-swap procedure should be limited to one trellis group at one time. After the procedure is completed for one trellis group. In one embodiment it is also suggested that the trellis bands be limited to a single transmission band. This helps to smoothly transit to low power modes. During the low power modes the transmission could be restricted to only few bands. Hence if it is allowed to smoothly transit to one mode of transmission to another, it allows for smooth transition of the mode. In low power mode, the same set of tone ordering table, bit gain tables could be used with the exception that the bands not used for transmission during the low power modes may not be used. The trellis pairs also will not be affected and hence allows for seamless transition to low power mode. This mode of operation provides efficient implementation.

Appendix 2

1. Here a new low power mode is defined for VDSL2 system. Since the DSL system is always on system. During non-usage of the traffic, the system is always transmitting over the complete frequency spectrum and hence significant amount of power is wasted. For DMT system, FFT is used for modulation of the carriers with information to be transmitted over the line. During the low power mode it is required only to transmit very limited information related to management and status. The system could be generally capable of transmitting 100 Mbps of data. In low power modes, the data to be transmitted could be fairly low in the range of 100 kbps to 1 Mbps. In this case, transmission over the complete spectrum 4k tones may not be required. Also it could be possible that based on the usage or data traffic the system could dynamically limit the tones to be transmitted and hence reduce the power consumption.
2. Current system lowers the energy by power spectral density mask shaping to be transmitted. Thereby reducing the power to be transmitted over the line. Though in the current DMT system, it is still required to do the modulation over the complete 4k tones. Hence in the digital physical layer device, there is significant power consumed by the digital modulating device.
3. Power consumption is a major issue which limits the power feeding, limits also the line card density. Here with the proposed solution, it is suggested to reduce the power consumption from 2 Watts to less than 1 Watt.
4. The invention suggests the transmission to switch the mode of transmission from VDSL to ADSL mode of transmission. Hence instead of power consumption of 2 Watts during the full traffic mode of VDSL, during the low power mode the power consumption will be much less than 1 Watt. Also it is suggested that the during the low traffic modes, the bands of transmission be limited to say 2k tones(8Mhz spectrum) or 1k tones(4Mhz spectrum) likewise. The data in the remaining tones is all zeros. The FFT/IFFT does not have to process the samples with value zero, and some stages in the processing can be skipped or reduced significantly.




Figure 1: Shows one band plan for VDSL2-DMT with 17 Mhz spectrum.




Figure 2: Transmission path for DMT system

The idea is that after the transmission could switch to low power mode. The command message to switch to low power mode be communicated by the ATU-C. Once the ATU-R receives the message and may either accept the command or reject the same. If ATU-R in its acknowledgement sends a accept message, after a fixed duration coupled either with sync symbol, the ATU-C may change the state from VDSL to low band mode either of M0, M1, M2, M3, M4 or M5 mode of transmission. It will be a coordinated transition based on the command protocol. Since the ATU-C knows the end to end delay, it may either synchronize the transition to new state after a fixed duration of transmitted frame or could also use the synchronization frame as a reference. The ATU-C may also switch to different mode of transmission after a fixed synchronization frame delay. A command message is appropriately defined to provide with the information for the mode of transmission which could consist of mode of transmission(the bands used, or limiting the FFT and frequency spectrum, the delay after which the switch could happen if required. The frames with a fixed reference could be used to specify the delay.

Hence, the advantage is a reduction of the power usage during low-power mode. The power saving can be more than 70% depending on the mode that are used to carry the relevant data. Consequently, it lowers the heat dissipation requirements. This is particularly important for remote DSL equipment, where heat is a challenging problem.




Figure 3: Implementation for low power mode with VDSL data rate.

Here in the Figure 3 it is suggested to use the transmission limited to 256 tones but to maintain the data rate on the interface. To ensure the same a simple interpolator (16x) could be used. The interpolator could assist to bring the frequency back to desired 17Mhz spectrum. Other kind of circuitry could also be used. The interpolator could be also located after the filter or before the filter. Hence significant power savings could be achieved for the frequency domain processing. The above shows the scenario for transmit but the same is also valid for receive direction.

Appendix 3

1. What technical problem(s) is (are) to be solved by your invention?

The Trellis modulation has been incorporated to improve the performance of the DSL line system. In VDSL-DMT with multiple bands, if traditional trellis encoding is done to span over the complete band will be efficient since error will propagate from one band to another with traditional trellis implementation. Error correction capability of the Viterbi decoder hence is reduced.

2. How has the problem been solved to date?
What are the disadvantages of the prior art?

The problem is new and has not been addressed in past. New system development is on the way.

3. How does your invention solve this problem and what advantages does this solution offer?
(please give a detailed description)

The idea to solve the above problem is to limit the trellis state not to span across the bands. The trellis state machine should cover only single band. And restart the same while processing the next band of Downstream or upstream. Hence it ensures the error propagation is not allowed across several bands hence achieving higher performance.

4. What is the core of your invention? (i.e. the basic concept of your invention)

5. Please provides precise embodiments and/or implementation examples with emphasis on one preferred embodiment. (Full description, preferably accompanied with drawings)

The first embodiment has the feature that in at least one, and more preferably both, of the upstream and downstream directions the tones are grouped into groups 10, 11 (in Fig. 4 the downstream groups are labelled 10 and the upstream groups 11). In various versions of the first embodiment within the scope of the invention this grouping is done in various ways. A first possibility, for example, is to predetermine the number of tones in each group 10, 11 (e.g. such that all groups are the same number N of tones, such as 4 tones), and then to allocate tones to groups according to a simple scheme, for example such that the groups 10 are respective sets of N consecutive tones.

The groups are used for encoding the data, in one or both of the following ways:

Trellis encoding is used only to combine the N tones within the groups 10, 11. The decoding of such data limits the error propagation across all bands and provides the benefits of the DMT technology. At the same time the decoding is much cheaper computationally than if trellis encoding is performed encoding together data which is to be placed in all the tones (as in ADSL). The idea behind is that the Trellis code spans only across the a certain group of tones. So the complete tones and frequency bands be divided into certain number of tones, for ex. each group contains maximum of 256 tones. 4-dimensional 16-state Wei's encoder widely used in ADSL2 is shown in the Figure1 where by the conversion of bits of tone pairs are encoded to $(v$ and $w)$ computed from u . The convolution encoder is shown in the Figure 2. For the encoding, two tone pairs are taken. The Trellis diagram associated with the encoder is shown in the Figure3. Other types of encoder may also be used. Block processing of Wei's 16-state 4-dimensional trellis code shall be supported to improve system performance. An algorithmic constellation encoder shall be used to construct constellations with a maximum number of bits defined in the standard. Data bits from the data frame buffer shall be extracted according to the bit allocation table $\underline{b'_i}$, least significant bit first.

Because of the 4-dimensional nature of the code, the extraction is based on pairs of consecutive $\underline{b'_i}$, rather than on individual ones. Furthermore, due to the constellation expansion associated with coding, the bit allocation table, $\underline{b'_i}$, specifies the number of coded bits per subcarrier, which can be any integer from 2 to 15. Trellis coding shall be performed on pairs of consecutive b' values, $(x = \underline{b'_{2 \times i}}, y = \underline{b'_{2 \times i + 1}})$, in the order $i = 0$ to $(NSC/2) - 1$. The value $\underline{b'_0}$ is prepended to the reordered bit table b' to make an integer number of pairs and shall be set to 0. Within each group in order to force the final state to the zero state $(0, 0, 0, 0)$, the 2 LSBs u_1 and u_2 of the final two 4-dimensional symbols within the group of the DMT symbol are constrained to $u_1 = S_1 \oplus S_3$, and $u_2 = S_2$.

The last two 4-dimensional symbols in the DMT symbol shall be chosen to force the convolutional encoder state to the zero state. For each of these symbols, the 2 LSBs of u are predetermined, and only $(x + y - 3)$ bits shall be extracted from the data frame buffer and shall be allocated to $t_3, t_4, \dots, t_z$. The above requirements imply a minimum size of the $\underline{b'_i}$ table of 4 non-zero entries.

The minimum number of non-zero entries in the corresponding b_i table could be higher. Other types of the Encoder could also be used.

The groups of the Trellis bands could be defined. By limiting the span of the Trellis, and forcing the reset after a certain number of tones, it makes it computational less expensive while the Viterbi decoding is performed. This provision provides several advantages for the decoding process:

- a) The invention helps to improve the overall performance of the system, reducing the error propagation only within certain groups of tones and frequency bands due to Trellis. Inherently it makes the system very much independent of the complete frequency bands. It enhances the capabilities of the DMT system making it more robust and provides better noise correction capabilities.
- b) It helps to reduce the dynamics of the Path metrics which is required to store while the decoding is performed. While performing the traceback during the Viterbi decoding, depending on the number of states, the branch metrics are calculated and at each input the accumulated path metrics are calculated for a given state. In the current ADSL system, depending on the frequency tones, the path metrics will be accumulated for all the tones (namely 256 tones for ADSL2 and 512 tones for ADSL2+). In VDSL system, where the number of tones are very high (approx for 4096 tones or higher), and considering 2.5k tones in one direction, the cumulative path metrics requires large bit-widths without providing any benefit.
- c) The proposal helps to reduce the total memory requirements for the Viterbi decoder state memory. During the viterbi decoding process, depending on the constraint length the survivor path, XY decode as well as the cumulated path metrics is required to be stored for decoding process depending on the constraint length. Hence this proposal allows for

The figure 4 shows the proposed downstream and upstream spectrum for VDSL1 services. There could be multiple of downstream and upstream frequency bands. It is proposed that each of the downstream and upstream frequency bands be further grouped in several bands GD1, .. GDn And GU1,... GU_m. Where n and m are number of bands for downstream and upstream transmission. These bands maybe limited by maximum frequency tones within each group. The frequency tones within the group GD1 may also span across complete frequency spectrum which could be more efficient for trellis and could be also dependent on the latency paths or tone ordering, the same is represented in the last presentation of the Figure4. It is also possible that the frequency bands DS1 be covered with a single trellis group and DS2 with second trellis group and similarly for the upstream bands.




Figure 1. Trellis Encoder for conversion of bits to the tone pairs.




Figure2: Convolutional encoder




Figure32: Trellis diagram.

48




Figure 4: VDSL Downstream and Upstream frequency bands and Bands of Trellis coding

Appendix 4

1. The proposal is to define the band plans for the Downstream and upstream so as to reduce the power of the FFT/IFFT module and have efficient computations. With the band plans appropriately defined, only small number of computations will be required and saving significant power. The idea here is to propose a scheme such that the FFT/IFFT only processes the non-trivial data and hence saving large amount of power.
2. Current implementation is to process all data regardless of the fact that the FFT/IFFT is processing large amount of trivial data. The disadvantage here is that the FFT/IFFT requires significant power.
3. The invention here suggests defining the bands for the downstream and upstream within consideration to have efficient IFFT/FFT implementation. The FFT/IFFT does not have to process the samples with value zero, and some stages in the processing can be skipped or reduced significantly.
Hence, the advantage is a reduction of the power usage in the FFT/IFFT module during low-power mode. The power saving can be more than 50% depending on the number of tones that are used to carry the relevant data. Consequently, it lowers the heat dissipation requirements. This is particularly important for remote DSL equipment, where heat is a challenging problem.
4. Basic idea here is that the data in one direction is constrained to the tones which allow efficient implementation. If the downstream band is configured to tones which are power of 2, i.e. efficient implementation could be drawn and significant savings could be achieved. Due to the architecture of the FFT/IFFT, the samples taken for each butterfly computation in stage 0 involves at most one non-trivial sample. This significantly reduces the computation since the samples with value zero do not contribute to the computation and can be bypassed.
5. For butterflies of different radices there are potential savings. As an illustration, consider radix-2 butterfly. In the pre-processing stage, for IFFT each butterfly processing tone N takes the other sample from tone (N-4096). With the proposed band plan 997, each butterfly at the pre-processing stage, has at most one non-trivial input. Thus no complex addition is needed in these butterflies, only one complex multiplication is needed. There could be many other butterflies where both inputs are trivial inputs and could be skipped.

If the number of non-zero tones is 1k in the downstream direction, then only 1k complex multiplications are required. Other 3k butterfly processing stages could be skipped. Refer to Fig. 1. no more butterfly computation is needed in this stage. In the next subsequent stages, if the band separation is the number of such reduced radix-2 butterflies doubles from one stage to another as the stride taken halved from one stage to another.

[illegible]

Figure1: Band Plan 997

[illegible]

Figure2: Pre-processing




Figure 3: Radix-2 butterfly

The following explanation is based on 4K tones and radix-2 butterfly. There are 8 stages of FFT/IFFT processing and a pre-processing or post-processing stage. In the VDSL band plans, tones 513 to 575 are used for all downstream transmission profile (TS 101 270-2 V2.0.8). Assume that instead the band plan we adjusted as one possible illustration below whereby the tones configures in direction are aligned to get more non-trivial computations. As in Figure 4 for the data distribution in the memory. The radix-2 butterfly takes in 2 samples in stride of 2k for the first stage, for example $\{X_1, X_{2048}\}$. Since only X_{2048} is zero, the two outputs of the radix-2 butterfly are $\{X_1, X_1W_1\}$ where W_1 is the associated twiddle factors. From Figure 4, it is clear that

all the radix-2 butterflies in the pre-processing phase and also at Stage0 have only one non-zero value each. It is possible that the pre-processing and Stage0 twiddle factors be combined together into a single stage. These twiddle factors are a subset of the set of twiddle factors needed for the 4K tones IFFT, so no pre-computation and additional storage are required. The outputs are again the product of the non-zero input and a twiddle factor. This process saves one complete stage and significant reduction in computational cycles as well as power consumption. Those butterflies with all inputs of zero value do not have to be processed at all. The radix-2 butterflies in the first stage have either 1 non-zero input or all inputs of value zero. After the first stage of radix-2 processing, some samples still have zero value if the frequency band allocation is done appropriately. If the frequency band groups are chosen with steps of two to power n , 2^n , where n is an integer values most computational savings could be achieved while IFFT/FFT processing. The pre-processing stage of the IFFT computation could be combined with the processing of 1st stage of the IFFT processing hence saving computational cycles, memory bandwidth and power consumption.

Figure 4: With optimized band plan

[illegible]

↓ Preprocessing

[illegible]

↓ Stage 0

[illegible]

↓ Stage 1

[illegible][illegible][illegible]

↓ Stage 0

0	256	512	768	1024	1280	1536	1792	2048	2304	2560	2816	3072	3328	3584	3840
1	257	513	769	1025	1281	1537	1793	2049	2305	2561	2817	3073	3329	3585	3841
2	258	514	770	1026	1282	1538	1794	2050	2306	2562	2818	3074	3330	3586	3842
		575													
63															
							1984								
255	511	767	1023	1279	1535	1791	2047	2303	2559	2815	3071	3327	3583	3839	4095

↓ Stage 1

[illegible]

Figure 5: Example of optimized band plan 2 for IFFT processing

Appendix 5

1. The proposal is to reduce the power of the FFT/IFFT module during low-power mode in VDSL2. During L2 low-power mode, only small amount of data needs to be communicated between the ATU-C and the ATU-R. However, the FFT/IFFT processing cannot be turned off completely as the data still has to be processed. The idea here is to propose a scheme to use specific tones during L2 mode such that the FFT/IFFT only processes the non-trivial data and hence saving large amount of power.
2. Current implementation is to process all data regardless of the fact that the FFT/IFFT is processing large amount of trivial data during low-power mode. The disadvantage here is that even in low-power mode, the FFT/IFFT is still operating with full power.
3. The invention here limits the useful data to a small number of specific tones. The data in the remaining tones is all zeros. The FFT/IFFT does not have to process the samples with value zero, and some stages in the processing can be skipped or reduced significantly. Hence, the advantage is a reduction of the power usage in the FFT/IFFT module during low-power mode. The power saving can be more than 50% depending on the number of tones that are used to carry the relevant data. Consequently, it lowers the heat dissipation requirements. This is particularly important for remote DSL equipment, where heat is a challenging problem.
4. Basic idea here is that during L2 low-power mode, the vital data is constrained to just a small number of tones. Due to the architecture of the FFT/IFFT, the samples taken for each butterfly computation in stage 0 involves at most one non-trivial sample. This significantly reduces the computation since the samples with value zero do not contribute to the computation and can be bypassed.
5. The following explanation is based on 4K tones and radix-16 butterfly which is implementation using radix-4 (See Figure 1). There are 3 stages of FFT/IFFT processing and a pre-processing or post-processing stage. In the VDSL band plans, tones 513 to 575 are used for all downstream transmission profile (TS 101 270-2 V2.0.8). Assume these 63 tones are used in the low-power mode, and the remaining tones are zeros. After the pre-processing, tones {513...575} and {3521...3583} are non-zeros. See Figure 2 for the data distribution

in the memory. The radix-16 butterfly takes in 16 samples in stride of 256, for example $\{X_{11}, X_{257}, X_{513}, \dots, X_{3585}\}$. Each radix-4 butterfly in the first phase of the radix-16 butterfly takes the samples in stride of 1024. Samples $X_{513}, X_{1537}, X_{2561}$ and X_{3585} are taken for the butterfly processing. Since only X_{513} is non-zero, the four outputs of the radix-4 butterfly are $\{X_{513}, X_{513}W_1, X_{513}W_2, X_{513}W_3\}$ where W_1, W_2 and W_3 are the associated twiddle factors. From Figure 1, it is clear that the four radix-4 butterflies in the second phase have only one non-zero value each. The outputs are again the product of the non-zero input and a twiddle factor. Hence, the radix-16 butterfly outputs are $\{X_{513}, X_{513}T_1, X_{513}T_2, X_{513}T_3, X_{513}T_4, X_{513}T_5, X_{513}T_6, X_{513}T_7, X_{513}T_8, X_{513}T_9, X_{513}T_{10}, X_{513}T_{11}, X_{513}T_{12}, X_{513}T_{13}, X_{513}T_{14}, X_{513}T_{15}\}$, where T_i ($i=1, \dots, 15$) are the combined twiddle factors. These twiddle factors are a subset of the set of twiddle factors needed for the 4K tones IFFT, so no pre-computation and additional storage are required. The radix-16 butterflies in the first stage have either 1 non-zero input or all inputs of value zero. Those butterflies with all inputs of zero value do not have to be processed at all. After the first stage of radix-16 processing, some samples still have zero value as shown in Figure 2.




Figure 1 Radix 16 Butterfly

[illegible][illegible][illegible]

Figure 2 Data Distribution for IFFT processing

The second stage of the IFFT processing takes data in stride of 16, for example $\{Y_1, Y_{17}, Y_{33}, Y_{49}, Y_{65}, Y_{81}, Y_{97}, Y_{113}, Y_{129}, Y_{145}, Y_{161}, Y_{177}, Y_{193}, Y_{209}, Y_{225}, Y_{241}\}$. The first 4 samples and the last 4 samples are non-zero. Therefore, for each radix-4 butterfly in the first segment, two samples are non-zero and two are zeros. The number of complex additions is halved. There would be no saving in the second segment since the inputs to the 4 radix-4 butterflies are all non-zero.

There is no potential saving in the third stage. Table 1 shows the breakdown of the computation saving in each stage as a factor of the number of non-trivial tones. For the number of tones from 48 to 63, the pre-processing stage, stage 0 and stage 1 have a computation reduction of approximately 97%, 80% and 10% respectively (1 CMAC $\approx$ 5 CADD). Overall saving is about 40%. For number of tones less than 16, even stage 1 can be reduced by up to 40%.

There is no saving in the IFFT processing, if the number of tones exceed 127. Hence, we can limit the tones to the range from 513 to 639.

Table 1 Computation Complexity Table for IFFT

Number of non-zero tones	Pre-processing	Stage 0	Stage 1	Stage 2
0-15	0-15 BF2	0-30 BF16'	$256 \times (2 \text{ BF}_{4_1} + 4 \text{ BF}_{4_2})$ (1 trivial BF_{4_1})	256 BF16
16-31	16-31 BF2	32-62 BF16'	$256 \times (4 \text{ BF}_{4_1} + 4 \text{ BF}_4)$ (1 trivial BF_{4_1}) $256 \times (2 \text{ BF}_{4_1} + 2 \text{ BF}_{4_2} + 4 \text{ BF}_4)$ (1 trivial BF_{4_1})	256 BF16
32-47	32-47 BF2	64-94 BF16'		256 BF16
48-63	48-63 BF2	96-126 BF16'	$256 \times (4 \text{ BF}_{4_2} + 4 \text{ BF}_4)$	256 BF16
64-79	64-79 BF2	128-158 BF16'	$256 \times (2 \text{ BF}_4 + 2 \text{ BF}_{4_2} + 4 \text{ BF}_4)$	256 BF16
80-95	80-95 BF2	160-190 BF16'	256 BF16	256 BF16
96-127	96-127 BF2	192-254 BF16'	256 BF16	256 BF16
128-143	128-143 BF2	256 BF16	256 BF16	256 BF16
4K	2K BF2	256 BF16	256 BF16	256 BF16
		BF16'=15 CMAC, 0 CADD	BF _{4₁} =3 CMAC, 0 CADD (room for optimization) BF _{4₂} =4 CADD, 3 CMAC BF ₄ =8 CADD, 3 CMAC	

The following analyses the power saving for the FFT processing. Assume that in low-power mode, the 64 tones from 960 to 1023 are non-trivial. Refer to Figure 3. These tones fit into the upstream spectrum for all transmission profiles. In stage 0, there are 64 radix-16 butterflies with a single non-zero input. These butterflies, similar to those discussed in the IFFT section, do not need any complex addition, only 15 complex multiplications are required. The remaining butterflies in this stage are trivial. In stage 1, all 256 radix-16 butterflies have to be processed, but with substantial saving in the computation. Each of the 4 radix-4 butterflies in the first segment has only a single non-zero input. Hence, they do not need any addition, only 3 multiplications are required. There is no potential saving in the next segment since all the samples are non-trivial. There are no saving in stage 2 and the post-processing. Computation saving in stage 0 and stage 1 are 84% and 18% respectively. Overall saving is 35%. The tones recommended for the low-power mode is from 960 to 1023. Only 64 tones are proposed since the upstream has lower traffic.

[illegible]

Stage 1

[illegible]

Stage 2

10	256	512	768	1024	1280	1536	1792	2048	2304	2560	2816	3072	3328	3584	3840
1	257	513	769	1025	1281	1537	1793	2049	2305	2561	2817	3073	3329	3585	3841
2	258	514	770	1026	1282	1538	1794	2050	2306	2562	2818	3074	3330	3586	3842
3	259	515	771	1027	1283	1539	1795	2051	2307	2563	2819	3075	3331	3587	3843
4	260	516	772	1028	1284	1540	1796	2052	2308	2564	2820	3076	3332	3588	3844
5	261	517	773	1029	1285	1541	1797	2053	2309	2565	2821	3077	3333	3589	3845
6	262	518	774	1030	1286	1542	1798	2054	2310	2566	2822	3078	3334	3590	3846
7	263	519	775	1031	1287	1543	1799	2055	2311	2567	2823	3079	3335	3591	3847
8	264	520	776	1032	1288	1544	1800	2056	2312	2568	2824	3080	3336	3592	3848
9	265	521	777	1033	1289	1545	1801	2057	2313	2569	2825	3081	3337	3593	3849
10	266	522	778	1034	1290	1546	1802	2058	2314	2570	2826	3082	3338	3594	3850
11	267	523	779	1035	1291	1547	1803	2059	2315	2571	2827	3083	3339	3595	3851
12	268	524	780	1036	1292	1548	1804	2060	2316	2572	2828	3084	3340	3596	3852
13	269	525	781	1037	1293	1549	1805	2061	2317	2573	2829	3085	3341	3597	3853
14	270	526	782	1038	1294	1550	1806	2062	2318	2574	2830	3086	3342	3598	3854
15	271	527	783	1039	1295	1551	1807	2063	2319	2575	2831	3087	3343	3599	3855
16	272	528	784	1040	1296	1552	1808	2064	2320	2576	2832	3088	3344	3600	3856
17	273	529	785	1041	1297	1553	1809	2065	2321	2577	2833	3089	3345	3601	3857
18	274	530	786	1042	1298	1554	1810	2066	2322	2578	2834	3090	3346	3602	3858
19	275	531	787	1043	1299	1555	1811	2067	2323	2579	2835	3091	3347	3603	3859
20	276	532	788	1044	1300	1556	1812	2068	2324	2580	2836	3092	3348	3604	3860
21	277	533	789	1045	1301	1557	1813	2069	2325	2581	2837	3093	3349	3605	3861
22	278	534	790	1046	1302	1558	1814	2070	2326	2582	2838	3094	3350	3606	3862
23	279	535	791	1047	1303	1559	1815	2071	2327	2583	2839	3095	3351	3607	3863
24	280	536	792	1048	1304	1560	1816	2072	2328	2584	2840	3096	3352	3608	3864
25	281	537	793	1049	1305	1561	1817	2073	2329	2585	2841	3097	3353	3609	3865
26	282	538	794	1050	1306	1562	1818	2074	2330	2586	2842	3098	3354	3610	3866
27	283	539	795	1051	1307	1563	1819	2075	2331	2587	2843	3099	3355	3611	3867
28	284	540	796	1052	1308	1564	1820	2076	2332	2588	2844	3100	3356	3612	3868
29	285	541	797	1053	1309	1565	1821	2077	2333	2589	2845	3101	3357	3613	3869
30	286	542	798	1054	1310	1566	1822	2078	2334	2590	2846	3102	3358	3614	3870
31	287	543	799	1055	1311	1567	1823	2079	2335	2591	2847	3103	3359	3615	3871
32	288	544	800	1056	1312	1568	1824	2080	2336	2592	2848	3104	3360	3616	3872
33	289	545	801	1057	1313	1569	1825	2081	2337	2593	2849	3105	3361	3617	3873
34	290	546	802	1058	1314	1570	1826	2082	2338	2594	2850	3106	3362	3618	3874
35	291	547	803	1059	1315	1571	1827	2083	2339	2595	2851	3107	3363	3619	3875
36	292	548	804	1060	1316	1572	1828	2084	2340	2596	2852	3108	3364	3620	3876
37	293	549	805	1061	1317	1573	1829	2085	2341	2597	2853	3109	3365	3621	3877
38	294	550	806	1062	1318	1574	1830	2086	2342	2598	2854	3110	3366	3622	3878
39	295	551	807	1063	1319	1575	1831	2087	2343	2599	2855	3111	3367	3623	3879
40	296	552	808	1064	1320	1576	1832	2088	2344	2600	2856	3112	3368	3624	3880
41	297	553	809	1065	1321	1577	1833	2089	2345	2601	2857	3113	3369	3625	3881
42	298	554	810	1066	1322	1578	1834	2090	2346	2602	2858	3114	3370	3626	3882
43	299	555	811	1067	1323	1579	1835	2091	2347	2603	2859	3115	3371	3627	3883
44	300	556	812	1068	1324	1580	1836	2092	2348	2604	2860	3116	3372	3628	3884
45	301	557	813	1069	1325	1581	1837	2093	2349	2605	2861	3117	3373	3629	3885
46	302	558	814	1070	1326	1582	1838	2094	2350	2606	2862	3118	3374	3630	3886
47	303	559	815	1071	1327	1583	1839	2095	2351	2607	2863	3119	3375	3631	3887
48	304	560	816	1072	1328	1584	1840	2096	2352	2608	2864	3120	3376	3632	3888
49	305	561	817	1073	1329	1585	1841	2097	2353	2609	2865	3121	3377	3633	3889
50	306	562	818	1074	1330	1586	1842	2098	2354	2610	2866	3122	3378	3634	3890
51	307	563	819	1075	1331	1587	1843	2099	2355	2611	2867	3123	3379	3635	3891
52	308	564	820	1076	1332	1588	1844	2100	2356	2612	2868	3124	3380	3636	3892
53	309	565	821	1077	1333	1589	1845	2101	2357	2613	2869	3125	3381	3637	3893
54	310	566	822	1078	1334	1590	1846	2102	2358	2614	2870	3126	3382	3638	3894
55	311	567	823	1079	1335	1591	1847	2103	2359	2615	2871	3127	3383	3639	3895
56	312	568	824	1080	1336	1592	1848	2104	2360	2616	2872	3128	3384	3640	3896
57	313	569	825	1081	1337	1593	1849	2105	2361	2617	2873	3129	3385	3641	3897
58	314	570	826	1082	1338	1594	1850	2106	2362	2618	2874	3130	3386	3642	3898
59	315	571	827	1083	1339	1595	1851	2107	2363	2619	2875	3131	3387	3643	3899
60	316	572	828	1084	1340	1596	1852	2108	2364	2620	2876	3132	3388	3644	3900
61	317	573	829	1085	1341	1597	1853	2109	2365	2621	2877	3133	3389	3645	3901
62	318	574	830	1086	1342	1598	1854	2110	2366	2622	2878	3134	3390	3646	3902
63	319	575	831	1087	1343	1599	1855	2111	2367	2623	2879	3135	3391	3647	3903
64	320	576	832	1088	1344	1600	1856	2112	2368	2624	2880	3136	3392	3648	3904
65	321	577	833	1089	1345	1601	1857	2113	2369	2625	2881	3137	3393	3649	3905
66	322	578	834	1090	1346	1602	1858	2114	2370	2626	2882	3138	3394	3650	3906
67	323	579	835	1091	1347	1603	1859	2115	2371	2627	2883	3139	3395	3651	3907
68	324	580	836	1092	1348	1604	1860	2116	2372	2628	2884	3140	3396	3652	3908
69	325	581	837	1093	1349	1605	1861	2117	2373	2629	2885	3141	3397	3653	3909
70	326	582	838	1094	1350	1606	1862	2118	2374	2630	2886	3142	3398	3654	3910
71	327	583	839	1095	1351	1607	1863	2119	2375	2631	2887	3143	3399	3655	3911
72	328	584	840	1096	1352	1608	1864	2120	2376	2632	2888	3144	3400	3656	3912
73	329	585	841	1097	1353	1609	1865	2121	2377	2633	2889	3145	3401	3657	3913
74	330	586	842	1098	1354	1610	1866	2122	2378	2634	2890	3146	3402	3658	3914
75	331	587	843	1099	1355	1611	1867	2123	2379	2635	2891	3147	3403	3659	3915
76	332	588	844	1100	1356	1612	1868	2124	2380	2636	2892	3148	3404	3660	3916
77	333	589	845	1101	1357	1613	1869	2125	2381	2637	2893	3149	3405	3661	3917
78	334	590	846	1102	1358	1614	1870	2126	2382	2638	2894	3150	3406	3662	3918
79	335	591	847	1103	1359	1615	1871	2127	2383	2639	2895	3151	3407	3663	3919
80	336	592	848	1104	1360	1616	1872	2128	2384	2640	2896	3152	3408	3664	3920
81	337	593	849	1105	1361	1617	1873	2129	2385	2641	2897	3153	3409	3665	3921
82	338	594	850	1106	1362	1618	1874	2130	2386	2642	2898	3154	3410	3666	3922
83	339	595	851	1107	1363	1619	1875	2131	2387	2643	2899	3155	3411	3667	3923
84	340	596	852	1108	1364	1620	1876	2132	2388	2644	2900	3156	3412	3668	3924
85	341	597	853	1109	1365	1621	1877	2133	2389	2645	2901	3157	3413	3669	3925
86	342	598	854	1110	1366	1622	1878	2134	2390	2646					

Figure 3 Data Distribution for FFT

Table 2 Computation Complexity Table for FFT

Number of non-zero tones	Stage 0	Stage 1	Stage 2	Post-processing
0-16	0-16 BF16'	$256*(BF4_1+4 BF4_1)$	256 BF16	2K BF2
17-32	17-32 BF16'	$256*(2 BF4_1+4 BF4_2)$	256 BF16	2K BF2
33-48	33-48 BF16'	$256*(3 BF4_1+4 BF4_3)$	256 BF16	2K BF2
49-64	49-64 BF16'	$256*(4 BF4_1+4 BF4)$	256 BF16	2K BF2
65-80	65-80 BF16'	$256*(3 BF4_1 + BF4_2 + 4 BF4)$	256 BF16	2K BF2
81-96	81-96 BF16'	$256*(2 BF4_1 + 2 BF4_2 + 4 BF4)$	256 BF16	2K BF2
97-112	97-112 BF16'	$256*(BF4_1 + 3 BF4_2 + 4 BF4)$	256 BF16	2K BF2
113-128	113-128 BF16'	$256*(4 BF4_2 + 4 BF4)$	256 BF16	2K BF2
129-144	129-144 BF16'	$256*(BF4_3 + 3 BF4_2 + 4 BF4)$	256 BF16	2K BF2
145-160	145-160 BF16'	$256*(2 BF4_3 + 2 BF4_2 + 4 BF4)$	256 BF16	2K BF2
161-176	161-176 BF16'	$256*(3 BF4_3 + BF4_2 + 4 BF4)$	256 BF16	2K BF2
177-192	177-192 BF16'	$256*(4 BF4_3 + 4 BF4)$	256 BF16	2K BF2
193-208	193-208 BF16'	$256*(3 BF4_3 + 5 BF4)$	256 BF16	2K BF2
209-224	209-224 BF16'	$256*(2 BF4_3 + 6 BF4)$	256 BF16	2K BF2
225-240	225-240 BF16'	$256*(1 BF4_3 + 7 BF4)$	256 BF16	2K BF2
241-256	241-256 BF16'	256 BF16	256 BF16	2K BF2
4K	256 BF16	256 BF16	256 BF16	2K BF2

BF16'=15 CMAC,
0 CADD

BF4₁=0 CADD, 3 CMAC
BF4₂=4 CADD, 3 CMAC
BF4₃=6 CADD, 3 CMAC
BF4=8 CADD, 3 CMAC

Claims

1. A method of two-directional communication of data over a line, the communication employing a bandwidth partitioned into a number of bands including at least one band associated with each of the two directions, the
5 method including encoding the data carried in each direction by modulation of a plurality of tones defined within the at least one band associated with that direction, the tones being associated into a plurality of groups of tones, each group of tones including a plurality of the tones, the method including for each group of tones:
 - 10 obtaining data which is to be encoded by modulation of that group of tones;
 - modulating the tones of the group to encode the data.
2. A method according to claim 1 further including trellis encoding together the data to be encoded by modulation of that group of tones, the
15 group of tones being modulated using the trellis encoded data.
3. A method according to claim 1 further including:
 - deriving a measure of the quality of transmission of each of the tones;
 - and
 - the relative transmission quality of the tones of the group determining
20 how the tones of that group are modulated using the data.
4. A method according to any preceding claim in which a selection is made from the tones of the group to determine which of at least two processing paths is used to encode data on each tone.

5. A method according to claim 4 in which the two processing paths include paths have a different latency.
6. A method according to any preceding claim in which different ones of the tones of a group are used to carry different respective proportions of the data.
7. A method according to claim 6 further including a re-determination of the transmission quality of the tones on the line, according to the re-determination the respective proportions of the data encoded using the respective tones being altered.
8. A method according to any preceding claim in which a selection is made from the tones of the group to determine a respective gain for each of the tones.
9. A method according to any preceding claim in which there are multiple groups within each of the bands, all the tones of any given group being within a single one of the bands.
10. A method according to any of claims 1 to 8 in which one or more of the groups include tones within at least two different respective bands.
11. A method according to claim 10 in which each of the groups includes tones from multiple bands.
12. A method of two-directional communication of data over a line, the communication employing a bandwidth partitioned into a number of bands including at least one band associated with each of the two directions, the method including encoding the data carried in each direction by modulation of a plurality of tones defined within the at least one band associated with that direction, the bands being selected within the bandwidth such that all the bands in a given direction have a number of tones which is substantially equal

to a power of 2, and the bands are spaced apart by a number of tones which is substantially equal to a power of 2.

13. A method of two-directional communication of data over a line, the communication employing a bandwidth partitioned into a number of bands including at least one band associated with each of the two directions, the method including encoding the data carried in each direction by modulation of a plurality of tones defined within the at least one band associated with that direction, the modulation being implemented using an IFFT of data indicating the amplitudes of the respective tones,
- 10 the method including operating selectively in one of at least two modes,
- in a first said mode the data to be transmitted in each direction being encoded on into all the tones of the bands associated with that direction, and
- in a second said mode the data to be transmitted in at least one direction being encoded into only a subset of the tones of the bands associated with that direction, the IFFT operating in the second mode by an algorithm which is computationally less expensive than in the first mode.
- 15
14. A method of two-directional communication of data over a line, the communication employing a bandwidth partitioned into a number of bands including at least one band associated with each of the two directions, the method including encoding the data carried in each direction by modulation of a plurality of tones defined within the at least one band associated with that direction, the method including operating selectively in one of at least two modes,
- 20
- in a first said mode, the data to be transmitted in each direction being encoded on into all the tones of the bands associated with that direction, and
- 25

the modulation being implemented using a first IFFT unit which receives data describing the amplitudes of the respective tones,

5 in a second said mode, the data to be transmitted in at least one direction being encoded into only a low-frequency subset of the tones of the bands associated with that direction, the modulation being performed by a second IFFT which receives only data describing the amplitudes of the respective tones, and the output of the second IFFT being up-sampled to a frequency rate as high as that in the first mode of operation.

10 15. A method of two-directional communication of data over a line, the communication employing a bandwidth partitioned into a number of bands including at least one band associated with each of the two directions, the method including encoding the data carried in each direction by modulation of a plurality of tones defined within the at least one band associated with that direction, the method including operating selectively in one of at least two
15 modes,

in a first said mode, the data in each direction being encoded in frequencies up to a first respective upper frequency,

20 in a second said mode, the data to be transmitted in at least one direction being encoded in frequencies up to a second upper frequency lower than the respective first upper frequency.

16. A method according to claim 15 in which the second mode of operation is according to an ADSL standard.

25 17. A method of two-directional communication of data over a line, the communication employing a bandwidth partitioned into a number of bands including at least one band associated with each of the two directions, the method including encoding the data carried in each direction by modulation of

a plurality of tones defined within the at least one band associated with that direction, the method including operating selectively in one of at least two modes,

5 in a first said mode, the data transmitted in at least one direction being encoded in frequencies including a high frequency range;

in a second said mode, the data transmitted in that direction being encoded in frequencies including the high frequency range;

the rate of data transmission in the high frequency range being higher in the first mode than in the first mode.