

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
25 August 2011 (25.08.2011)

(10) International Publication Number
WO 2011/101443 A1

(51) International Patent Classification:
H03M 7/40 (2006.01)

(21) International Application Number:
PCT/EP2011/052449

(22) International Filing Date:
18 February 2011 (18.02.2011)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
61/306,381 19 February 2010 (19.02.2010) US
12/838,118 16 July 2010 (16.07.2010) US

(71) Applicant (for all designated States except US): **SKYPE LIMITED** [IE/IE]; Arthur Cox Building, Earlsfort Centre, Earlsfort Terrace, Dublin, 2 (IE).

(72) Inventors; and

(75) Inventors/Applicants (for US only): **VAFIN, Renat** [EE/EE]; Raagu 33-7, EE-13417 Tallinn (EE). **BIVOLARSKY, Lazar** [BG/US]; 11024 Holly Tree Lane, Cupertino, CA 95014 (US). **NILSSON, Mattias** [SE/SE]; Egnahemsvagen 5, S-17446 Sundbyberg (SE). **ANDERSEN, Soren Vang** [DK/LU]; 38 Rue Wilson, L-2732 Luxembourg (LU).

(74) Agents: **DRIVER, Virginia, Rozanne** et al.; Page White & Farrer, Bedford House, John Street, London Greater London WC1N 2BF (GB).

(81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IS, JP, KE, KG, KM, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LT, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PE, PG, PH, PL, PT, RO, RS, RU, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

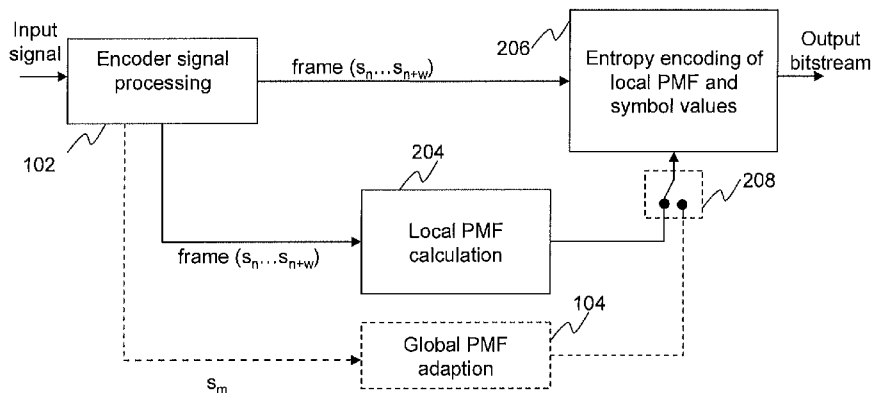
(84) Designated States (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, GH, GM, KE, LR, LS, MW, MZ, NA, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

Published:

— with international search report (Art. 21(3))

(54) Title: ADAPTIVE ENTROPY ENCODING

Figure 2a



(57) Abstract: An encoder stage, and corresponding encoded bitstream and decoder. The encoder stage comprises: a variable length encoder for encoding an input signal; and a counter configured to dynamically detect an observed frequency at which different symbols are found to occur within each of a plurality of predetermined portions of the input signal, prior to the symbols of each respective portion being encoded by the variable length encoder. The variable length encoder is configured to encode the symbols of each portion using variable length coding performed in dependence on the observed frequencies detected within the respective portion of the input signal, to generate an encoded bitstream comprising the encoded symbols along with an additional element indicating information regarding the observed frequencies detected for each portion, and to output the encoded bitstream to at least one of a storage medium and a transmission medium for supply to a decoder.

WO 2011/101443 A1

ADAPTIVE ENTROPY ENCODING

Field of the Invention

5 The present invention relates to variable length coding (VLC), whereby a bitstream comprising a plurality of symbols is encoded using variable length codes. This is sometimes referred to as entropy encoding, because it allows the average number of bits per symbol to be reduced to closer to the theoretical minimum (the "entropy") for a particular bitstream.

10

Background

A simple coding scheme would encode all symbol values with the same number of bits. The minimum required number of bits per value is $\log_2(n_v)$, where n_v is the number of possible values a symbol can take. The average number of bits over
15 all values is therefore also $\log_2(n_v)$.

A more efficient coding scheme encodes symbol values that are expected to occur more frequently with fewer bits, and symbol values that are expected to occur less frequently with more bits. There are a number of ways to do this, but
20 for illustration one way is to let the number of leading zeros in the encoded symbol indicate the length of the encoded symbol. Different possible symbol values v are then mapped to the different length symbol codes c , e.g.:

25	Code number	Symbol code (c)	Maps to symbol value (v)
	0	1	v_a
	1	010	v_b
	2	011	v_c
	3	00100	v_d
30	4	00101	v_e
	5	00110	v_f

2

6	00111	v_g
7	0001000	v_h
8	0001001	v_i
etc...	etc...	etc...

5

The encoder and decoder both have a pre-stored copy of this VLC look-up table. So if the decoder receives a 1 as the first bit of an encoded symbol, it knows this is a one-bit encoded symbol representing code number 0, which maps to symbol value v_a . If the decoder receives one leading 0 before receiving the first 1, it
 10 knows this is a three-bit encoded symbol representing code number 1 or 2 depending on the third bit, which map to symbol values v_b and v_c respectively. If the decoder receives two leading 0s before receiving the first 1, it knows this is a five-bit encoded symbol representing one of code numbers 3 to 6 depending on the fourth and fifth bits, which map to symbol values v_d to v_g respectively; and so
 15 forth.

This is an example of a set of "prefix" codes, whereby no shorter code is equal to the prefix of any longer code. As mentioned however, this is only one example given for illustrative purposes. Other sets of prefix codes are in common use,
 20 such as Huffman codes for example. Other types of variable length coding schemes are also possible, such as those not based on this principle of prefixes, or those that are not in themselves uniquely decodable and rely on some additional known constraint on the bitstream in order to decode. Another type of
 25 variable length encoding is arithmetic encoding, which does not map each symbol value onto an individual respective code, but instead encodes a group of symbols together with one overall code.

Such variable length schemes result in a lower average number of bits over all values, compared to assigning the same number of bits to all symbol values. The
 30 theoretically minimum possible average number of bits is called entropy and is defined as $H = -\sum_i (p_i \cdot \log_2(p_i))$, where p_i is the probability of the i^{th} symbol value.

Therefore such a coding scheme is sometimes called an entropy coder.

Entropy coding relies on availability of some information regarding the relative probabilities of symbol values. Such information may be referred to as the
5 “probability mass function” (PMF), i.e. in this case the likely frequency at which the different possible symbol values are expected to occur relative to one another within the bitstream. Information about the PMF is required at both the encoder and the decoder side (the latter is needed for correct decoding). For example, this information may take the form of a VLC look-up table pre-stored at the
10 encoder and encoder, mapping of symbol values v to the different codes c .

For efficient entropy coding, the available PMF should represent the expected frequencies as accurately as possible. To achieve this, the PMF is conventionally pre-trained at the design stage by using a large set of data that represent
15 symbols to be encoded. This “globally” trained PMF is then pre-stored at the encoder and the decoder.

To further increase coding efficiency, the pre-stored PMF is often adapted during the coding and decoding process. In this case, a small number of predetermined
20 alternative PMFs are pre-trained for different scenarios, and each alternative is pre-stored at both the encoder and decoder. The encoder and decoder then select the relevant PMF for encoding and decoding a current symbol by observing the previously-coded symbol values.

25 Examples of such adaptive entropy coding schemes are context-adaptive variable length coding (CAVLC) and context-adaptive binary arithmetic coding (CABAC) used in MPEG Video coding. These are currently defined in the H.264 standard of the Joint Video Team (JVT) of the International Organization for Standardization (ISO) and International Electrotechnical Commission (IEC),
30 comprising the Moving Picture Experts Group (MPEG) and International Telecommunications Union Video Coding Experts Group (ITU-T VCEG).

Figure 1a is a schematic block diagram of a conventional CAVLC video encoder. Each video frame (i.e. the image at each moment in time) is sub-divided into a plurality of smaller blocks. The input signal is first passed through earlier signal processing stages 102 which perform encoding such as intraframe prediction encoding and interframe motion prediction, to achieve an initial compression by exploiting properties of the video content such as correlation between nearby blocks and movement of blocks between frames. The input signal is then passed to an entropy encoder 106 for further compression based on variable length coding (VLC). A handful of different pre-trained PMFs are pre-stored at the entropy encoder 106, in the form of a plurality of predetermined VLC look-up tables each being a potential candidate for encoding a current symbol s_n .

To achieve adaption, a global PMF adaptation block 104 is arranged to select an appropriate one of the VLC look-up tables for encoding a symbol s_n of the current block. The selection is based on certain previously encoded symbols s_m . For example in CAVLC the selection is based on the number of non-zero coefficients in certain previously encoded blocks, and on previously encoded elements in the current block.

A corresponding decoder is illustrated in Figure 1b. The incoming bit stream is first passed to an entropy decoder 116, then to later stages 112 such as intraframe prediction decoding and interframe motion prediction decoding stages. The same few pre-trained VLC look-up tables as stored at the encoder are also pre-stored at the entropy decoder 116. The decoder also has the same knowledge of the previous symbol values $v[s_m]$ as the encoder (because it has already just decoded those symbols). Given this knowledge of previously decoded symbols s_m , a global PMF adaption block 114 in the decoder can therefore deterministically ascertain the VLC look-up table that the encoder must have used to encode a current encoded symbol $c(v[s_n])$, and thus control the entropy decoder 116 to decode the current symbol s_n accordingly.

CABAC works in a similar manner, but can also apply an additional adaptation by updating the global PMF based on a previously encoded block and previously encoded symbols in the current block. That is, CABAC adapts not only by
5 selecting between pre-stored look-up tables, but also by modifying the pre-stored look-up tables based on previously encoded symbol values. The decoder can determine the same modification to apply to the look-up tables as applied by the encoder, because the decoder has already previously decoded the same symbols that were used to determine that modification at the encoder.

10

Summary

As discussed, in prior systems the VLC look-up tables are pre-stored at the encoder and decoder, and so can only describe how frequently symbol values
15 are *expected* to occur. That is, the PMF is pre-trained based on a large set of test data at the design stage. Even in the adaptive case, this only involves selecting between a few pre-stored, pre-trained VLC look-up tables based on certain previously encoded symbols, and/or by updating a pre-trained PMF based on previously encoded symbols.

20

According to one aspect of the present invention on the other hand, there is provided an encoder stage for encoding an input signal comprising a plurality of different symbols, the encoder stage comprising: a variable length encoder for encoding the input signal; and a counter configured to dynamically detect an
25 observed frequency at which the different symbols are found to occur within each of a plurality of predetermined portions of the input signal, prior to the symbols of each respective portion being encoded by the variable length encoder; wherein the variable length encoder is configured to encode the symbols of each portion using variable length coding performed in dependence on the observed
30 frequencies detected within the respective portion of the input signal, to generate an encoded bitstream comprising the encoded symbols along with an additional

element indicating information regarding the observed frequencies detected for each portion, and to output the encoded bitstream to at least one of a storage medium and a transmission medium for supply to a decoder.

5 The present invention thus trains a probability function dynamically, "on the fly", based on actual observed symbol values in the input signal to be encoded. That is, it makes an *a posteriori* empirical measurement of how frequently symbols actually occur within a given portion of the input signal such as a certain time interval, and uses that information to encode the symbols of that same interval of
10 the input signal. It also outputs information about the observed frequency into the encoded bitstream for supply to the decoder, so that the decoder can determine how the symbols of each interval were encoded and therefore decode them accordingly.

15 Preferably the input signal comprises at least one of a video signal, a speech signal or another kind of audio signal. In this case the time intervals used in the above approach are preferably video frames (each representing the video image at a different respective moment in time), speech frames (each representing a speech interval such as predetermined number of voice pitch periods), and/or
20 audio frames of an audio standard such as an MPEG audio standard. The invention may thus allow the PMF to be adapted on a per frame basis, with the probability mass function (PMF) estimated on the current audio or video frame being used to facilitate efficient entropy coding of symbol values within the frame. The PMF is transmitted along with entropy-coded symbol values to the decoder
25 as side information.

Referred to in the context of the present invention therefore, the PMF does not describe the likely frequency at which symbol values are expected to occur based on a pre-trained system, but rather the actual frequency at which they did
30 occur within the frame or interval in question. This may be referred to herein as a local PMF.

While a (possibly adaptive) globally pre-trained PMF may typically provide a good representation of occurrence of symbol values on average, the present invention improves on this by providing a high-degree of adaptation to the current symbol values (e.g. within a frame). The result is a lower average number of bits per symbol in the output bit stream.

The present invention does require some extra information about the PMF to be included in the output bitstream. This is unlike the global adaptive PMF techniques such as CAVLC and CABAC, which do not insert information about the PMF into the output bitstream. Instead, because these conventional global techniques adapt the VLC coding based on previously encoded symbols, then the decoder can determine the adaptation applied to a current symbol based on the previously decoded symbols using the same criterion as the encoder, without needing explicit information about the PMF to be provided in the bitstream.

However, because a frame or other short interval typically only contains few symbols, the inventors have recognised that the resulting PMF will usually be sparse and therefore can be efficiently encoded with a small number of bits (unlike the global PMF trained on a large set of data which would not be sparse and would not be included in the encoded output bitstream). So although the present invention does require some extra information about the PMF to be included in the output bitstream, the inventors have recognised that this is outweighed by the reduced number of average bits per symbol achieved by the local adaptation based on actual observed symbol values. Therefore, the present invention can still result in a lower average number of bits for coding of symbol values compared to the global PMF techniques, with the sparse local PMF itself being efficiently coded with a low number of bits of side information.

If the number of bits required to encode symbol values using the conventional global PMF is $R(\text{symbols, global})$; the number of bits required to encode symbol

values using the improved, highly-adaptive local PMF of the present invention is $R(\text{symbols, local})$; and the number of bits required to transmit side information for such local adaption is $R(\text{side info})$; then the present invention will in many cases achieve: $R(\text{symbols, local}) + R(\text{side info}) < R(\text{symbols, global})$.

5

In embodiments, the variable length encoder may be arranged to encode more frequent symbols of the input signal with codes having fewer bits, and to encode less frequent symbols of the input signal with codes having more bits; wherein the variable length encoder may be configured to map the different length codes to the symbols based on the observed frequencies detected within said portion of the input signal, and encode the symbols of said portion of the input signal using said mapping.

In further embodiments, the variable length encoder may be configured to switch between (i) a first mode of operation in which the symbols of each portion are encoded in dependence on the observed frequencies detected for the respective portion, and (ii) a second mode of operation in which the variable length encoding is either non-adaptive or adapted based on previously encoded symbols.

The variable length encoder may be configured to determine whether to switch between said modes of operation once per number of said portions, in dependence on a relative efficiency metric of the modes for the respective number of portions.

Said number may be one, such that the variable length encoder may be configured to determine whether to switch between said modes of operation once per portion, in dependence on a relative efficiency metric of the modes for the respective portion.

The variable length encoder may be configured to encode said information on the observed frequencies for inclusion in the output bitstream, based on a known

constraint on a total number of occurrences within each portion.

The variable length encoder may be configured to encode said information on the observed frequencies for inclusion in the output bitstream by predicting the
5 occurrences of less frequent symbols based on the occurrences of more frequent symbols.

The variable length encoder may be configured to encode said information on the observed frequencies for inclusion in the output bitstream using run-length
10 coding.

The variable length encoder may be configured to encode said information on the observed frequencies for inclusion in the output bitstream by differentially encoding the information of one portion relative to one of: the information of a
15 preceding portion, and a global PMF.

The variable length encoder may be configured to encode said information on the observed frequencies for inclusion in the output bitstream, using a non-lossless encoding.
20

The counter may be configured to update the observed frequencies for a portion of the input signal after encoding each of a plurality of successive ones of said symbols within the respective portion; and the variable length encoder may be configured to encode each successive symbol of the respective portion using
25 variable length coding performed in dependence on the updated frequencies.

According to another aspect of the present invention, there is provided a bitstream arranged into a plurality of portions, each comprising: a plurality of different encoded symbols, encoded according to an adaptive variable length coding scheme, and an additional element indicating information regarding an observed frequency at which the different symbols occur within the respective
5 portion; wherein the symbols of each portion are encoded according to a variable length encoding scheme adapted in dependence on the observed frequencies of the different symbols within the respective portion, said information being for use by a decoder in decoding each portion.

10

In embodiments, the bitstream may be encoded according to any of the above features of the encoder.

15

According to another aspect of the present invention, there is provided a network equipment comprising a transmission medium conveying the bitstream.

20

According to another aspect of the present invention, there is provided a decoder stage, comprising: a variable length decoder for decoding a bitstream arranged into a plurality of portions, and each portion comprising a plurality of different
20 encoded symbols and an additional element indicating information regarding an observed frequency at which the different symbols occur within the respective portion, wherein the symbols of each portion are encoded according to a variable length encoding scheme adapted in dependence on the observed frequencies of the different symbols within the respective portion; and wherein the variable
25 length decoder is configured to decode the symbols of each portion in dependence on the received information regarding the observed frequencies, and to thereby generate an output signal for output to a user device.

25

30

In embodiments, the variable length decoder may be further configured to decode a bitstream encoded according to the above features of the encoder.

Brief Description of the Drawings

For a better understanding of the present invention and to show how it may be put into effect, reference is made by way of example to the accompanying drawings in which:

Figure 1a is a schematic block diagram of an encoder,
Figure 1b is a schematic block diagram of a decoder,
Figure 2a is a schematic block diagram of an encoder,
10 Figure 2b is a schematic representation of an output bitstream, and
Figure 2c is a schematic block diagram of a decoder.

Detailed Description of Preferred Embodiments

15 As discussed, the present invention provides a high-degree of adaptability of an entropy-coding scheme to the actual frequency of occurrence of symbol values within a short interval (e.g., within one audio or video frame), thereby reducing the average number of bits per symbol required for encoding. In preferred embodiments, the PMF is estimated for each frame based on the symbol values
20 in that same frame. This PMF is used to entropy-code the symbol values within the frame and is transmitted along with the entropy-coded symbol values to the decoder as side information to facilitate decoding.

Figure 2a shows an encoder according to an exemplary embodiment of the present invention. The encoder comprises earlier encoder stages 102 arranged
25 to receive an input signal such as a video signal, speech signal or other audio signal. These stages 102 are arranged to perform prediction coding such as intraframe prediction coding and interframe motion prediction in the case of video coding, or short- and long-term prediction coding in the case of speech encoding,
30 as will be familiar to a person skilled in the art. Following this the encoder comprises an entropy encoder 206 arranged to receive the (now partially

processed) input signal after passing through the earlier encoder stages 102. The entropy encoder is arranged to perform variable length entropy encoding on this received input signal, to thereby generate an output bitstream for output to a transmission medium (e.g. for transmission over a network such as the internet) or storage medium (e.g. for storage on a hard drive or flash memory).

The encoder further comprises a local PMF calculation block 204, having an input also arranged to receive the (partially processed) input signal from the earlier stages 102, and having an output arranged to supply control signals to the entropy encoder 206. However, this local PMF calculation block is not like the global PMF adaption block 104 of Figure 1a, which selects an adaption based on symbols of the input signal that have already been received and encoded by the entropy encoder 106 (e.g. selected based on previous symbols s_m whilst the current block being encoded by the entropy encoder 106 is a subsequent symbol s_n). That is, the adaptation by the global PMF block 104 is obtained by observing the previously encoded symbol values. In contrast, the local PMF calculation block 204 of Figure 2a comprises a counter which counts the actual number of occurrences of the different possible symbol values within each frame (e.g. comprising symbols $s_n \dots s_{n+w}$ of a current block), in order to determine a PMF for use in adapting the variable length encoding of that same frame (e.g. comprising the same symbols $s_n \dots s_{n+w}$ of the same current block).

This does not involve reference to previously encoded symbols or blocks as in the conventional CAVLC or CABAC techniques or such like. The local PMF calculation block 204 thus looks ahead to forthcoming symbols yet to be encoded within the current frame or other current time interval (and only within that frame or interval), rather than looking back at previously encoded symbols.

The local PMF calculation block 204 thus empirically observes the actual *a posteriori* number of occurrences of each symbol within the frame or interval about to be encoded by the entropy encoder 206, as opposed to relying on a pre-

trained PMF which only describes how frequently the symbols are expected to occur based on a pre-trained simulation or test. The local PMF calculation block 204 then outputs a control signal to the entropy encoder 206 providing information on the observed PMF, which the entropy encoder 206 then uses to adapt the variable length encoding accordingly.

For example, the local PMF calculation block 204 can dynamically determine an ad-hoc, bespoke VLC look-up table mapping of symbol values v to codes c , populated for the frame based on the frequency at which the different symbol values v are found to actually occur within that same frame.

The PMF itself is also inserted into the output bitstream as side information. The PMF side information can be encoded efficiently using a number of methods. The main preferred principles are explained below, any or all of which can be used to encode the PMF side information.

- (i) $\sum_i (p_i)$ is known to be 1 (can be a different known value in computer representation). That is, the total number of symbols must add up to a certain number determined by the size of the frame or interval.
- (ii) By encoding first the probabilities of the most-probable symbol values, the probabilities of less-probable symbol values become more predictable (often zeros in a sparse PMF).
- (iii) The sparse PMF can be coded efficiently using run-length-coding techniques.
- (iv) The PMF of the short interval can be coded differentially with respect to the PMF of previous short interval(s) and/or the global PMF (when available).

- (v) The PMF of the short interval can be coded with reduced number of bits and reduced precision (i.e. encoding in manner that loses some information about the local PMF), since it is often anyway more efficient for the interval than the global PMF. This would reduce the bitrate incurred by the side information at a cost of only a small increase of rate for the symbols. It is important to note that then the reduced-precision PMF has to be used to both encode and decode the symbol values.

Figure 2b shows an example bitstream as output by the encoder of Figure 2a.

- 10 The bitstream comprises a plurality of frames, and each frame comprises a plurality of codes c plus compressed information about the local observed PMF for that same respective frame. The local PMF is transmitted as side information, i.e. it is a part of the bitstream.

- 15 As discussed, conventional global adaptive PMF techniques such as CAVLC and CABAC do not transmit the PMF as side information because the decoder determines the VLC adaption based on previously decoded symbols, in the same manner as the encoder. Although the present invention does require information about the PMF to be included in the output bitstream, the reduced average number of bits per symbol achieved by the local *a posteriori* adaption outweighs the extra bits required to transmit the PMF as side information.

- Figure 2c shows an exemplary decoder according to embodiments of the present invention. This comprises an entropy decoder 216 arranged to receive the encoded symbols and the PMF side information of each frame from the encoded bitstream. For each frame, the entropy decoder 216 uses the received local PMF for the respective frame to determine the relative frequency of occurrence of the symbols within that frame, and therefore to derive the adapted manner in which they were encoded by the entropy encoder 206 (e.g. to reconstruct the bespoke VLC look-up table used to encode the symbols). The entropy decoder 216 then decodes the symbols accordingly and outputs the decoded symbol values to the

later stages 112 of the decoder to produce an output signal for output to a user. E.g. a video signal for display on a video screen, and/or speech or other audio signal for playing out of a speaker.

5 As illustrated in Figure 2a, the described technique may optionally be augmented by comparing the efficiency of the highly-adaptive PMF with that of the global PMF. That is, comparing $R(\text{symbols, local}) + R(\text{side info})$ with $R(\text{symbols, global})$ for each frame, and choosing the best option for the frame. To this end, the encoder comprises a switch 208 configured to switch between the local adaption
10 204 and global adaption 104, making the switching decision once per frame in dependence on the comparison. The choice is indicated by a flag f supplied to the decoder as additional side information in the bitstream for each frame (the additional side information to transmit this flag is negligible). This flag f is shown in Figure 2b. Alternatively the switching decision could be performed once per
15 number of frames instead of necessarily being decided every frame.

In preferred applications of the present invention, the input signal to the encoder comprises at least one of a video signal (e.g. from a webcam), a speech signal (e.g. from a microphone), or another type of audio signal. In the case of a video
20 signal, each frame represents the video image at a different instant in time. In the case of a speech signal, each frame represents a voice interval such as a predetermined number of pitch periods in the case of vocative sounds (pitch period is a term of art in speech encoding, referring to the period of a quasi-periodic waveform modelling an underlying pitch of the voice box). In the case of
25 another type of audio signal, the frame is a predetermined time interval determined according to some particular standard such as MPEG.

Also in preferred applications, the present invention may be used to encode and decode live voice or video over IP (VoIP) calls conducted over a packet-based
30 network such as the Internet. Furthermore, the call may be established using a peer-to-peer (P2P) topology.

It will be appreciated that the above embodiments have been described only by way of example.

5 For instance, the principle of entropy encoding can apply to encoding the symbols in any bit stream based on the relative frequency of occurrence of those symbols within the bit stream, and the present invention is not limited to video encoding, speech encoding or audio encoding, nor any particular standards for such encoding. Further, the present invention is not limited to VoIP or P2P.

10

Further, to encode the PMF into the output bitstream, it is not necessarily required to transmit the absolute probabilities of the different symbol values. In another embodiment for example, the encoded PMF may comprise an indication of the relative probabilities within the interval, e.g. the relative order from most to
15 least likely such as an indication that v_c is the most likely, v_a is the next most likely, etc., without actually enumerating the probabilities in the side information. So for example the ordered list could be used by the decoder to reconstruct the appropriate VLC look-up table. The invention is not limited to any particular form of PMF, as long as it provides empirical information about the relative
20 frequencies at which the different symbol values are found to occur within the local time interval in question.

Furthermore, the invention is not limited to adaptation performed on the basis of any particular video, audio or speech codec frame. Generally speaking, the PMF
25 can be observed and encoded in relation to a plurality of frames, or indeed any portion of the input signal whether it be a discrete time window or otherwise. For example, to achieve further local adaptation it would be possible to transmit two or more local PMFs per time frame, e.g., one PMF for upper part of a video frame image and one PMF for lower part of the image, or one PMF for odd row of pixels
30 and one PMF for even row of pixels. Transmitting more PMFs increases the rate required for side information, but this may be compensated by lower rate of

encoding symbol values.

If the local PMF is transmitted in terms of number of occurrences (counts) of each symbol value, then the local PMF may be transmitted prior to
5 encoding/decoding symbol values. In a further embodiments, after encoding/decoding each symbol value, the local PMF may be further updated at both the encoder and the decoder by subtracting 1 from the count of this encoded/decoded symbol. Thus, the local PDF is made even more adaptive for the next symbols.

10

The invention can also apply to other types of variable length encoding scheme such as arithmetic encoding that does not use a one-to-one mapping between symbol values v and respective codes c . In this case, the manner of encoding the overall codeword could be adapted based on observed frequencies.

15

The encoder components 102, 204, 206, 104 and 208; and decoder components 216 and 112 each preferably comprise software modules embodied on a computer-readable storage medium such as hard drive or flash memory, arranged for execution on one or more processors. However, in alternative
20 embodiments some of all of these components may be implemented at least partially in dedicated hardwired circuitry. The encoder may comprise an output means such as a network interface arranged to transmit the output bitstream to the decoder over a packet-based network such as the internet. In this case the bitstream in transit is embodied on a transmission medium of the network
25 equipment, e.g. in the packet queue(s) of one or more routers. Alternatively or additionally, the encoder output means may comprise memory access circuitry arranged to store the bitstream on a storage medium such as a hard drive or flash memory. The decoder would comprise a corresponding network interface for receiving the bitstream over the network, and/or memory access circuitry for
30 retrieving the bitstream from the storage medium.

Other variants or applications of the invention may be apparent to a person skilled in the art given the disclosure herein. The scope of the invention is not limited by the described embodiments, but only by the appendant claims.

Claims:

1. An encoder stage for encoding an input signal comprising a plurality of different symbols, the encoder stage comprising:
 - a variable length encoder for encoding the input signal; and
 - a counter configured to dynamically detect an observed frequency at which the different symbols are found to occur within each of a plurality of predetermined portions of the input signal, prior to the symbols of each respective portion being encoded by the variable length encoder;wherein the variable length encoder is configured to encode the symbols of each portion using variable length coding performed in dependence on the observed frequencies detected within the respective portion of the input signal, to generate an encoded bitstream comprising the encoded symbols along with an additional element indicating information regarding the observed frequencies detected for each portion, and to output the encoded bitstream to at least one of a storage medium and a transmission medium for supply to a decoder.
2. The encoder of claim 1, wherein the encoder is configured to encode the symbols of said portion without reference to any symbols already encoded by the variable length encoder.
3. The encoder of claim 1 or 2, wherein the variable length encoder is arranged to encode more frequent symbols of the input signal with codes having fewer bits, and to encode less frequent symbols of the input signal with codes having more bits; and
 - wherein the variable length encoder is configured to map the different length codes to the symbols based on the observed frequencies detected within said portion of the input signal, and encode the symbols of said portion of the input signal using said mapping.
4. The encoder stage of any preceding claim, wherein the variable length

encoder is configured to switch between (i) a first mode of operation in which the symbols of each portion are encoded in dependence on the observed frequencies detected for the respective portion, and (ii) a second mode of operation in which the variable length encoding is either non-adaptive or adapted based on previously encoded symbols.

5. The encoder stage of claim 4, wherein the variable length encoder is configured to determine whether to switch between said modes of operation once per number of said portions, in dependence on a relative efficiency metric of the modes for the respective number of portions.

6. The encoder stage of claim 5, wherein the variable length encoder is configured to determine whether to switch between said modes of operation once per portion, in dependence on a relative efficiency metric of the modes for the respective portion.

7. The encoder stage of any preceding claim, wherein the variable length encoder is configured to encode said information on the observed frequencies for inclusion in the output bitstream, based on a known constraint on a total number of occurrences within each portion.

8. The encoder stage of any preceding claim, wherein the variable length encoder is configured to encode said information on the observed frequencies for inclusion in the output bitstream, by predicting the occurrences of less frequent symbols based on the occurrences of more frequent symbols.

9. The encoder stage of any preceding claim, wherein the variable length encoder is configured to encode said information on the observed frequencies for inclusion in the output bitstream, using run-length coding.

10. The encoder stage of any preceding claim, wherein the variable length

encoder is configured to encode said information on the observed frequencies for inclusion in the output bitstream, by differentially encoding the information of one portion relative to one of: the information of a preceding portion, and a global PMF.

11. The encoder stage of any preceding claim, wherein the variable length encoder is configured to encode said information on the observed frequencies for inclusion in the output bitstream, using a non-lossless encoding.

12. The encoder stage of any preceding claim, wherein the input signal comprises a video signal, and each of a plurality of said portions is a video frame representing a different respective moment in time.

13. The encoder stage of any preceding claim, wherein the input is an audio signal, and each of a plurality of said portions is an audio frame.

14. The encoder stage of claim 13, wherein said audio frames are MPEG audio frames.

15. The encoder stage of any preceding claim, wherein the input signal comprises a speech signal, and each of a plurality of said portions is a speech frame representing a predetermined number of voice pitch periods.

16. The encoder stage of any preceding claim, wherein:
the counter is configured to update the observed frequencies for a portion of the input signal after encoding each of a plurality of successive ones of said symbols within the respective portion; and
the variable length encoder is configured to encode each successive symbol of the respective portion using variable length coding performed in dependence on the updated frequencies.

17. A bitstream arranged into a plurality of portions, each comprising:
a plurality of different encoded symbols, encoded according to an adaptive variable length coding scheme, and
an additional element indicating information regarding an observed frequency at which the different symbols occur within the respective portion;
wherein the symbols of each portion are encoded according to a variable length encoding scheme adapted in dependence on the observed frequencies of the different symbols within the respective portion, said information being for use by a decoder in decoding each portion.
18. The bitstream of claim 17, encoded according to any of claims 2 to 16.
19. A network equipment comprising a transmission medium conveying the bitstream of claim 17 or 18.
20. A decoder stage, comprising:
a variable length decoder for decoding a bitstream arranged into a plurality of portions, and each portion comprising a plurality of different encoded symbols and an additional element indicating information regarding an observed frequency at which the different symbols occur within the respective portion, wherein the symbols of each portion are encoded according to a variable length encoding scheme adapted in dependence on the observed frequencies of the different symbols within the respective portion; and
wherein the variable length decoder is configured to decode the symbols of each portion in dependence on the received information regarding the observed frequencies, and to thereby generate an output signal for output to a user device.
21. A decoder according to claim 20, wherein the variable length decoder is configured to decode a bitstream according to any of claims 2 to 16.

Figure 1a

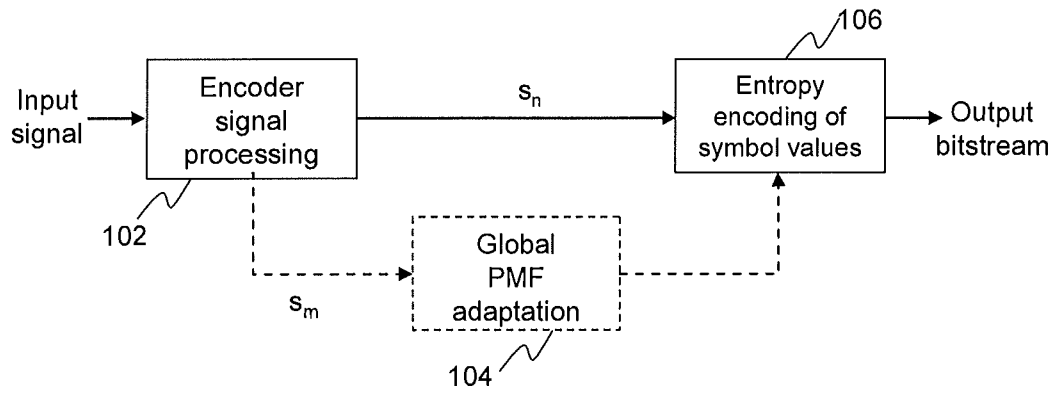


Figure 1b

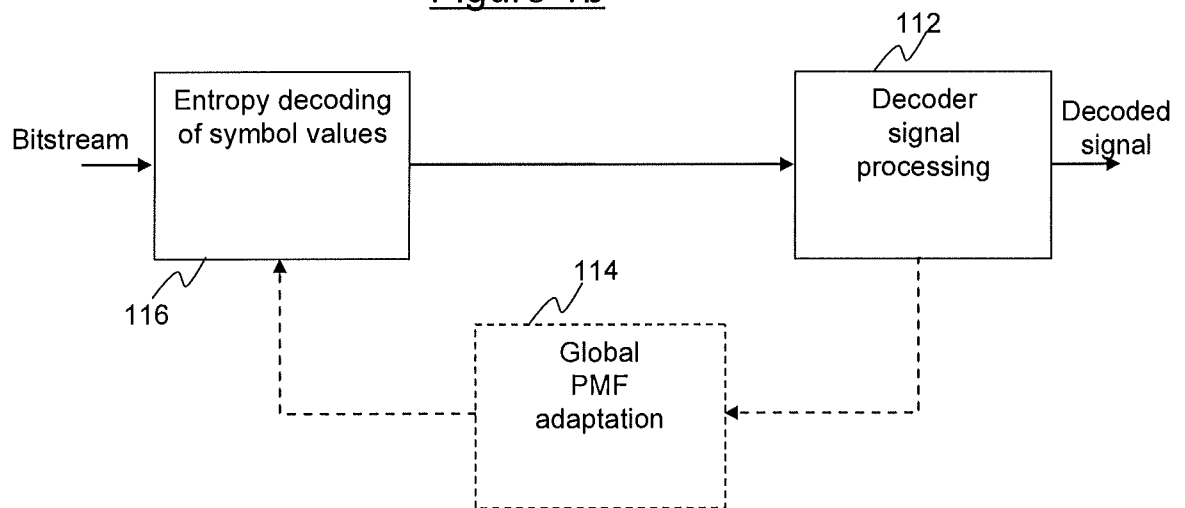


Figure 2b




Figure 2a

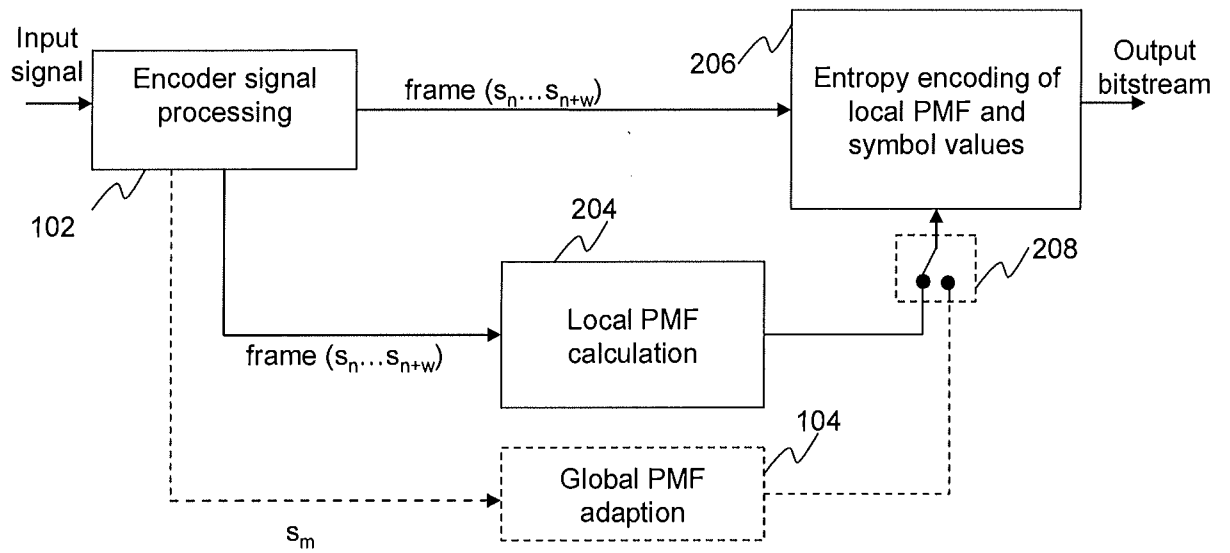



Figure 2c



INTERNATIONAL SEARCH REPORT

International application No
PCT/EP2011/052449

A. CLASSIFICATION OF SUBJECT MATTER
 INV. H03M7/40
 ADD.
 According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED
 Minimum documentation searched (classification system followed by classification symbols)
 H03M

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practical, search terms used)
 EPO-Internal

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	VITTER J S: "DESIGN AND ANALYSIS OF DYNAMIC HUFFMAN CODES", JOURNAL OF THE ASSOCIATION FOR COMPUTING MACHINERY, ACM, NEW YORK, NY, US, vol. 34, no. 4, 1 October 1987 (1987-10-01), pages 825-845, XP000912265, ISSN: 0004-5411, DOI: DOI:10.1145/31846.42227 page 826, line 13 - line 17 -----	1-21

Further documents are listed in the continuation of Box C.

See patent family annex.

* Special categories of cited documents :

- "A" document defining the general state of the art which is not considered to be of particular relevance
- "E" earlier document but published on or after the international filing date
- "L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)
- "O" document referring to an oral disclosure, use, exhibition or other means
- "P" document published prior to the international filing date but later than the priority date claimed

- "T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
- "X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
- "Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.
- "&" document member of the same patent family

Date of the actual completion of the international search
 13 April 2011

Date of mailing of the international search report
 20/04/2011

Name and mailing address of the ISA/
 European Patent Office, P.B. 5818 Patentlaan 2
 NL - 2280 HV Rijswijk
 Tel. (+31-70) 340-2040,
 Fax: (+31-70) 340-3016

Authorized officer
 Nowbakht Irani, Ali