




(86) **Date de dépôt PCT/PCT Filing Date:** 2007/10/18
 (87) **Date publication PCT/PCT Publication Date:** 2008/04/24
 (45) **Date de délivrance/Issue Date:** 2015/10/13
 (85) **Entrée phase nationale/National Entry:** 2009/09/30
 (86) **N° demande PCT/PCT Application No.:** CA 2007/001821
 (87) **N° publication PCT/PCT Publication No.:** 2008/046203
 (30) **Priorité/Priority:** 2006/10/18 (US60/862,029)

(51) **Cl.Int./Int.Cl. H04L 9/32** (2006.01)
 (72) **Inventeurs/Inventors:**
 VESTERGAARD, STEVEN ERIK, CA;
 TSUI, CHE-WAI, CA
 (73) **Propriétaire/Owner:**
 DESTINY SOFTWARE PRODUCTIONS INC., CA
 (74) **Agent:** OYEN WIGGS GREEN & MUTALA LLP

(54) **Titre : FILIGRANAGE DE DONNEES DE MEDIA**
 (54) **Title: METHODS FOR WATERMARKING MEDIA DATA**



(57) **Abrégé/Abstract:**

Methods are provided for encoding watermark information into media data containing a series of digital samples in a sample domain. The method involves: dividing the series of digital samples into a plurality of sections in the sample domain, each section

(57) Abrégé(suite)/Abstract(continued):

comprising a corresponding plurality of samples; processing the corresponding plurality of samples in each section to obtain a single energy value associated with each section; grouping the sections into groups, each group containing three or more sections; assigning a nominal bit value to each group according to a bit assignment rule, the bit assignment rule based on the energy values of the sections in the group; and assigning a watermark bit value to each group. The methods also involve, for each group, comparing the watermark bit value to the nominal bit value and, if the nominal bit value and the watermark bit value of the watermark information bit do not match, modifying one or more energy values of one or more corresponding sections in the group such that re-application of the bit assignment rule would assign the watermark bit value to the group. The bit assignment rule may comprise: a categorization rule for categorizing each group into one of a plurality of categories; and for each category, a unique category bit assignment rule for assigning a nominal zero bit value or a nominal one bit value to each group.

(12) INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
24 April 2008 (24.04.2008)

PCT

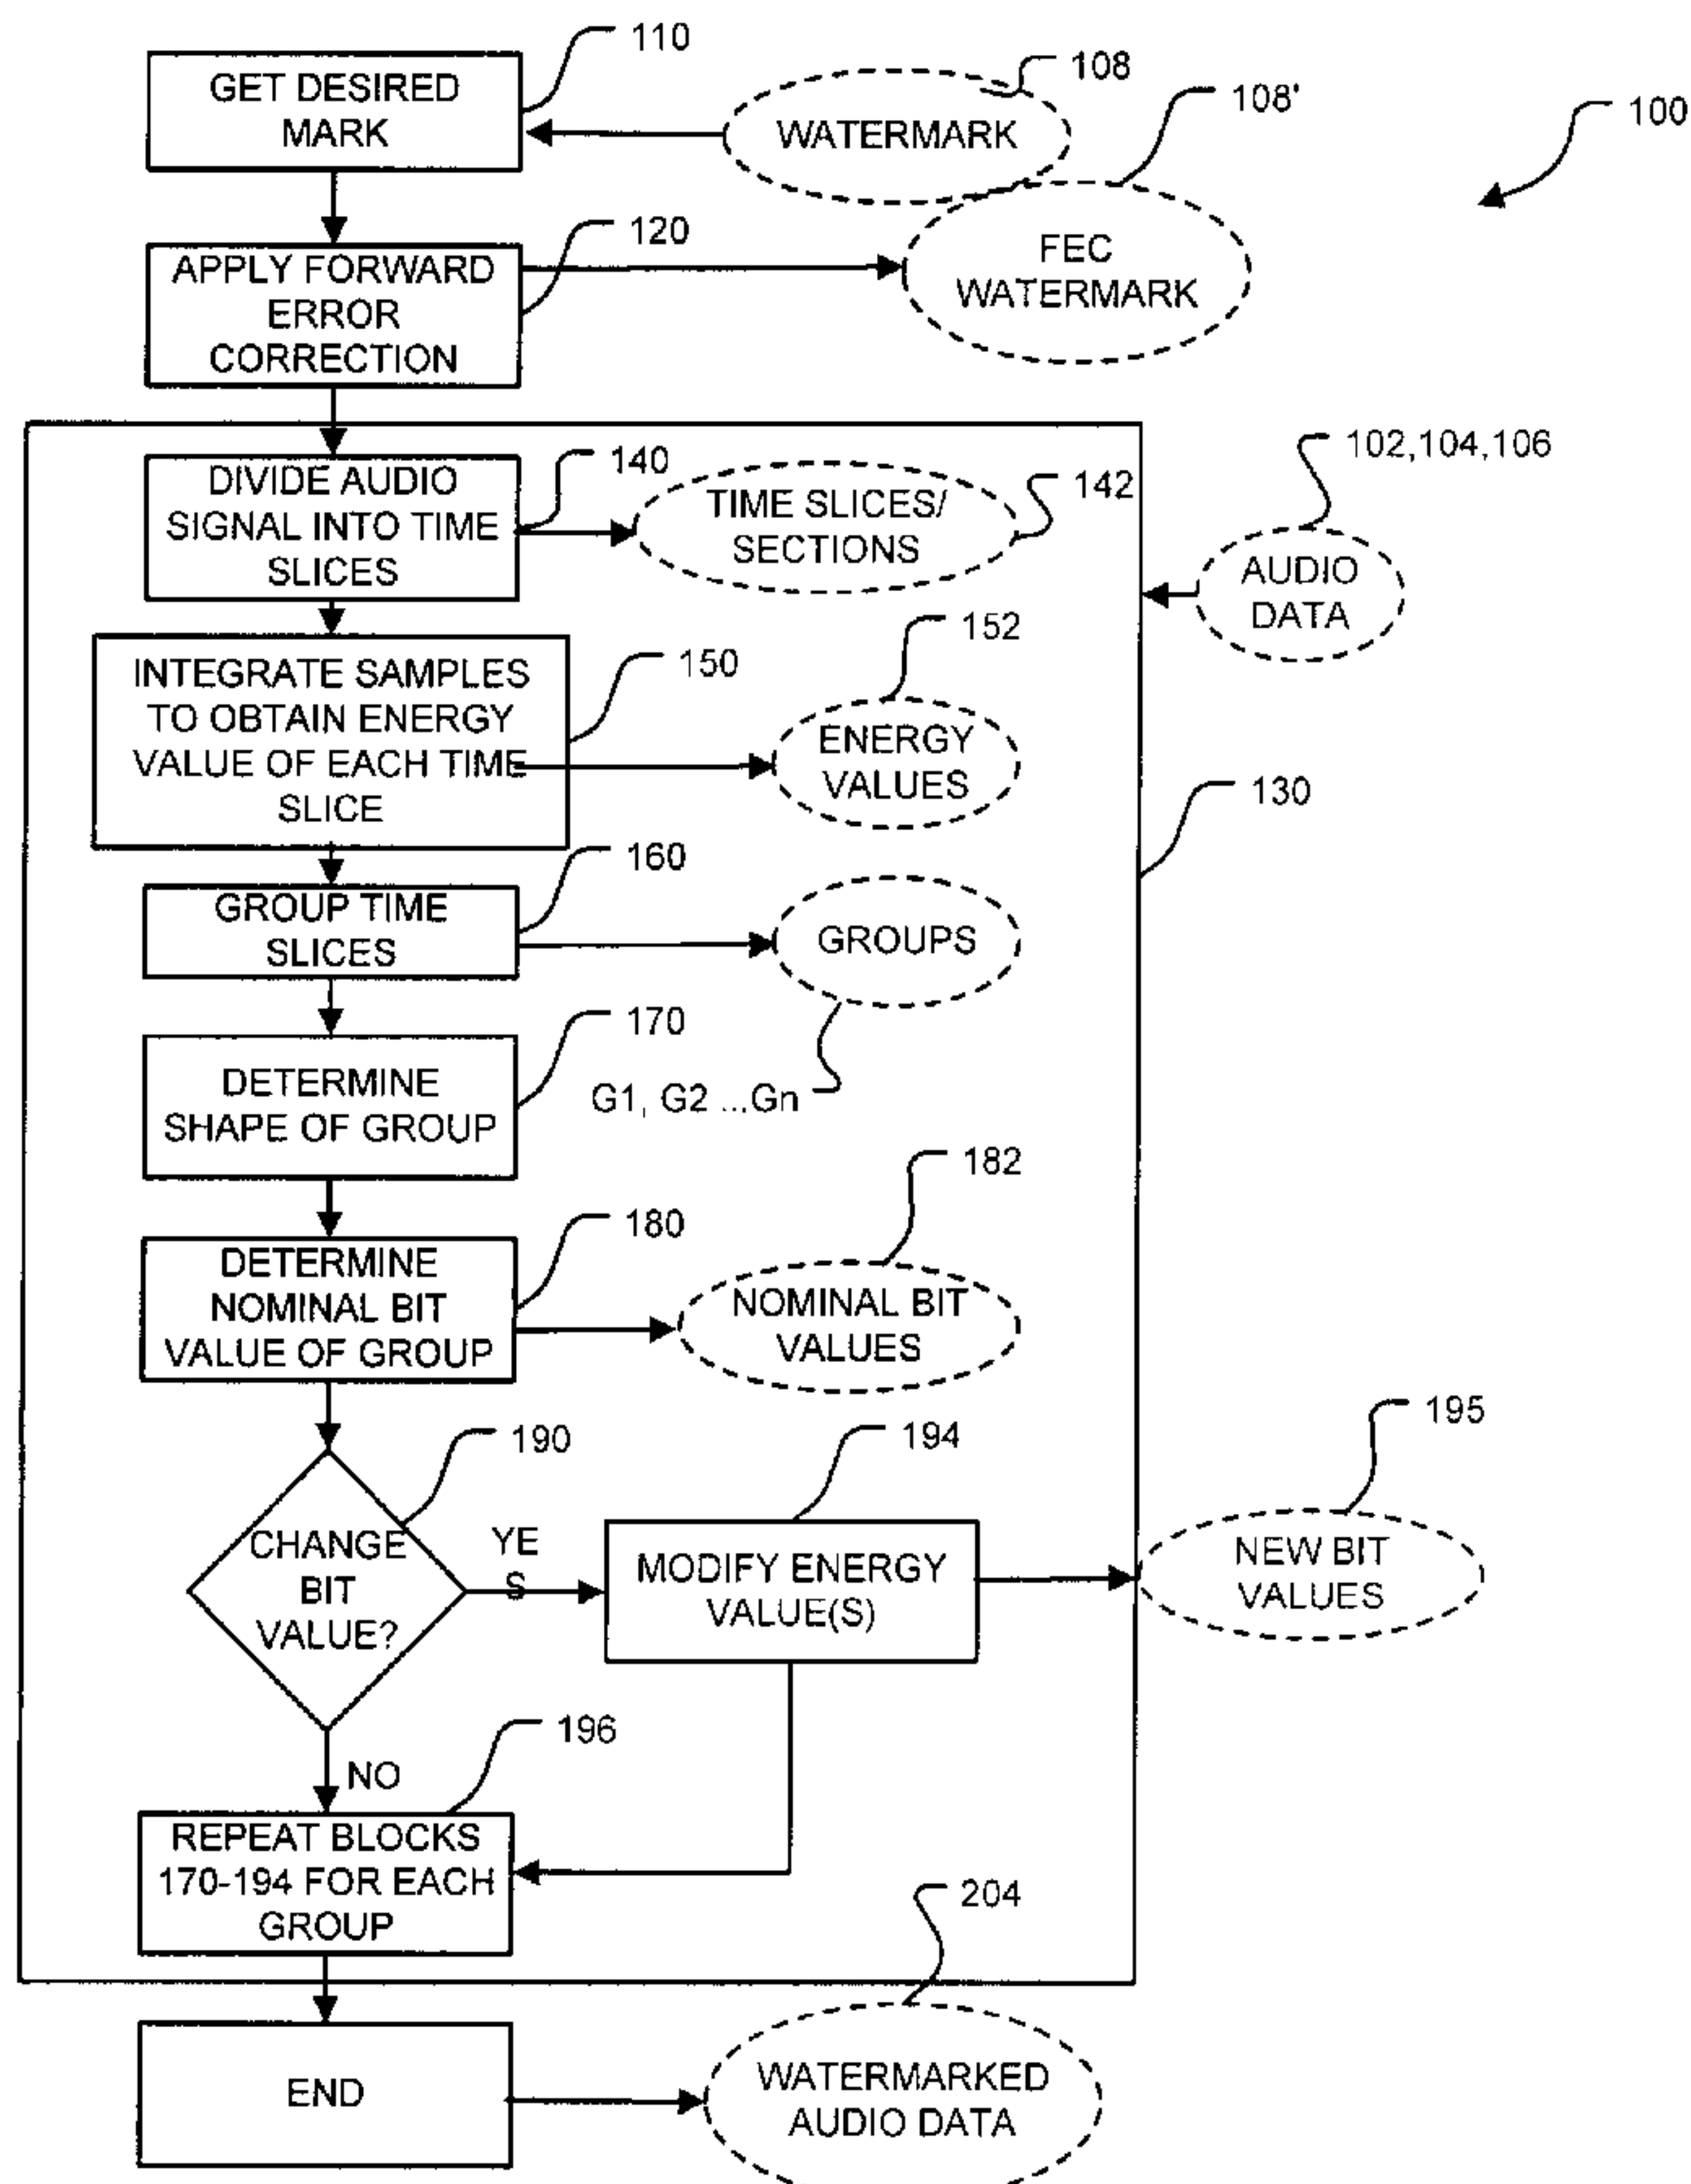
(10) International Publication Number
WO 2008/046203 A1

- (51) International Patent Classification:
H04L 9/32 (2006.01)
- (21) International Application Number:
PCT/CA2007/001821
- (22) International Filing Date: 18 October 2007 (18.10.2007)
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data:
60/862,029 18 October 2006 (18.10.2006) US
- (71) Applicant (for all designated States except US): **DESTINY SOFTWARE PRODUCTIONS INC.** [CA/CA];
800 - 570 Granville Street, Vancouver, British Columbia V6C 3P1 (CA).
- (72) Inventors; and
- (75) Inventors/Applicants (for US only): **VESTERGAARD, Steven Erik** [CA/CA]; c/o 800 - 570 Granville Street, Vancouver, British Columbia V6C 3P1 (CA). **TSUI, Che-Wai** [CA/CA]; 1398 El Camino Drive, Coquitlam, British Columbia V3E 2W6 (CA).

- (74) Agents: **RATTRAY, Todd** et al.; 480-The Station, 601 West Cordova Street, Vancouver, British Columbia V6B 1G1 (CA).
- (81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BH, BR, BW, BY, BZ, CA, CH, 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, PG, PH, PL, PT, RO, RS, RU, SC, SD, SE, SG, SK, SL, SM, SV, SY, 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, LS, MW, MZ, NA, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HU, IE, IS, IT, LT, LU, LV, MC, MT, NL, PL,

[Continued on next page]

(54) Title: METHODS FOR WATERMARKING MEDIA DATA



(57) Abstract: Methods are provided for encoding watermark information into media data containing a series of digital samples in a sample domain. The method involves: dividing the series of digital samples into a plurality of sections in the sample domain, each section comprising a corresponding plurality of samples; processing the corresponding plurality of samples in each section to obtain a single energy value associated with each section; grouping the sections into groups, each group containing three or more sections; assigning a nominal bit value to each group according to a bit assignment rule, the bit assignment rule based on the energy values of the sections in the group; and assigning a watermark bit value to each group. The methods also involve, for each group, comparing the watermark bit value to the nominal bit value and, if the nominal bit value and the watermark bit value of the watermark information bit do not match, modifying one or more energy values of one or more corresponding sections in the group such that re-application of the bit assignment rule would assign the watermark bit value to the group. The bit assignment rule may comprise: a categorization rule for categorizing each group into one of a plurality of categories; and for each category, a unique category bit assignment rule for assigning a nominal zero bit value or a nominal one bit value to each group.

WO 2008/046203 A1

WO 2008/046203 A1



PT, RO, SE, SI, SK, TR), OAPI (BF, BJ, CF, CG, CI, CM,
GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

— *before the expiration of the time limit for amending the
claims and to be republished in the event of receipt of
amendments*

Published:

— *with international search report*

METHODS FOR WATERMARKING MEDIA DATA

5

Technical Field

[0002] The invention relates to watermarking media data. Particular embodiments of the invention provide methods for embedding watermarks into media data and/or extracting watermarks from media data.

10

Background

[0003] Watermarking media data (e.g. media files and/or media signals) generally involves embedding or otherwise adding information to the media data (i.e. "marking" the media data). In many applications, it is preferable that the information embedded in the media data does not substantially alter the human perception of playback of the media content. A watermark may generally comprise any information, such as, for example: encryption keys, authentication information, indications of ownership, indications of user ID or the like.

15

20

[0004] In the example of audio media data, watermarking processes typically exploit the human audio perception's inability to detect slight changes in the playback of audio content. When a watermark is embedded into audio data, it is generally desirable that when the audio data is played back, it is difficult for a human to perceive differences between playback of the watermarked audio content and playback of the original non-watermarked audio content.

25

[0005] It is generally desirable that, once embedded, a watermark be relatively robust to manipulation of the resultant watermarked media data. By way of non-limiting examples, it is desirable that the watermark be recoverable after: compression of the watermarked

30

- 2 -

data; addition of noise to the watermarked data; D/A and A/D conversion of the watermarked data; resampling of the watermarked data; filtering of the watermarked data; changing the volume of the watermarked data; and/or other manipulation of the watermarked data.

5

[0006] It is also generally desirable that it be relatively difficult to filter or otherwise remove the watermarking from the watermarked media data without significantly degrading the playback quality of the media content.

10 [0007] Prior art watermarking techniques have yet to achieve some or all of the aforementioned desirable characteristics to a level sufficient to gain widespread commercial acceptance.

- 3 -

Summary

[0008] One aspect of the present invention provides a method for encoding watermark information into media data containing a series of digital samples in a sample domain. The method comprises: dividing the series of digital samples into a plurality of sections in the sample domain, each section comprising a corresponding plurality of samples; processing the corresponding plurality of samples in each section to obtain a single energy value associated with each section; grouping the sections into groups, each group containing three or more sections; assigning a nominal bit value to each group according to a bit assignment rule, the bit assignment rule based on the energy values of the sections in the group; assigning a watermark bit value to each group; and for each group, comparing the watermark bit value to the nominal bit value and, if the nominal bit value and the watermark bit value of the watermark information bit do not match, modifying one or more energy values of one or more corresponding sections in the group such that re-application of the bit assignment rule would assign the watermark bit value to the group.

[0009] Assigning the nominal bit value to each group according to the bit assignment rule may comprise: categorizing each group into one of a plurality of categories according to a categorization rule; and, for each category, assigning a nominal zero bit value or a nominal one bit value to each group according to a unique category bit assignment rule associated with the category.

[0010] Another aspect of the invention provides a method for extracting watermark information from media data containing a series of digital samples in a sample domain. The method comprises: dividing the series of digital samples into a plurality of sections in the sample domain, each section comprising a corresponding plurality of samples; processing the corresponding plurality of samples in each section to obtain a single energy value associated with each section; grouping the sections into groups, each group

- 4 -

containing three or more sections; assigning a nominal bit value to each group according to a bit assignment rule, the bit assignment rule based on the energy values of the sections in the group. Assigning the nominal bit value to each group according to the bit assignment rule may comprise: categorizing each group into one of a plurality of
5 categories according to a categorization rule; and, for each category, assigning a nominal zero bit value or a nominal one bit value to each group according to a unique category bit assignment rule associated with the category.

10 [0011] Other aspects and features of specific embodiments of the invention are described in more detail below.

Brief Description of the Drawings

[0012] In drawings which depict non-limiting embodiments of the invention:

15 Figure 1 is a schematic depiction of digital audio data which may be watermarked in accordance with the invention;

Figure 2 depicts a method for embedding watermark data into a media file in accordance with a particular embodiment of the invention;

20 Figure 3 shows how media data is divided into sections/time slices containing a plurality of samples and how the sections/time slices are grouped into groups;

Figure 4 depicts a number of shape categories into which the groups G may be categorized according to a particular embodiment of the invention;

25 Figure 5A-5C schematically depicts a number of methods which may be used to modify one or more energy values of a group G to change the bit value that would be assigned to the group G on re-application of the bit assignment rule;

Figure 6 shows a method for extracting a watermark from watermarked media data according to a particular embodiment of the invention; and

- 5 -

Figure 7 is a schematic depiction of digital image data which may be watermarked in accordance with the invention.

Figure 8 depicts a method for energy value modification according to a particular embodiment of the invention.

Detailed Description

[0013] Throughout the following description, specific details are set forth in order to provide a more thorough understanding of the invention. However, the invention may be practiced without these particulars. In other instances, well known elements have not been shown or described in detail to avoid unnecessarily obscuring the invention. Accordingly, the specification and drawings are to be regarded in an illustrative, rather than a restrictive, sense.

[0014] Particular embodiments of the invention provide methods for watermarking media data which involve minor modification of the relative energy the media data samples. The watermarking techniques of the invention may generally be applied to a variety of types of media data, including, by way of non-limiting example, audio data, image data and video data. For the purposes of illustrating particular exemplary embodiments of the invention, the following description describes methods for watermarking audio data and then describes how these techniques can be modified for watermarking other types of media data.

Audio Data

[0015] Figure 1 schematically depicts a segment of typical digital audio data 10 which comprises a series 12 of samples 14 in the time domain t . Samples 14 may comprise pulse-code modulation (PCM) samples, for example. This is not necessary, however, and the invention could be practiced using other sampling schemes. Samples 14 are spaced from one another by a sampling period T_s , which defines a corresponding sampling frequency f_s . In the particular case of CD audio, the sampling frequency f_s is typically 44.1 kHz, although the invention should be understood to incorporate digital audio data 10

- 6 -

having any suitable sampling frequency f_s . Samples 14 may be digitally quantized in a binary scheme. By way of non-limiting example, CD audio currently uses a 16-bit resolution for each sample 14.

5 [0016] Audio data 10 may be maintained in a suitable file (not shown), which facilitates storage, transmission, processing and the like of audio data 10. Those familiar with audio technology will appreciate that digital audio data 10 of Figure 1 may represent one audio channel and that a plurality of channels may be combined (e.g. by time division multiplexing (TDM) or any other suitable technique) to form multi-channel audio data
10 10. In such cases, it will be assumed for the purposes of the description below that the sampling period T_s and the sampling frequency f_s are characteristics of the multi-channel audio data 10. In such cases, the watermarking methods described below may be performed on multi-channel audio data 10. In other embodiments, the watermarking methods described below could be applied separately to each channel and then the
15 channels may be subsequently multiplexed.

Embedding a watermark

[0017] Figure 2 schematically depicts a method 100 for embedding a watermark 108 into media data 102 in accordance with a particular embodiment of the invention. In the
20 exemplary embodiment of Figure 2, media data 102 comprises audio data 104. Audio data 104 may be provided in the form of an audio file 106. Audio data 104 may have the characteristics described above, including a sampling period T_s and a sampling frequency f_s . Audio data 104 may comprise multi-channel audio data or single channel audio data. In some embodiments, audio file 106 may comprise additional data (e.g. metadata) in
25 addition to audio data 104.

[0018] Method 100 commences in block 110 which involves obtaining the desired watermark data 108. In one particular embodiment, watermark data 108 comprises 32 bits

- 7 -

of information, although watermark data 108 may generally comprise a different number of bits. Watermark data 108 could generally comprise any data. By way of non-limiting example, watermark data 108 may comprise user ID information, encryption keys, authentication information designating ownership of media data 102 or the like. In some
5 embodiments, watermark data 108 could be encrypted (e.g. using one or more encryption keys and/or pseudo randomization processes). A variety of suitable encryption techniques are known in the art.

[0019] Method 100 then proceeds to block 120, where watermark data 108 is expanded
10 using one or more forward error correction (FEC) encoding schemes to provide FEC watermark data 108'. In some embodiments, the block 120 forward error correction encoding scheme is based on Low-Density Parity-Check (LDPC) codes, although other error correction encoding techniques may additionally or alternatively be used. In one particular embodiment of the invention, block 120 involves expanding the 32 bit
15 watermark data 108 using forward LDPC error correction to FEC watermark data 108' of 600 bits - i.e. a redundancy factor of more than 15 times. As discussed further below, a redundancy factor of this magnitude permits watermark data 108 to be recovered even if up to 1/3 of the bits of FEC watermark data 108' are corrupted.

20 [0020] Method 100 then proceeds to block 130, where FEC watermark data 108' is embedded into audio data 104.

[0021] In the illustrated embodiment, the block 130 watermark embedding procedure starts in block 140, where time-domain digital audio data 104 is parsed into a number of
25 sections 142. This block 140 sectioning process is shown schematically in Figure 3 for a portion of audio data 104, where sections 142 are demarcated by lines 144. In the exemplary embodiment, where media data 102 comprises audio data 104 and audio data 104 comprises a series of samples in the time domain (see Figure 1), the block 140

- 8 -

sections 142 may be referred to as time slices 142. Each time slice 142 may comprise a plurality of samples of audio data 104. As shown in Figure 3, sections 142 may have equal numbers of samples - e.g. time slices 142 may have equal durations. For example, in some embodiments, each time slice 142 may have duration $t_{section}$ of ~5-100 ms and
5 may comprise ~10-1000 samples. The duration of time slices 142 and/or the number of samples in each time slice 142 may vary as between implementations of method 100. For example, where audio data 104 comprises a short audio clip, it may be desirable to have relatively small time slices 142 to ensure that a sufficient number of the bits of FEC watermark 108' can be embedded into audio data 104. Conversely, if audio data 104 is
10 two hours long, time slices 142 may be relatively large while still enabling a sufficient number of time slices 142 to encode a desired number of bits of FEC watermark 108'.

[0022] In block 150, the samples in each time slice 142 are integrated to obtain a single value 152 for each time slice 142. In one particular embodiment, the block 150
15 integration comprises adding the values of the samples of audio data 104 in each corresponding time slice 142. The value 152 obtained from the block 150 integration may represent the energy associated with audio data 104 for each associated time slice 142 and may be referred to herein as the energy value 152 of the associated time slice 142. Since time slices 142 are the same length, each time slice 142 will have the same number of
20 samples and each energy value 152 will be representative of the total magnitude of the samples in its corresponding time slice 142.

[0023] Block 160 involves assembling time slices 142 and their corresponding energy values 152 into groups. This block 150 grouping process is also shown schematically in
25 Figure 3, which depicts groups G_1 and G_2 . Depending on the length of audio data 104, the block 150 groups may comprise groups $G_1, G_2 \dots G_n$ (referred to individually and collectively as groups G). In one particular embodiment, each group G of time slices 142 comprises three time slices 142 and three corresponding energy values 152.

- 9 -

[0024] Method 100 then proceeds to block 170 which involves categorizing the shapes of the energy values 152 associated with each group G of time slices 142. In the illustrated embodiment, each group G of time slices 142 comprises an ordered triplet of three energy values 152. The ordered triplet of energy values 152 comprises:

- a first energy value 152A associated with a first time slice 142A;
- a second energy value 152B associated with a second time slice 142B, wherein the second time slice 142B is adjacent to the first time slice 142A; and
- a third energy value 152C associated with a third time slice 142C, wherein the third time slice 142C is adjacent to the second time slice 142B and is spaced apart from the first time slice 142A.

Block 170 may involve categorizing each group G of time slices 142 (and its corresponding energy values 152) into one of the shape categories shown in Figure 4.

[0025] Figure 4 show a number of shape categories and a number of examples of groups G which fit into each category. The Figure 4 shape categories include:

- Shape 1: where second energy value 152B is greater than first energy value 152A and third energy value 152C is greater than second energy value 152B;
- Shape 2: where the second energy value 152B is less than the first energy value 152A and the third energy value 152C is less than the second energy value 152B;
- Shape 3: where the second energy value 152B is greater than the first energy value 152A and the third energy value 152C is between the first and second energy values 152A, 152B;
- Shape 4: where the second energy value 152B is greater than the first energy value 152A and the third energy value 152C is less than the first energy value 152A;
- Shape 5: where the second energy 152B value is less than the first energy value 152A and the third energy value 152C is between the first and second energy values 152A, 152B;

- 10 -

- Shape 6: where the second energy value 152B is less than the first energy value 152A and the third energy value 152C is greater than the first energy value 152A; and
- Shape 7: the rare case that the first, second and third energy values 152A, 152B, 152C are the same.

5
10
15
[0026] The seven shape categories shown in Figure 4 represent a range of possible shape categories where there are three time slices 142 and three corresponding energy values 152 in each group G of time slices 142. In some embodiments, the number of time slices 142 in each group G can differ. With a different number of time slices 142 in each group G, the range of possible shape categories and the rules for assigning shape categories may also differ. However, those skilled in the art will appreciate that even with a different number of time slices 142 in each group G, it is possible to create shape categories based on the relative energy values 152 of the time slices 142.

15
20
[0027] Returning to Figure 2, method 100 proceeds to block 180 which involves determining a nominal bit value 182 for each group G of time slices 142. A nominal bit value 182 is assigned to each of the block 170 shape categories according to a bit assignment rule that may be particular to that shape category. The bit assignment rules may be based on the energy values 152 and/or the slopes between successive energy values 152A, 152B, 152C and/or functions of energy values 152 and the slopes between successive energy values 152A, 152B, 152C, for example.

25
[0028] Figure 4 schematically depicts a set of bit assignment rules for the Figure 4 shape categories according to one particular embodiment. The set of bit assignment rules associated with the seven shapes shown in Figure 4 is:

- Shape 1: if the magnitude of the slope between the first and second energy values 152A, 152B is greater than or equal to the magnitude of the slope between the

- 11 -

second and third energy values 152B, 152C, then assign a nominal bit value 182 of 0; else if the magnitude of the slope between the first and second energy values 152A, 125B is less than the magnitude of the slope between the second and third energy values 152B, 152C, then assign a nominal bit value 182 of 1;

- 5
- Shape 2: if the magnitude of the slope between the first and second energy values 152A, 152B is less than or equal to the magnitude of the slope between the second and third energy values 152B, 152C, then assign a nominal bit value 182 of 0; else if the magnitude of the slope between the first and second energy values 152A, 152B is greater than the magnitude of the slope between the second and third energy values 152B, 152C, then assign a nominal bit value 182 of 1;
- 10
- Shape 3: if the magnitude of the slope between the second and third energy values 152B, 152C is less than or equal to $\frac{1}{2}$ of the magnitude of the slope between the first and second energy values 152A, 152B, then assign a nominal bit value 182 of 0; else if the magnitude of the slope between the second and third energy values 152B, 152C is greater than $\frac{1}{2}$ of the magnitude of the slope between the first and second energy values 152A, 152B, then assign a nominal bit value 182 of 1;
- 15
- Shape 4: if the magnitude of the slope between the second and third energy values 152B, 152C is less than or equal to 2 times the magnitude of the slope between the first and second energy values 152A, 152B, then assign a nominal bit value 182 of 0; else if the magnitude of the slope between the second and third energy values 152B, 152C is greater than 2 times the magnitude of the slope between the first and second energy values 152A, 152B, then assign a nominal bit value 182 of 1;
- 20
- Shape 5: if the magnitude of the slope between the second and third energy values 152B, 152C is less than or equal to $\frac{1}{2}$ of the magnitude of the slope between the first and second energy values 152A, 152B, then assign a nominal bit value 182 of 1; else if the magnitude of the slope between the second and third energy values
- 25

- 12 -

- 152B, 152C is greater than $\frac{1}{2}$ of the magnitude of the slope between the first and second energy values 152A, 152B, then assign a bit value 182 of 0;
- Shape 6: if the magnitude of the slope between the second and third energy values 152B, 152C is less than or equal to 2 times the magnitude of the slope between the first and second energy values 152A, 152B, then assign a nominal bit value 182 of 1; else if the magnitude of the slope between the second and third energy values 152B, 152C is greater than 2 times the magnitude of the slope between the first and second energy values 152A, 152B, then assign a nominal bit value 182 of 0; and
 - Shape 7: assign no nominal bit value 182.

[0029] In the illustrated embodiment, where the separation between the first and second energy values 152A, 152B is the same as the separation between the second and third energy values 152B, 152C, the slopes referred to in the block 180 nominal bit assignment rules may be proportional to the difference between the first and second energy values 152A, 152B and the difference between the second and third energy values 152B, 152C. This is not necessary. In some embodiments, the separation between energy values 152 used in the bit assignment rule may not be equal (e.g. where each group G contains more than three time slices 142). In such cases, slope calculations may involve division by the relative distances between energy values 152.

[0030] The above-described rules assigned to the seven shape categories shown in Figure 4 represent only one possible set of nominal bit assignment rules. Other suitable sets of nominal bit assignment rules may be used in other embodiments of the invention. Such additional or alternative sets of nominal bit assignment rules may depend on the number of time slices 142 in each group G, for example. Such additional or alternative sets of nominal bit assignment rules may use different parameters and/or different functions to assign the nominal bit values 182, for example. Preferably, the nominal bit assignment

- 13 -

rules include the possibility of assigning a nominal bit value 182 of 1 and a nominal bit value 182 of 0 to at least three or more shape categories.

[0031] Returning to Figure 2, block 190 involves an inquiry into whether the block 180
5 nominal bit value 182 for the next group G accords with the next bit value for FEC
watermark 108'. Block 190 may be carried out sequentially for each group G. If a
particular group G has a nominal bit value 182 which corresponds to the bit value of FEC
watermark 108' (block 190 NO output), then the energy values 152 of the particular group
G need not be changed and method 100 proceeds to block 196. If, on the other hand, the
10 particular group G has a nominal bit value 182 of 0 when it is desired to mark the group
with a 1 bit value (i.e. the corresponding bit of FEC watermark 108' is a 1) or the
particular group G has a nominal bit value 182 of 1 when it is desired to mark the group
with a 0 bit value (i.e. the corresponding bit of FEC watermark 108' is a 0), then method
100 changes the nominal bit value 182. This latter situation corresponds to the block 190
15 YES output, where method 100 proceeds to block 194. Block 194 involves changing the
nominal bit value 182 of the particular group G by adjusting one or more of its energy
value(s) 152 in such a manner that re-application of the block 180 bit assignment rule
would assign the group G a new bit value 195 - i.e. the new bit value 195 that would be
assigned to the group G is different than the nominal bit value 182 assigned to the group
20 G. The new bit value 195 corresponds with the FEC watermark 108' bit value that it is
desired to encode in group G.

[0032] The block 194 adjustment of the energy value(s) 152 of one or more time slices
142 may be accomplished using a variety of techniques. A number of example bit value
25 changes and their corresponding energy value 152 change(s) are presented in Figures
5A-5C. The illustrated examples of Figures 5A-5C begin with data that is determined (in
block 170) to have a shape category of shape 1 and determined (in block 180) to have a
nominal bit value of 0. For the examples in Figures 5A-5C, it is desired to adjust one or

- 14 -

more of the energy values 152 so as to encode a FEC watermark 108' bit value of 1 rather than the nominal bit value 182 of 0 - i.e. for the examples in Figures 5A-5C, after the block 194 adjustment of energy values, it is desired that re-application of the block 180 bit assignment rule would assign the group G a new bit value 195 of 1.

5

[0033] In some embodiments, the block 194 adjustment of energy value(s) 152 of one or more time slices 142 may involve adjustment of a nominal bit value 182 of 0 to a new bit value of 1 or a nominal bit value 182 of 1 to a new bit value of 0 within a particular one of the block 170 shape categories. Examples of this technique are shown in Figures 5A and 5B. Figure 5A involves modifying one energy value 152 (second energy value 152B) within the group G. In the particular example of Figure 5A, second energy value 152B is decreased from its original value 152B to a modified value 152B'. It can be seen from Figure 5A and from the above-discussed block 180 bit assignment rule that after reducing second energy value 152B to modified value 152B', re-application of the block 180 bit assignment rule would cause the bit value assigned to the group G to change from its original nominal bit value 182 of 0 to a new bit value 195 of 1. That is, the block 194 energy value modification depicted in Figure 5A causes a change in the bit value that would be assigned by the block 180 bit assignment rule.

20 [0034] Figure 5B shows another technique for the block 194 adjustment of energy values(s) 152 which involves modifying a plurality of energy values 152 (e.g. the first and second energy values 152A, 152B in the illustrated example of Figure 5B). In the particular example of Figure 5B, first energy value 152A is increased from its original value 152A to a modified value 152A' and second energy value 152B is decreased from its original value 152B to a modified value 152B'. Again, it can be seen from Figure 5B and from the above-discussed block 180 bit assignment rule that after increasing the first energy value 152A to modified value 152A' and decreasing the second energy value 152B to modified value 152B', re-application of the block 180 bit assignment rule would cause

25

- 15 -

the bit value assigned to the group G to change from its original nominal bit value 182 of 0 to a new bit value 195 of 1. That is, the block 194 energy value modification depicted in Figure 5B causes a change in the bit value that would be assigned by the block 180 bit assignment rule.

5

[0035] The block 194 adjustment of energy value(s) 152 may additionally or alternatively involve adjustment of a group G's nominal bit value 182 of 0 to a new value 195 of bit 1 or a group G's nominal bit value 182 of 1 to a new bit value 195 of 0 together with a corresponding change in the group G's block 170 shape category. An example of this technique is shown schematically in Figure 5C. In the Figure 5C example, the group G starts in the block 170 shape category 1, with a nominal bit value 182 of 0. Block 194 involves increasing second energy value 152B to modified second energy value 152B' and decreasing third energy value 152C to modified third energy value 152C'. In this manner, group G is changed such that re-application of the block 170 shape categorization process would assign the group G to a new shape category and re-application of the block 180 bit assignment rule would assign the group G a new bit value 195. In the particular example shown in Figure 5C, after adjustment of the second and third energy values 152B, 152C to modified values 152B', 152C', re-application of the block 170 shape categorization process would assign the group G to shape category 3 and re-application of the block 180 bit assignment rule would assign the group G a new bit value 195 of 1.

[0036] Figure 8 schematically depicts a method 400 for the block 194 energy value modification according to a particular embodiment of the invention. Method 400 may involve modifying the energy values 152A, 152B, 152C associated with each of the time slices 142 in a group G. Method 400 is general to any of the block 170 (Figure 4) shape categories and to the situation where it is desired to modify energy values 152 so as to encode a 1 bit or to modify energy values 152 so as to encode a 0 bit.

25

- 16 -

[0037] Method 400 is explained first for the circumstance where the block 170 shape category is shape 1, the block 180 nominal bit value 182 is a bit 0 and it is desired to encode a bit 1. Method 400 starts in block 410 which involves evaluation of a terminal condition. The block 410 terminal condition may be expressed in terms of an error margin δ . The error margin δ may be a percentage $x\%$ of the sum of the energy values 152 of a group G -i.e. $\delta = x(e_0 + e_1 + e_2)/100$ where e_0 is the first energy value 152A, e_1 is the second energy value 152B and e_2 is the third energy value 152C. In some embodiments, the error margin δ may be on the order of 0.2%-10% of the sum of the energy values $e_0 + e_1 + e_2$, for example. The size of the error margin may reflect a tradeoff between the robustness of the watermark and degradation of audio quality.

[0038] The block 410 terminal condition may be similar to the block 170 shape categorization rules. In the particular case of shape category 1, where it is desired to modify a nominal bit value 182 of bit 0 to bit 1, the block 410 terminal condition may be: $(e_2 + e_0)/2 \geq e_1 + \delta$. Typically, if method 100 (Figure 2) has arrived at block 194, then the block 410 terminal condition will not be satisfied on the first iteration of method 400. Consequently, method 400 will proceed to block 420 via the block 410 NO output.

[0039] In block 420 method 400 involves modifying a first energy value 152. In the particular case of shape category 1 where it is desired to modify a nominal bit value 182 of 0 to a bit value of 1, block 420 may involve increasing the third energy value 152C (e_2). Block 420 may involve modifying the corresponding energy value 152 by a percentage $y\%$ of its previous value. In some embodiments, y is less than 2%. In one particular embodiment, y is 1%.

[0040] Method 400 then proceeds to block 430 which involves evaluating a first intermediate condition. In the illustrated embodiment, the block 430 first intermediate condition comprises a condition that ensures that subsequent energy value modifications

- 17 -

do not cause the block 170 shape category of the group G to change. In the particular case of shape category 1 where it is desired to modify a nominal bit value 182 of 0 to a bit value of 1, the block 430 first intermediate condition may be if $e_1 > e_0$. If the block 430 first intermediate condition is positive, then method 400 proceeds to block 440 which involves modifying another energy value. In the particular case of shape category 1 where it is desired to modify a nominal bit value 182 of 0 to a bit value of 1, block 440 may involve decreasing the second energy value 152B (e_1). Block 440 may involve modifying an energy value 152 by a percentage $k\%$ of its previous value. In some embodiments, k is less than 2%. In one particular embodiment, k is 1%.

10

[0041] After block 440 (or in the case where the block 430 inquiry is negative), method 400 proceeds to a second intermediate inquiry in block 450. In the illustrated embodiment, the block 450 second intermediate condition comprises a condition that ensures that subsequent energy value modifications do not cause the block 170 shape category of the group G to change. In the particular case of shape category 1 where it is desired to modify a nominal bit value 182 of 0 to a bit value of 1, the block 450 second intermediate condition may be if $e_0 < e_1$. If the block 450 second intermediate condition is positive, then method 400 proceeds to block 460 which involves modifying another energy value. In the particular case of shape category 1 where it is desired to modify a nominal bit value 182 of 0 to a bit value of 1, block 460 may involve increasing the first energy value 152A (e_0). Block 460 may involve modifying an energy value 152 by a percentage $l\%$ of its previous value. In some embodiments, l is less than 2%. In one particular embodiment, l is 1%.

25 [0042] After block 460 (or in the case where the block 450 inquiry is negative), method 400 proceeds to block 470 where method 400 returns to block 410 to re-evaluate the terminal condition. Method 400 continues to loop in this manner until the block 410 terminal condition is satisfied, in which case method 400 proceeds to block 480 and ends.

- 18 -

By potentially modifying all three energy values 152 in a group G, method 400 may achieve the desired change in bit value relatively quickly (i.e. with fewer iterations) and may therefore consume a relatively low amount of processing resources.

- 5 [0043] Method 400 may also be used for other block 170 shape categories and bit value changes. The various blocks in a particular embodiment of method 400 are shown below for each of the other block 170 shape category changes and each of the other bit value changes.
- 10 *Shape category 1, changing a nominal bit value of 1 to a bit value of 0*
 Block 410 terminal condition: $(e_2 + e_0)/2 \leq e_1 - \delta$
 Block 420: decrease e_0 by $y\%$
 Block 430 first intermediate condition: $e_1 < e_2$
 Block 440: increase e_1 by $k\%$
- 15 Block 450 second intermediate condition: $e_2 > e_1$
 Block 460: decrease e_2 by $l\%$
- Shape category 2, changing a nominal bit value of 0 to a bit value of 1*
 Block 410 terminal condition: $(e_2 + e_0)/2 \geq e_1 + \delta$
- 20 Block 420: increase e_0 by $y\%$
 Block 430 first intermediate condition: $e_1 > e_2$
 Block 440: decrease e_1 by $k\%$
 Block 450 second intermediate condition: $e_2 < e_1$
 Block 460: increase e_2 by $l\%$
- 25 *Shape category 2, changing a nominal bit value of 1 to a bit value of 0*
 Block 410 terminal condition: $(e_2 + e_0)/2 \leq e_1 - \delta$
 Block 420: decrease e_2 by $y\%$

- 19 -

Block 430 first intermediate condition: $e_1 > e_0$

Block 440: increase e_1 by $k\%$

Block 450 second intermediate condition: $e_0 > e_1$

Block 460: decrease e_0 by $l\%$

5

Shape category 3, changing a nominal bit value of 0 to a bit value of 1

Block 410 terminal condition: $(e_1 + e_0)/2 \geq e_2 + \delta$

Block 420: increase e_1 by $y\%$

Block 430 first intermediate condition: $e_2 > e_0$

10

Block 440: decrease e_2 by $k\%$

Block 450 second intermediate condition: $e_0 < e_2$

Block 460: increase e_0 by $l\%$

Shape category 3, changing a nominal bit value of 1 to a bit value of 0

15

Block 410 terminal condition: $(e_1 + e_0)/2 \leq e_2 - \delta$

Block 420: decrease e_0 by $y\%$

Block 430 first intermediate condition: $e_2 < e_1$

Block 440: increase e_2 by $k\%$

Block 450 second intermediate condition: $e_1 > e_2$

20

Block 460: decrease e_1 by $l\%$

Shape category 4, changing a nominal bit value of 0 to a bit value of 1

Block 410 terminal condition: $(e_2 + e_1)/2 \leq e_0 - \delta$

Block 420: decrease e_2 by $y\%$

25

Block 430 first intermediate condition: $e_0 < e_1$

Block 440: increase e_0 by $k\%$

Block 450 second intermediate condition: $e_1 > e_0$

Block 460: decrease e_1 by $l\%$

- 20 -

Shape category 4, changing a nominal bit value of 1 to a bit value of 0

Block 410 terminal condition: $(e_2 + e_1)/2 \geq e_0 + \delta$

Block 420: increase e_1 by $y\%$

Block 430 first intermediate condition: $e_0 > e_2$

5 Block 440: decrease e_0 by $k\%$

Block 450 second intermediate condition: $e_2 < e_0$

Block 460: increase e_2 by $l\%$

Shape category 5, changing a nominal bit value of 0 to a bit value of 1

10 Block 410 terminal condition: $(e_1 + e_0)/2 \geq e_2 + \delta$

Block 420: increase e_0 by $y\%$

Block 430 first intermediate condition: $e_2 > e_1$

Block 440: decrease e_2 by $k\%$

Block 450 second intermediate condition: $e_1 < e_2$

15 Block 460: increase e_1 by $l\%$

Shape category 5, changing a nominal bit value of 1 to a bit value of 0

Block 410 terminal condition: $(e_1 + e_0)/2 \leq e_2 - \delta$

Block 420: decrease e_1 by $y\%$

20 Block 430 first intermediate condition: $e_2 < e_0$

Block 440: increase e_2 by $k\%$

Block 450 second intermediate condition: $e_0 > e_2$

Block 460: decrease e_0 by $l\%$

25 *Shape category 6, changing a nominal bit value of 0 to a bit value of 1*

Block 410 terminal condition: $(e_2 + e_1)/2 \leq e_0 - \delta$

Block 420: decrease e_1 by $y\%$

Block 430 first intermediate condition: $e_0 < e_2$

- 21 -

Block 440: increase e_0 by $k\%$

Block 450 second intermediate condition: $e_2 > e_0$

Block 460: decrease e_2 by $l\%$

5 *Shape category 6, changing a nominal bit value of 1 to a bit value of 0*

Block 410 terminal condition: $(e_2 + e_1)/2 \geq e_0 + \delta$

Block 420: increase e_2 by $y\%$

Block 430 first intermediate condition: $e_0 > e_1$

Block 440: decrease e_0 by $k\%$

10 Block 450 second intermediate condition: $e_1 < e_0$

Block 460: increase e_1 by $l\%$

Those skilled in the art will appreciate that the technique described above represents a particular embodiment of the invention and that other techniques for implementing the
15 block 194 energy value modification are possible.

[0044] Block 194 may involve a wide variety of modifications to energy value(s) 152 to modify a particular group G's nominal bit value 182 of 0 to a new bit value 195 of 1 or from a particular group G's nominal bit value 182 of 1 to a new bit value 195 of 0 in order
20 to encode a desired bit of FEC watermark data 108'. However, it will be appreciated that the block 194 modification of energy value(s) 152 necessitates corresponding changes in audio data 104. That is, modification of energy value(s) 152 in block 194 comprises corresponding adjustment to the values on individual samples in audio data 104 such that re-application of the block 250 integration process would yield the modified energy
25 value(s) 152. For example, an increase in an energy value 152 of a particular time slice 142 may involve increasing the values of individual samples within the corresponding time slice 142 and a decrease in an energy value 152 of a particular time slice 142 may

- 22 -

involve decreasing the values of individual samples within the corresponding time slice 142.

[0045] The adjustment of the individual samples in audio data 104 may be made by
5 adding/subtracting corresponding offsets to the original sample values of audio data 104
or by multiplying the original samples of audio data 104 by suitable scaling factors, for
example. The block 194 modifications may be constrained by their effect on playback of
the audio content represented by audio data 104. If the block 194 modification(s) to
energy value(s) 152 are too significant, then the embedded FEC watermark data 180' will
10 impact playback of the audio content represented by audio data 104 in a manner that is
detectable to the human ear.

[0046] As discussed above and shown in the illustrated examples of Figures 5A-5C and
Figure 8, there are a variety of block 194 modifications to energy value(s) 152 that would
15 have the effect of changing a nominal bit value 182 of 1 to a new bit value 195 of 0 or of
changing a nominal bit value 182 of 0 to a new bit value 195 of 1. In some embodiments,
the block 194 modification process comprises techniques for selecting between a variety
of suitable adjustments to energy value(s) 152, so as to modify a nominal bit value 182 of
0 to a new bit value 195 of 1 or a nominal bit value 182 of 1 to a new bit value 195 of 0
20 using minimum changes to energy value(s) 152, so as to minimize the impact on playback
of the audio content represented by audio data 104. Where the number of time slices 142
in a group G is three (as is the case in the illustrated embodiment), such techniques for
minimizing changes to energy value(s) 152 may involve minimizing the function
 $f = \Delta E_1 + \Delta E_2 + \Delta E_3$, where ΔE_i represents the change in the i^{th} energy value 152, subject to
25 the constraint that the nominal bit value 182 must change to a new bit value 195, either by
changing the nominal bit value 182 within a block 170 shape category or by changing to a
new block 170 shape category.

- 23 -

[0047] Block 194 may comprise a thresholding process. That is, block 194 may have a limit to the amount of modification of any individual energy value 152 and/or to the aggregate modification of the energy values 152 within a particular group G. This thresholding process may be related to a percentage change in the energy value(s) 152, for example. In one particular example, a maximum change to any one particular energy value 152 within a group G is 10% and the maximum aggregate change (i.e. to the sum of energy values 152 within a group G) is 25%. In a small number of cases, it may not be possible to modify the nominal bit value 182 of a group G while complying with the block 194 thresholding process. In such cases, the energy values 152 of the group G need not be adjusted and the resultant corrupted bit of the group G can be reliably accommodated using forward error correction decoding as discussed in more detail below.

[0048] As discussed above, modification of an energy value 152 corresponding to a particular time slice 142 in block 194 also involves corresponding adjustment of the amplitudes of individual samples (see samples 14 in Figure 1) of audio data 104 within the corresponding time slice 142 - i.e. such that re-application of the block 150 integration to the time slice 142 would result in the modified energy value 152. The amount by which each sample in a particular time slice 142 is varied may represent a fraction of the amount by which the energy value 152 of the time slice 142 is modified. The change in the energy value 152 may be divided equally between the samples in the time slice 142.

[0049] In some embodiments, the adjustment of the individual samples may be performed such that there is no change (or minimal change) to the samples at (or near) the beginning and end of time slice 142. This can help to avoid a clipping effect at the boundaries of adjacent time slices 142. In some embodiments, the adjustment of the energy value 152 of a particular time slice 142 may be performed in such a manner that

- 24 -

the aggregate change in the energy value 152 is provided, but that the modifications to the individual samples in time slice 142 are performed (e.g. weighted) in accordance with a function that is relatively large for samples near the middle of time slice 142 and relatively small for samples at the edges of time slice 142. By way of non-limiting
5 example, the adjustment of the individual samples may be performed in accordance with a parabolic function having a vertex at or near the middle sample to the time slice 142. The function used to perform the adjustment of the individual samples of the time slice 142 may have zeroes at the samples corresponding to the edges of the time slice 142 (in the case where the block 194 modification is applied via an additive offset) or may have
10 unity values at the samples corresponding to the edges of the time slice 142 (in the case where the block 194 modification is applied via a multiplicative scaling factor). Again, such functions can help to avoid clipping effects at the boundaries of the time slice 142 that may be audible to a listener.

15 [0050] In addition to or in the alternative to thresholding processes or uneven distributions of adjustments to the individual samples, block 194 may involve other rules and or processes intended to limit the impact of the block 194 watermarking modifications on the quality of the playback of the audio content corresponding to audio data 104.

20 [0051] After block 194, method 100 proceeds to block 196 which involves repeating the procedures of blocks 170, 180, 190 and, if necessary, block 194 for each group G of time slices 142 until the desired FEC watermark 108' is embedded into audio data 104. Method 100 may terminate when FEC watermark 108' is fully embedded in audio data 104. In
25 some embodiments, block 196 involves repeating blocks 170-194 until the end of the available audio data 104 by looping back to the beginning of FEC watermark 108' each time that FEC watermark 108' has been completely embedded. Embedding multiple successive FEC watermarks 108' may permit even more significant compression and/or manipulation of the watermarked media data while permitting the watermark to be
30 successfully extracted when desired.

Extracting the Watermark

[0052] Figure 6 schematically depicts a method 200 for extracting a watermark from watermarked media data 204' in accordance with a particular embodiment of the invention. Method 200 starts in block 233 by obtaining watermarked audio data 204'. Watermarked audio data 204' obtained in block 233 may be identical to watermarked audio data 204 generated by method 100. This is not necessary, however, as watermarked audio data 204' may have been modified after the application of method 100 but prior to being received in block 233. By way of non-limiting example, audio data 204' may have:
5
10
15
20
25
30
been compressed (e.g. converted to MP3 format), been amplified (e.g. increasing or decreasing the magnitudes of individual samples), been re-sampled, converted to analog format, been manipulated in the frequency domain (e.g. non-uniform amplification of certain frequency components), undergone channel mixing, undergone channel reduction, had noise added to it or the like.

[0053] Method 200 then proceeds to optional block 235 where the watermarked audio data 204' is reformatted into reformatted audio data 204'' which is in a suitable format for extracting the watermark. The block 235 reformatting may not be required in the case where watermarked audio data 204' is already in a format suitable for extracting the watermark. In the illustrated embodiment of method 200, audio data 204' (or reformatted audio data 204'') will be in a format suitable for extracting the watermark if it is digitally sampled at the same rate as the audio data 104 into which the watermark was inserted by method 100. Accordingly, the block 235 reformatting may comprise, by way of non-limiting example: sampling audio data 204' (e.g. in the case where audio data 204' is in analog format), resampling audio data 204' (e.g. in the case where audio data 204' has been down-sampled), decompression (e.g. in the case where audio data 204' has been compressed), or the like. Watermark extraction method 200 is robust to amplification, frequency domain manipulation and channel mixing/reduction. Accordingly, the block 235 reformatting does not generally require reformatting to address these types of manipulations.

- 26 -

[0054] Method 200 then proceeds to block 237 which involves optionally aligning the watermarked audio data 204' (or the reformatted watermarked audio data 204") to generate aligned reformatter watermarked audio data 204"". In cases where method 200 has access to the original audio data 104 or to the original watermarked audio data 204, then block 237 may comprise aligning watermarked audio data 204' (or reformatted watermarked audio data 204") to the original audio data 104 or to the original watermarked audio data 204. The block 237 alignment between watermarked audio data 204' (or reformatted watermarked audio data 204") and the original audio data 104 (or the original watermarked audio data 204) may be done by visually comparing the waveform (e.g. by a human user). Additionally or alternatively, the block 237 alignment may comprise comparing the two data sets using a suitable metric between individual samples (e.g. by a least squares comparison or some other suitable metric). The watermarked audio data 204' (or reformatted watermarked audio data 204") may be translated in the time domain by adding zeroes or by removing samples at the beginning or the end of the data stream, so as to align watermarked audio data 204' (or reformatted watermarked audio data 204") more closely to the original audio data 104 (or the original watermarked audio data 204). Adding zeros and/or removing data may be repeated until the least squares comparison (or other suitable metric) is minimized or until the visual waveform comparison indicates that waveforms are aligned. The result of the block 237 alignment process is aligned reformatted watermarked audio data 204"". There may be other suitable techniques for aligning watermarked audio data 204' (or reformatted watermarked audio data 204") to the original audio data 104 (or the original watermarked audio data 204).

[0055] In some cases, watermarked audio data 204' (or reformatted audio data 204") may already be aligned with the original audio data 104 of the original watermarked audio data 204. In this circumstance, the block 237 alignment procedure may not be required. In some cases, neither the original audio data 104 nor the original watermarked audio data 204 are available. In this circumstance, method 200 may bypass the block 237 alignment

- 27 -

procedure. The situation where neither the original audio data 104 nor the original watermarked audio data 204 is available is discussed in more detail below.

5 [0056] For simplicity, watermarked audio data 204', reformatted watermarked audio data 204" and aligned reformatted audio data 204''' are referred to in the remainder of this description as watermarked audio data 204, it being understood that watermarked audio data 204 may comprise any of watermarked audio data 204', reformatted watermarked audio data 204" and aligned reformatted audio data 204'''.

10 [0057] Method 200 processes watermarked audio data 204 in blocks 240, 250, 260, 270 and 280. Blocks 240, 250, 260, 270 and 280 may respectively be substantially similar to blocks 140, 150, 160, 170 and 180 of method 100. Blocks 240, 250, 260, 270 and 280 involve dividing watermarked audio data 204 into sections/time slices 242, integrating the samples in each time slice 242 to obtain an energy value 252 for each time slice 242,
15 grouping the time slices 242 into groups G, categorizing each time slice 242 into a shape category and applying a bit assignment rule to each group G to determine a nominal bit value 282 for each group G. Because of the embedded watermark (i.e. because of the block 194 modification of energy values 152), the energy values 252 assigned to each group G in block 250 may be different than the energy values 152 assigned to each group
20 G in block 150. Similarly, the block 280 application of the bit assignment rule will typically result in different nominal bit values 282 than the nominal bit values 182 obtained when the bit assignment rule is applied in block 180.

25 [0058] The nominal bit values 282 extracted in the block 280 application of the bit assignment rule represent the method 200 estimate of the FEC watermark bit corresponding to a particular group G. The estimated FEC watermark estimate is referred to in Figure 6 as FEC watermark estimate 286. Block 285 involves repeating the processes of blocks 270 and 280 for each group G of time slices 242 to obtain an FEC watermark estimate 286 for each group G.

30

- 28 -

[0059] In some embodiments, block 285 may involve repeating blocks 270 and 280 until a single FEC watermark estimate 286 is extracted. In other embodiments, where multiple successive watermarks are encoded into watermarked media data 204', block 285 may comprise repeating blocks 270 and 280 until the end of watermarked media data 204'. In such cases, FEC watermark estimate 286 may comprise more than one single estimate of FEC watermark 108' and method 200 may comprise subsequently processing the multiple FEC watermark estimates in FEC watermark estimate 286 to obtain bit values for a single FEC watermark estimate. By way of non-limiting example, such subsequent processing may comprise averaging and/or filtering the data from the multiple FEC watermark estimates to obtain bit values for a single FEC watermark estimate. Encoding and extracting multiple successive repetitions of FEC watermark 108' into audio data 104 may permit even more significant compression and/or manipulation of watermarked media data 204 while permitting watermark 108 to be successfully extracted from watermarked media data 204' when desired.

15

[0060] Block 287 involves applying a forward error correction decoding scheme to FEC watermark estimate 286. The forward error correction decoding in block 287 is complementary to the forward correction encoding in block 120 of method 100. As discussed above, in one particular embodiment, blocks 120 and 287 make use of LDPC forward error correction techniques. The output of block 287 is watermark estimate 289 which represents the method 200 estimate of watermark data 108. Preferably, watermark estimate 289 is identical to watermark data 108. Advantageously, the forward error correction of blocks 120 and 287 allows the nominal bit values determined in blocks 270 and 280 to have a relatively large number of errors while still ensuring that watermark estimate 289 is identical to watermark 108. That is, the FEC watermark estimate 286 obtained in block 285 need not be identical to FEC watermark 108'. In preferred embodiments, as many as 1/3 or more of the data bits of watermark estimate 286 may contain errors and may still yield a watermark estimate 289 that is identical to watermark data 108.

30

- 29 -

[0061] Once watermark estimate 289 is obtained, it may be used for a wide variety of purposes including, without limitation, looking up additional data in a suitable database which may be indexed by watermark data 108.

5 [0062] As mentioned briefly above, in some cases, neither the original audio data 104 nor the original watermarked audio data 204 are available. In such circumstances, it is still possible to use a method similar to method 200 to extract a watermark estimate 289. Watermark estimate 289 can then be compared to watermark data 108, for example. If
10 neither the original audio data 104 nor the original watermarked audio data 204 are available, then it is not possible to implement the block 237 alignment procedures discussed above. Instead, the relevant parts of method 200 may be implemented several times with different time domain translations of the watermarked audio data 204 and the resulting watermark estimates 289 from each iteration may be compared to watermark data 108.

15 [0063] In one embodiment where neither the original audio data 104 nor the original watermarked audio data 204 are available, method 200 (without block 237) may be implemented a first time and the resulting watermark estimate 289 may be compared to watermark data 108. If the watermark estimate 289 matches (or is sufficiently close to
20 matching) watermark data 108, then it is concluded that watermarked audio data 204 was indeed watermarked with watermark data 108. If the watermark estimate 289 does not match watermark data 108, then the watermarked audio data 204 may be shifted by n samples in the time domain (e.g. by adding n samples at the beginning or end of watermarked audio data 204 or by removing n samples from the beginning or end of
25 watermarked audio data 204). In some embodiments, n may be greater than 100,000 samples, for example, which corresponds to a time shift of ~2.5 seconds for 44.1kHz audio sampling rate. In one particular embodiment, n is 3,000 samples.

[0064] After watermarked audio data 204 is shifted, blocks 240 through 287 of method
30 200 are repeated to obtain a second watermark estimate 289. This second watermark

- 30 -

estimate 289 is again compared to watermark data 108 to ascertain whether there is a match. If there is no match, then watermarked audio data 204 may be shifted by a suitable fraction of n samples in the time domain. For example, watermarked audio data 204 may be shifted by $n/2$ samples. Again, blocks 240 through 287 of method 200 are repeated to
5 obtain a third watermark estimate 289, which is compared to watermark data 108. If there is no match, then watermarked audio data 204 may be shifted by a smaller fraction of n samples in the time domain. For example, watermarked audio data 204 may be shifted by $n/4$ samples. Again, blocks 240 through 287 of method 200 are repeated to obtain a fourth watermark estimate 289, which is compared to watermark data 108. If there is no match,
10 then watermark audio data 204 may be shifted by a multiple of the smaller fraction, if this shift has not yet been evaluated. For example, watermarked audio data 204 may be shifted by $3n/4$ samples. Similarly, for a fraction $n/8$, watermarked audio data 204 may be shifted by $3n/8$, $5n/8$ and $7n/8$.

15 **[0065]** This process may be repeated until either it is concluded that watermark data 108 is embedded in watermarked audio data 204 or until the shift of watermarked audio data 204 approaches a sufficiently small number that it may be concluded that audio data 204 is, in fact, not watermarked. In some embodiments, this condition for discontinuing the watermark retrieval process is a shift of less than 200 samples.

20

[0066] The watermarking process of method 100 and/or the watermark extraction process of method 200 may easily be modified for application to other forms of digital media. For example, Figure 7 schematically depicts monochromatic (or grayscale) digital image data 310 which may be represented as a series 312 of samples 314 in a spatial domain
25 (referred to as x in the schematic illustration of Figure 7). Each of samples 314 may be referred to as a pixel 314 of image data 310. Typically image data 310 is two dimensional. Image data 310 may comprise pixels 314 corresponding to a portion of a two dimensional image. By way of non-limiting example, image data 310 may comprise pixels from one row of the image, one column of the image or one rectangular section of
30 the image, for example. Image data 310 may comprise a concatenation of the pixels

- 31 -

corresponding to the two dimensional image into a single array. By way of non-limiting example, image data 310 may comprise a concatenation of the pixels corresponding to two or more rows of the two dimensional image. In the illustrated example of Figure 7, the sampling period of image data 310 is X_s in the spatial domain x .

5

[0067] When applying methods 100 and 200 to digital image data 310, the block 140 and 240 time slices 142, 242 may actually represent spatial slices (in the spatial dimension x of the Figure 7 schematic illustration). Accordingly, references to “time slices” in the description above should be understood to include or otherwise refer to “sections” which may be in any suitable sampling domain. Each section comprises a plurality of samples which are adjacent to one another in the sampling domain.

10

[0068] Color image data may comprise a plurality of sample values for each pixel. For example, in the case of RGB representation, each pixel may comprise a red sample value and green sample value and a blue sample value. In such cases, watermark data may be encoded into one set of sample values (e.g. the green sample values) or into each set of sample values or into some pattern of alternating between the sets of sample values.

15

[0069] In other respects, the application of methods 100 and/or 200 to image data 310 may be similar to that described above for audio data.

20

[0070] The watermarking process of method 100 and/or the watermark extraction process of method 200 may also easily be modified for application to video media data. When applying methods 100, 200 to video media data, each frame (or particular portions of each frame) of the video media data may be treated as two dimensional images of the type described above and then the video data may be processed in the manner described above for still images. Additionally or alternatively, when applying methods 100, 200 to video media data, each frame (or particular portions of each frame) may be treated as a time slice/section.

25

30

- 32 -

[0071] Certain implementations of the invention comprise computer processors which execute software instructions which cause the processors to perform one or more methods of the invention. For example, methods 100 and/or 200 may be implemented by one or more processors which execute software instructions which cause the processor to

5 perform these methods. Such software instructions may be retrieved from a program memory accessible to the processors. The invention may also be provided in the form of a program product. The program product may comprise any medium which carries a set of computer-readable instructions which, when executed by a data processor, cause the data

10 processor to execute a method of the invention. Program products according to the invention may be in any of a wide variety of forms. The program product may comprise, for example, physical media such as magnetic data storage media including floppy diskettes, hard disk drives, optical data storage media including CD ROMs, DVDs, electronic data storage media including ROMs, flash RAM, or the like. The instructions may be present on the program product in encrypted and/or compressed formats.

15

[0072] Where a component (e.g. a software module, processor, assembly, device, circuit, etc.) is referred to above, unless otherwise indicated, reference to that component (including a reference to a "means") should be interpreted as including as equivalents of that component any component which performs the function of the described component

20 (i.e. that is functionally equivalent), including components which are not structurally equivalent to the disclosed structure which performs the function in the illustrated exemplary embodiments of the invention.

[0073] As will be apparent to those skilled in the art in the light of the foregoing

25 disclosure, many alterations and modifications are possible in the practice of this invention without departing from the spirit or scope thereof. For example:

- when grouping time slices 142 into groups G (blocks 160, 260), groups G may be selected such that one or more time slices/sections 142 on the edge(s) of group G overlap with an adjacent group G. For example, the last time slice/section 142 of

30 the first group G_1 may also be the first time slice/section 142 of the second group

- 33 -

G_2 . This type of overlapping is permissible, provided that the energy value(s) 152 of the overlapped time slice(s)/section(s) 142 is not varied in block 194.

Accordingly, block 194 may comprise a process for preventing the energy value(s) 152 corresponding to the overlapped time slice(s)/section(s) 142 from being

5 selected for modification. Alternatively, it may be possible to allow the energy value(s) 152 of the overlapped time slice(s)/section(s) 152 to be varied in block 194, but to impose limits on the variation of the energy value(s) 152 of overlapped

10 time slice(s) 152 based on the effect of such energy value variation on the shape category and/or nominal bit value 182 of the adjacent group G. For example, it may be possible to allow the energy value 152 of an overlapped time slice/section 142 to be varied in block 194, but to impose limits on the variation of the energy

15 value 152 of the overlapped time slice/section 142 based on a desire not to change the shape category of the adjacent group G. The overlapping of time slices 142 from adjacent groups G provides a larger number of groups G within a given set of media data and therefore allows a correspondingly greater number of watermarking bits to be encoded into the media data.

- The methods described above need not be implemented in the specific order illustrated in the schematic drawings and that the order of the operations described in the schematic block diagrams may vary while achieving the same effect. For example, those skilled in the art can envisage other looping structures, where a blocks 140, 150, 160, 170, 180, 190 and if necessary block 194 are performed for a single group G of time slices/sections 142 and then the process loops back to block 140 for the next time slice/section 142.

20

- The shapes described above and shown in Figure 4 represent one possible categorization of shapes when each group G includes three time slices/sections 142. Groups G may be defined to have a different number of time slices/section 142, in which case the available shape categories may differ. For example, if groups G are defined to comprise four time slices/sections 142, then there could be a large number of shape categories.

25

- 34 -

- Method 400 described above, comprises repetitively incrementally modifying energy values by a percentage of their previous value. In this manner, sections having larger energy values are modified relatively more than sections having smaller energy values. This is not necessary. In some embodiments, method 400 could be implemented by repetitively incrementally modifying energy values by absolute amounts.

- 35 -

WHAT IS CLAIMED IS:

1. A method for encoding watermark information into media data containing a series
5 of digital samples in a sample domain, the method comprising:
 - dividing the series of digital samples into a plurality of sections in the
sample domain, each section comprising a corresponding plurality of samples;
 - processing the corresponding plurality of samples in each section to obtain
a single energy value associated with each section;
 - 10 grouping the sections into groups, each group containing three or more
sections;
 - assigning a nominal bit value to each group according to a bit assignment
rule, the bit assignment rule based on the energy values of the sections in the
group;
 - 15 assigning a watermark bit value to each group;
 - for each group, comparing the watermark bit value to the nominal bit
value and, if the nominal bit value and the watermark bit value of the watermark
information bit do not match, modifying one or more energy values of one or
more corresponding sections in the group such that re-application of the bit
20 assignment rule would assign the watermark bit value to the group;
 - wherein assigning the nominal bit value to each group according to the bit
assignment rule comprises:
 - categorizing each group into one of a plurality of categories according to a
categorization rule; and
 - 25 for each category, assigning a nominal zero bit value or a nominal one bit
value to each group according to a unique category bit assignment rule associated
with the category.
2. A method according to claim 1 wherein each group comprises an ordered plurality
30 of three or more sections, wherein first and second sections are adjacent to one

- 36 -

another in the sample domain and a third section is adjacent to the second section in the sample domain, but spaced apart from the first section in the sample domain.

- 5 3. A method according to claim 2 wherein the categorization rule is based on: the energy value of the first section relative to the energy value of the second section; and the energy value of the third section relative to the energy values of the first and second sections.
- 10 4. A method according to claim 3 wherein at least one of:
 the media data comprises audio data and the sample domain comprises a time domain;
 the media data comprises image data and the sample domain comprises a spatial domain; and
15 the media data comprises video data and the sample domain comprises at least one of: a spatial domain and a time domain.
5. A method according to claim 3 wherein, for a particular group, the categorization rule comprises:
20 categorizing the particular group into a first shape category if the energy value of the second section is greater than an energy value of the first section and the energy value of the third section is greater than the energy values of the first and second sections;
 categorizing the particular group into a second shape category if the energy
25 value of the second section is less than an energy value of the first section and the energy value of the third section is less than the energy values of the first and second sections;
 categorizing the particular group into a third shape category if the energy value of the second section is greater than an energy value of the first section and

- 37 -

the energy value of the third section is between the energy values of the first and second sections;

5 categorizing the particular group into a fourth shape category if the energy value of the second section is greater than an energy value of the first section and the energy value of the third section is less than the energy values of the first and second sections;

10 categorizing the particular group into a fifth shape category if the energy value of the second section is less than an energy value of the first section and the energy value of the third section is between the energy values of the first and second sections; and

15 categorizing the particular group into a sixth shape category if the energy value of the second section is greater than an energy value of the first section and the energy value of the third section is greater than the energy values of the first and second sections.

15

6. A method according to claim 3 wherein, for each category, the corresponding category bit assignment rule is based on relative differences between the energy values of sections that are adjacent to one another in the sample domain.

20 7. A method according to claim 6 wherein, for each category, the corresponding category bit assignment rule is based on a difference between the energy values of the first and second sections relative to a difference between the energy values of the second and third sections.

25 8. A method according to claim 6 wherein, for each category, the category bit assignment rule is based on a slope between the energy values of the first and second sections relative to a slope between the energy values of the second and third sections.

- 38 -

9. A method according to claim 5 wherein, for each category, the category bit assignment rule is based on a difference between the energy values of the first and second sections relative to a difference between the energy values of the second and third sections and wherein:
- 5 for the first and second shape categories, the category bit assignment rule assigns a first bit value where a magnitude of a first slope between the energy values of the first and second sections is greater than a magnitude of a second slope between the energy values of the second and third sections and assigns a second bit value where the magnitude value of the first slope is less than the
- 10 magnitude of the second slope;
- for the third and fifth shape categories, the category bit assignment rule assigns a first bit value where a magnitude of a second slope between the energy values of the second and third sections is greater than one half of a magnitude of a first slope between the first and second energy values and assigns a second bit
- 15 value where the magnitude of the second slope is less than one half of the magnitude of the first slope; and
- for the fourth and sixth shape categories, the category bit assignment rule assigns a first bit value where a magnitude of a second slope between the energy values of the second and third sections is greater than twice a magnitude of a first
- 20 slope between the first and second energy values and assigns a second bit value where the magnitude of the second slope is less than twice the magnitude of the first slope.
10. A method according to claim 1 wherein modifying one or more energy values of
- 25 one or more corresponding sections in the group comprises modifying a single energy value corresponding to a single section in the group.
11. A method according to claim 1 wherein modifying one or more energy values of
- 30 one or more corresponding sections in the group comprises modifying the one or more energy values in the group in such a manner that re-application of the

- 39 -

categorization rule to the group would categorize the group into a different category.

- 5
12. A method according to claim 1 wherein modifying one or more energy values of one or more corresponding sections in the group comprises modifying the one or more energy values in the group in such a manner that re-application of the category bit assignment rule to the group would assign a different bit value to the group but that re-application of the categorization rule to the group would categorize the group into the same category.
- 10
13. A method according to claim 1 wherein modifying one or more energy values of one or more corresponding sections in the group comprises modifying three or more energy values corresponding to three or more sections in the group.
- 15
14. A method according to claim 13 wherein modifying the three or more energy values corresponding to three or more sections in the group comprises repetitively incrementally modifying at least one of the three or more energy values until satisfying a termination condition, the termination condition based on the three or more energy values.
- 20
15. A method according to claim 14 wherein repetitively incrementally modifying at least one of the three or more energy values comprises repetitively incrementally modifying a plurality of the energy values until satisfying one or more intermediate conditions, the one or more intermediate conditions based on two of
- 25
- the three or more energy values.
- 30
16. A method according to claim 14 wherein repetitively incrementally modifying at least one of the three or more energy values comprises, for each repetitive iteration, incrementally modifying the at least one of the three or more energy values by a fixed percentage of its previous value.

17. A method according to claim 14 wherein the termination condition comprises a condition where re-application of the bit assignment rule would assign the watermark bit value to the group modified by an error margin δ .
- 5
18. A method according to claim 17 wherein modifying three or more energy values corresponding to three or more sections in the group comprises minimizing a sum of the changes to the three or more energy values while satisfying the termination condition.
- 10
19. A method according to claim 1 wherein grouping the sections into groups comprises defining the groups to overlap one another such that the adjacent groups in the sampling domain include one or more sections in common.
- 15
20. A method according to claim 1 wherein the corresponding plurality of samples in each section is equal and wherein processing the corresponding plurality of samples in each section includes, for each section, integrating values of the corresponding plurality of samples in the section to obtain the energy value without performing a division operation.
- 20
21. A computer program product comprising a computer readable medium storing programmable instructions which, when executed by a processor, cause the processor to carry out the method of claim 1.
- 25
22. A method for extracting watermark information from media data containing a series of digital samples in a sample domain, the method comprising:
- dividing the series of digital samples into a plurality of sections in the sample domain, each section comprising a corresponding plurality of samples;
- processing the corresponding plurality of samples in each section to obtain
- 30 a single energy value associated with each section;

grouping the sections into groups, each group containing three or more sections;

assigning a nominal bit value to each group according to a bit assignment rule, the bit assignment rule based on the energy values of the sections in the group;

wherein assigning the nominal bit value to each group according to the bit assignment rule comprises:

categorizing each group into one of a plurality of categories according to a categorization rule; and

for each category, assigning a nominal zero bit value or a nominal one bit value to each group according to a unique category bit assignment rule associated with the category.

23. A method according to claim 1 comprising aligning the series of samples with an original version of the media data, wherein aligning the series of samples comprises adding zero samples to the series or removing samples from the series to translate the series in the sample domain.

24. A method according to claim 23 wherein aligning the series of samples with the original version of the media data comprises comparing each sample of series with a corresponding sample of the original version of the media data and adding zero samples or removing samples from the series to minimize a metric associated with a difference between the series and the original version of the media data.

25. A computer program product comprising a computer readable medium storing programmable instructions which, when executed by a processor, cause the processor to carry out the method of claim 22.

1/7




FIGURE 1




FIGURE 3

2/7

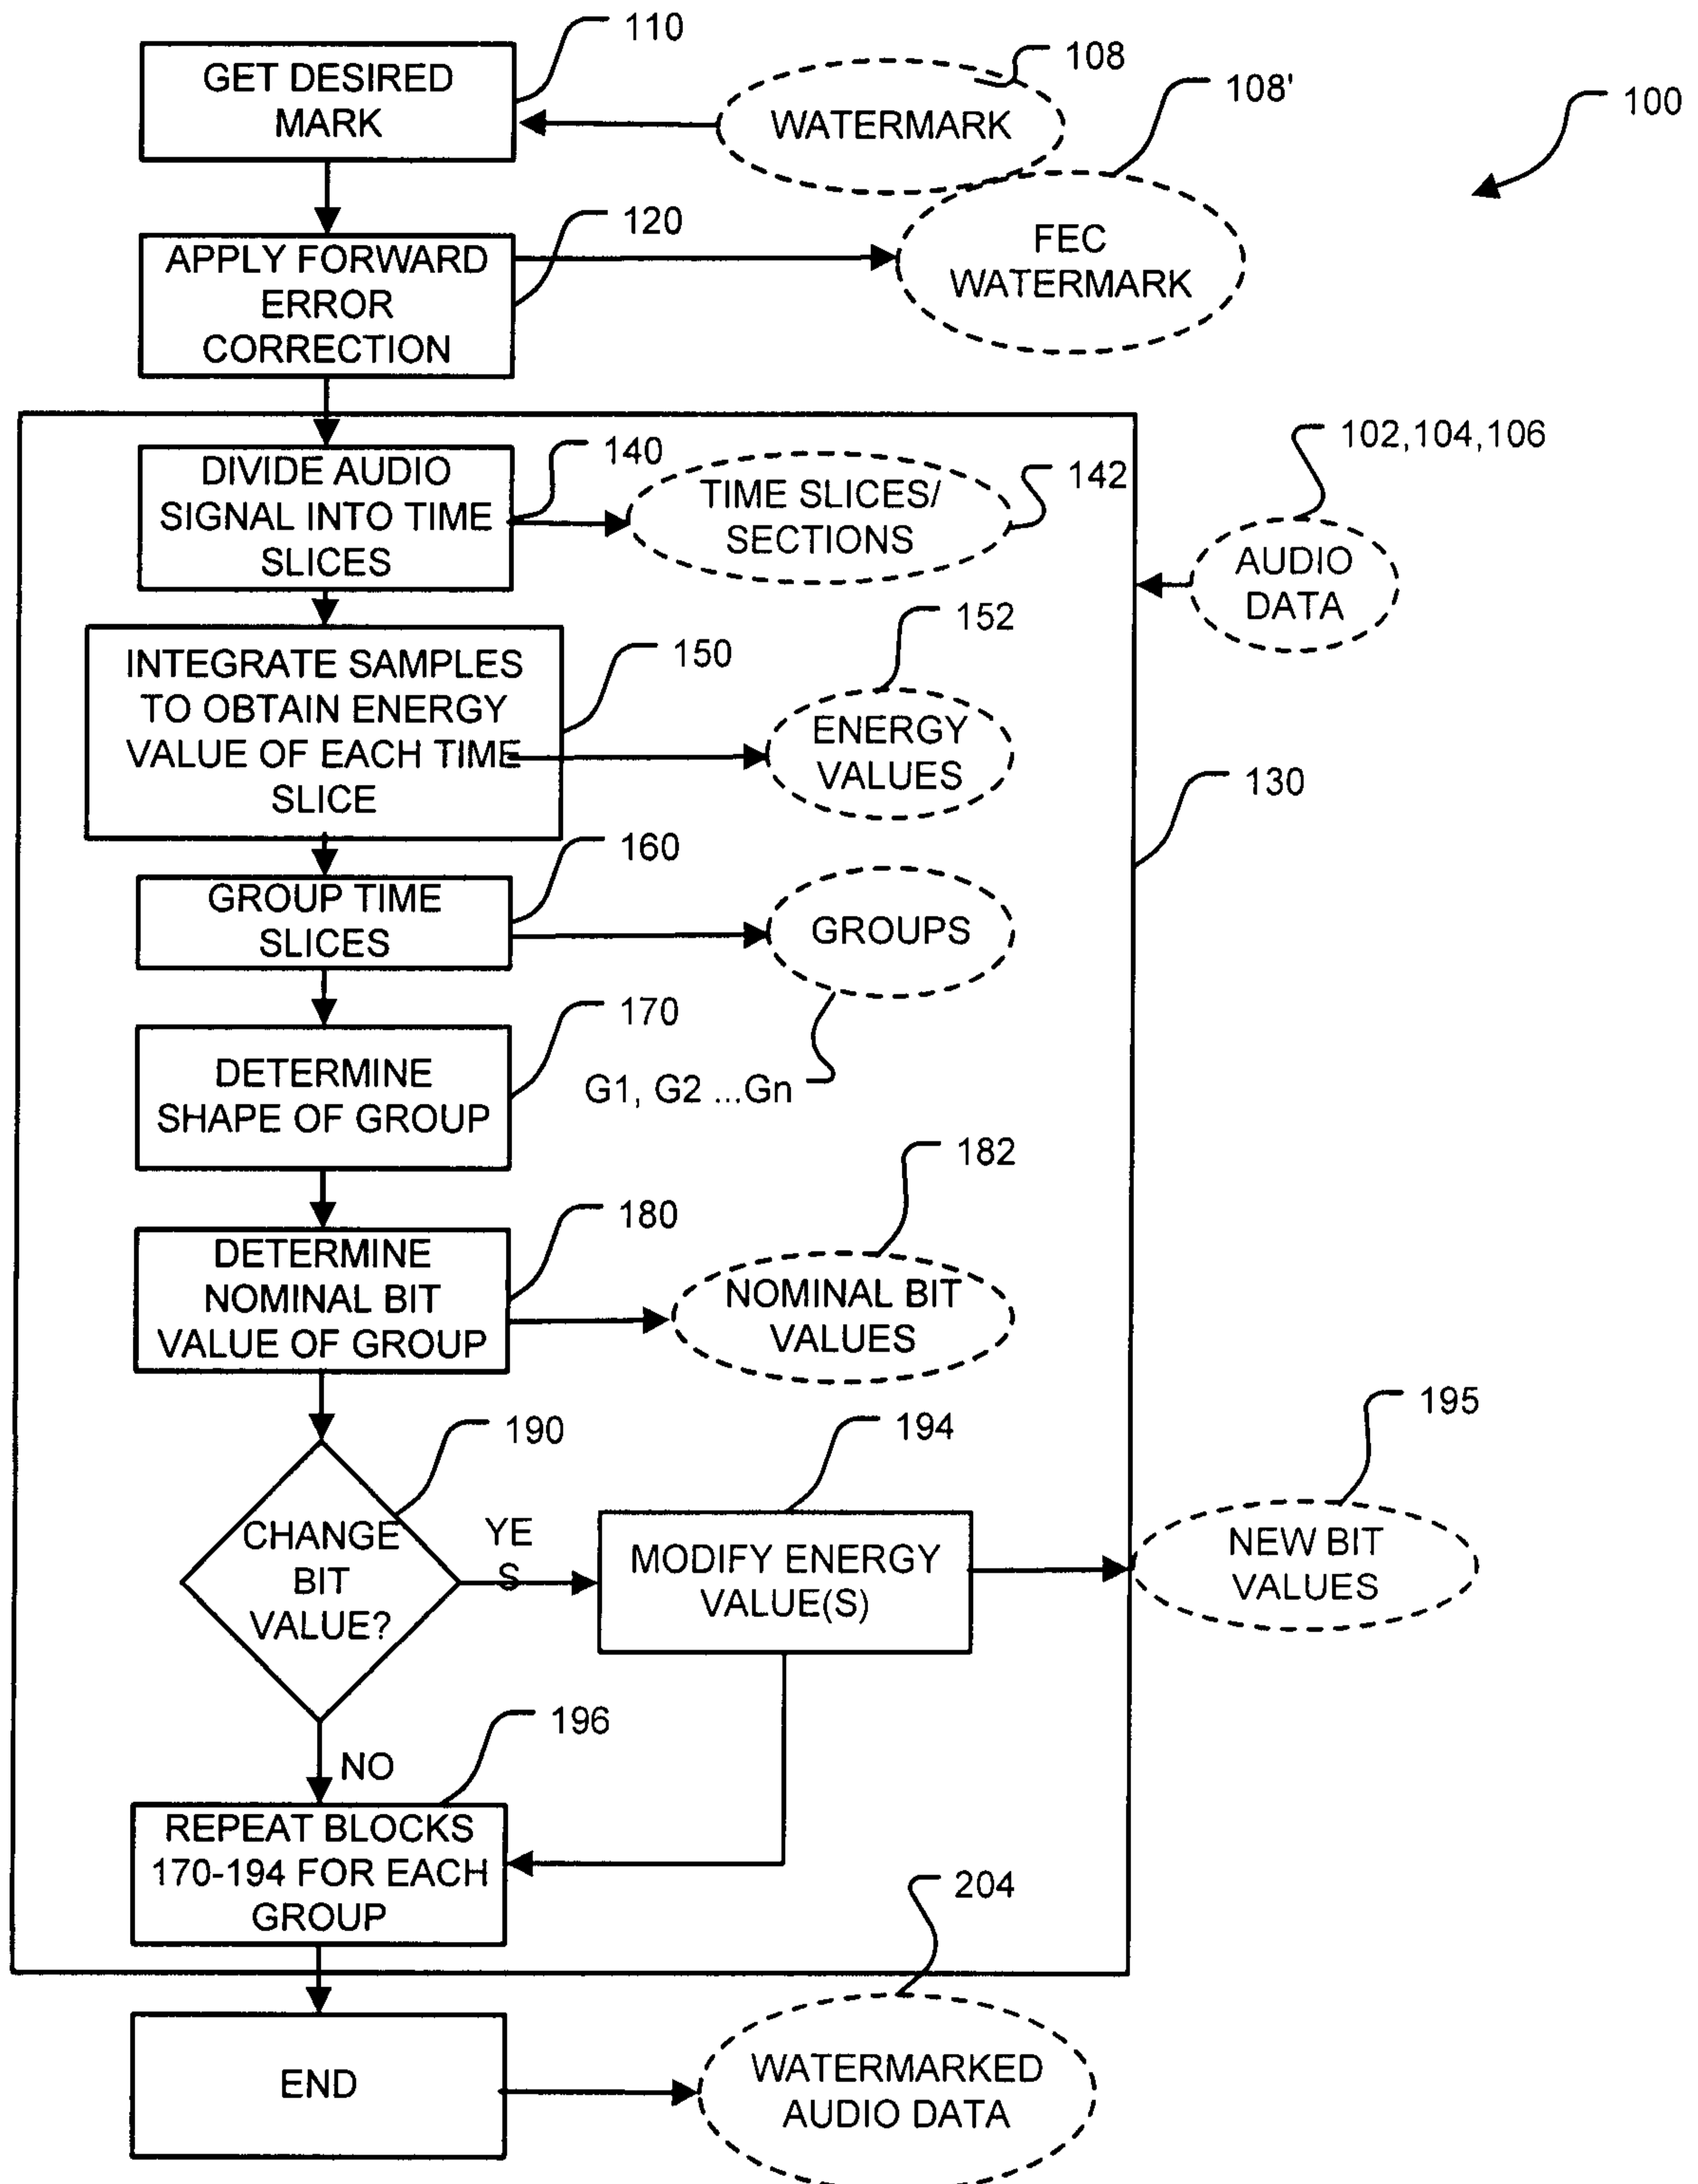


FIGURE 2

3/7

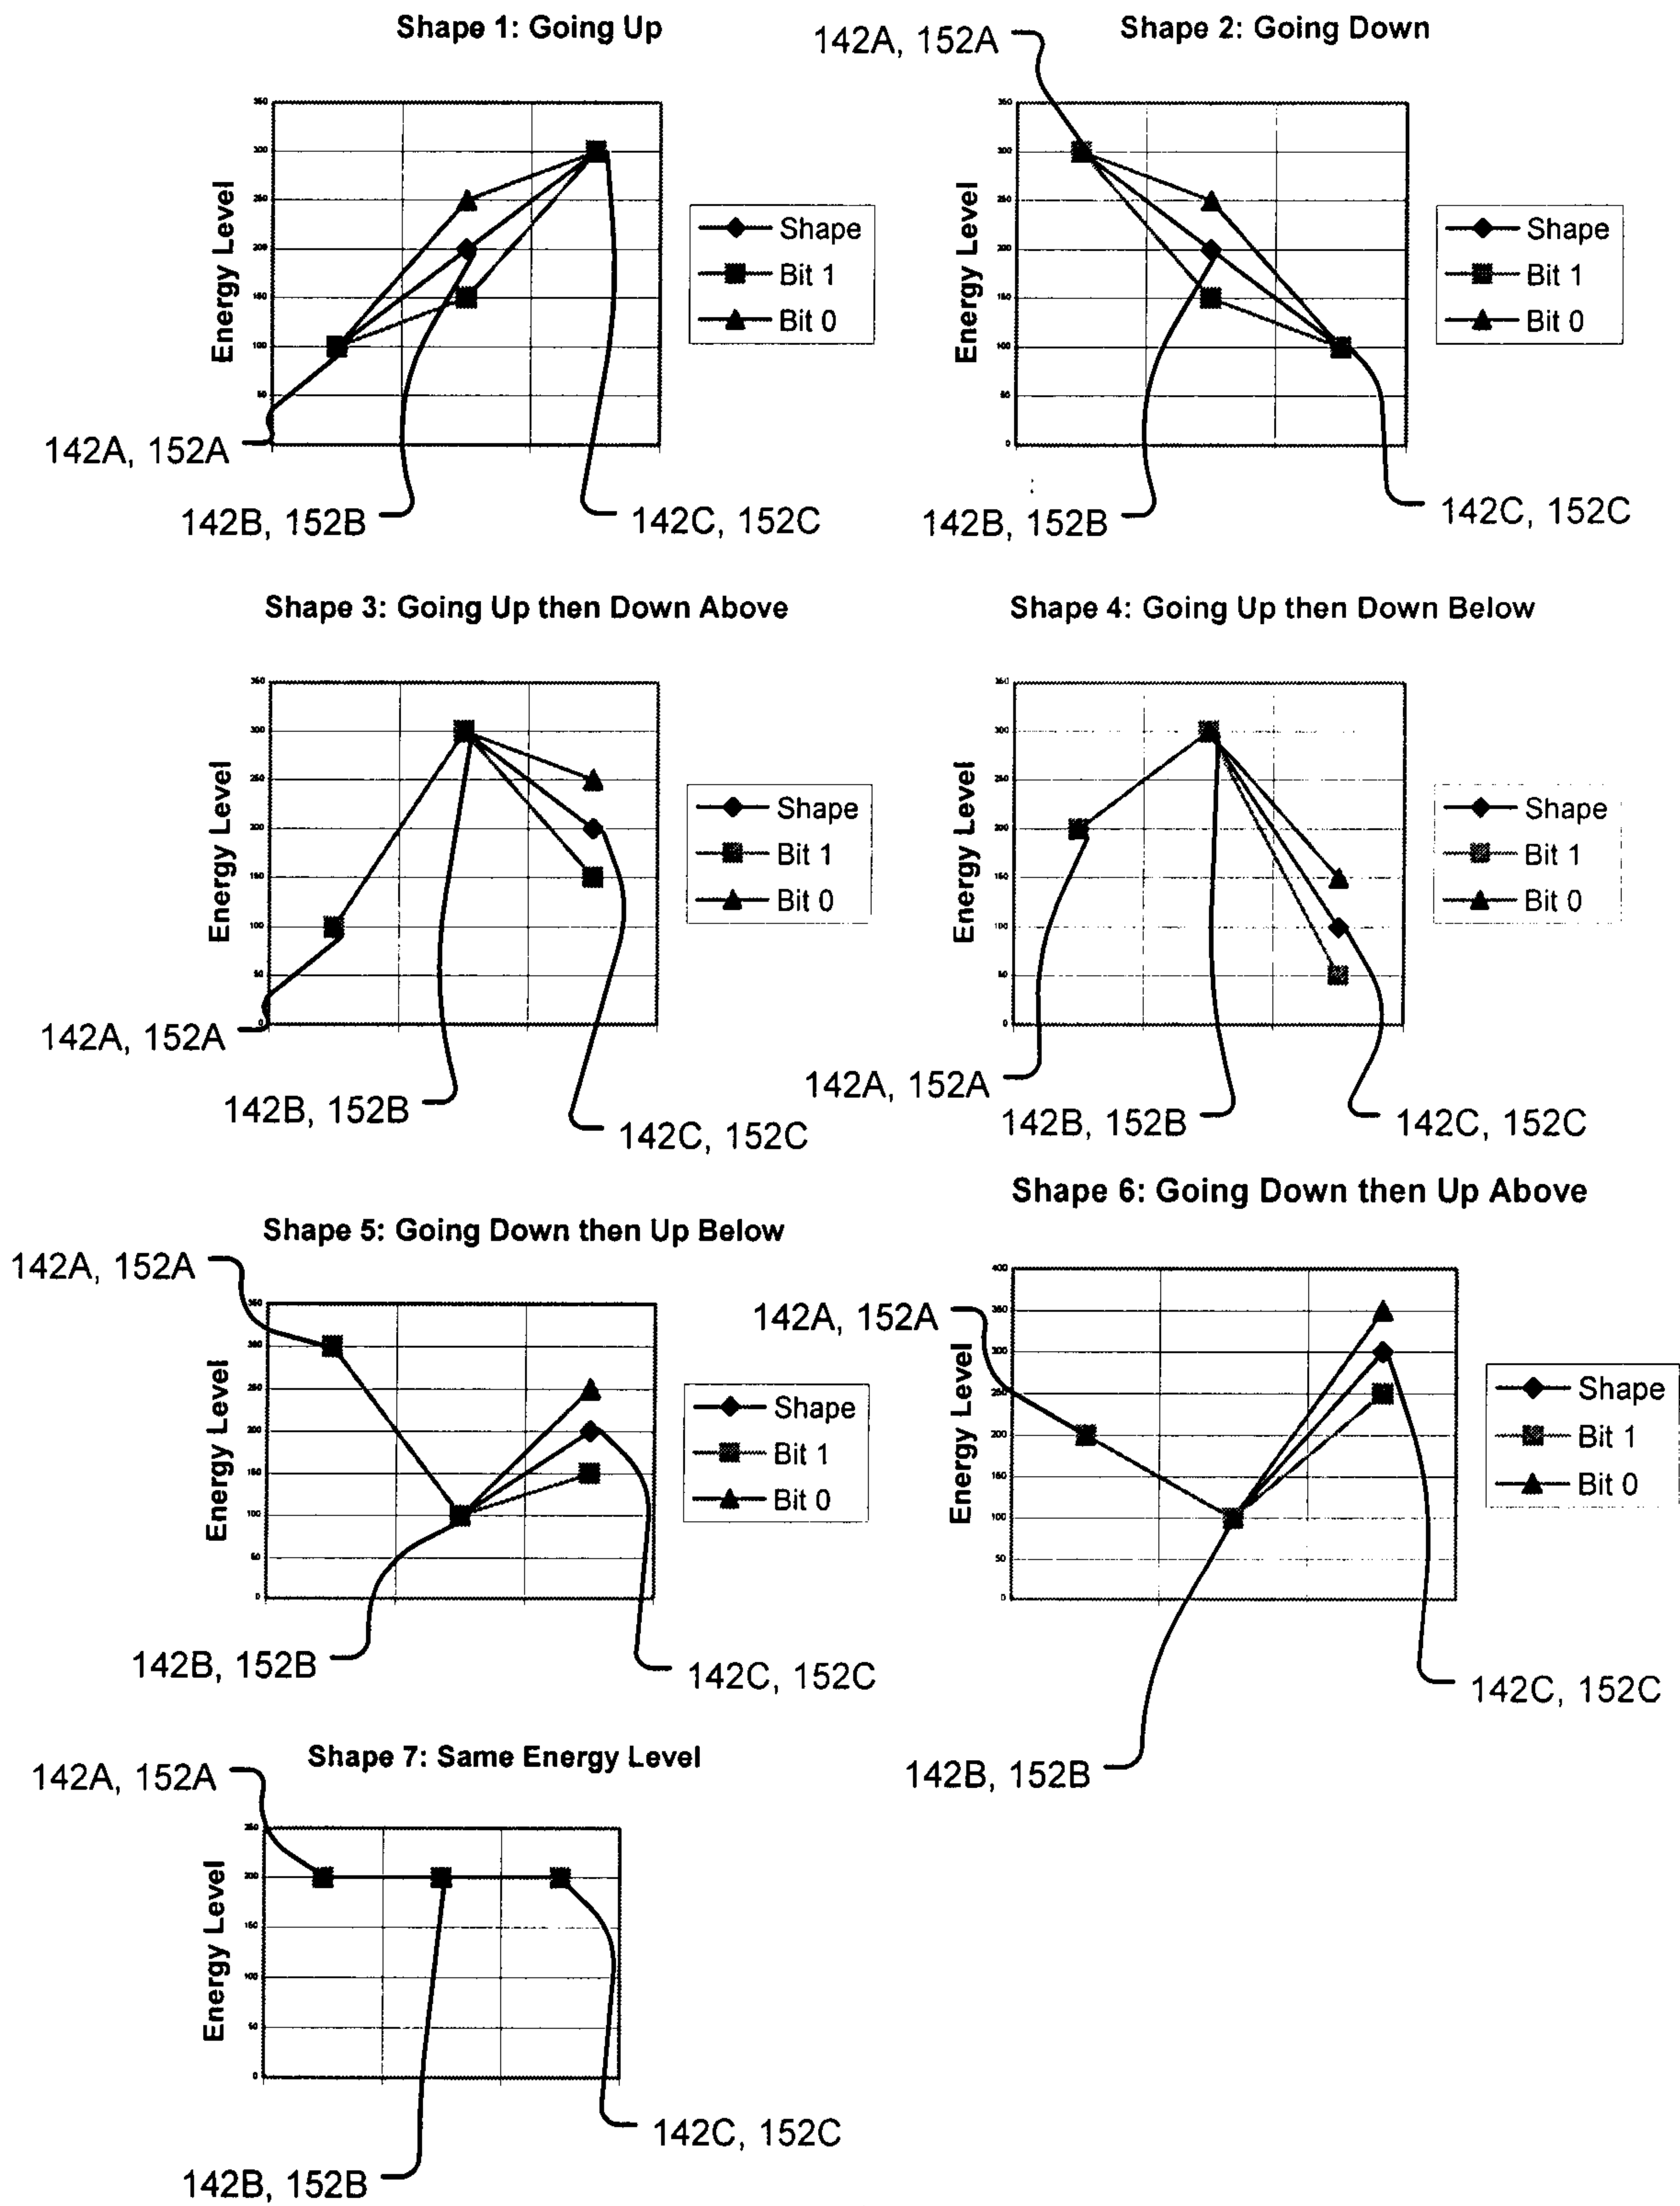





FIGURE 4

4/7



5/7

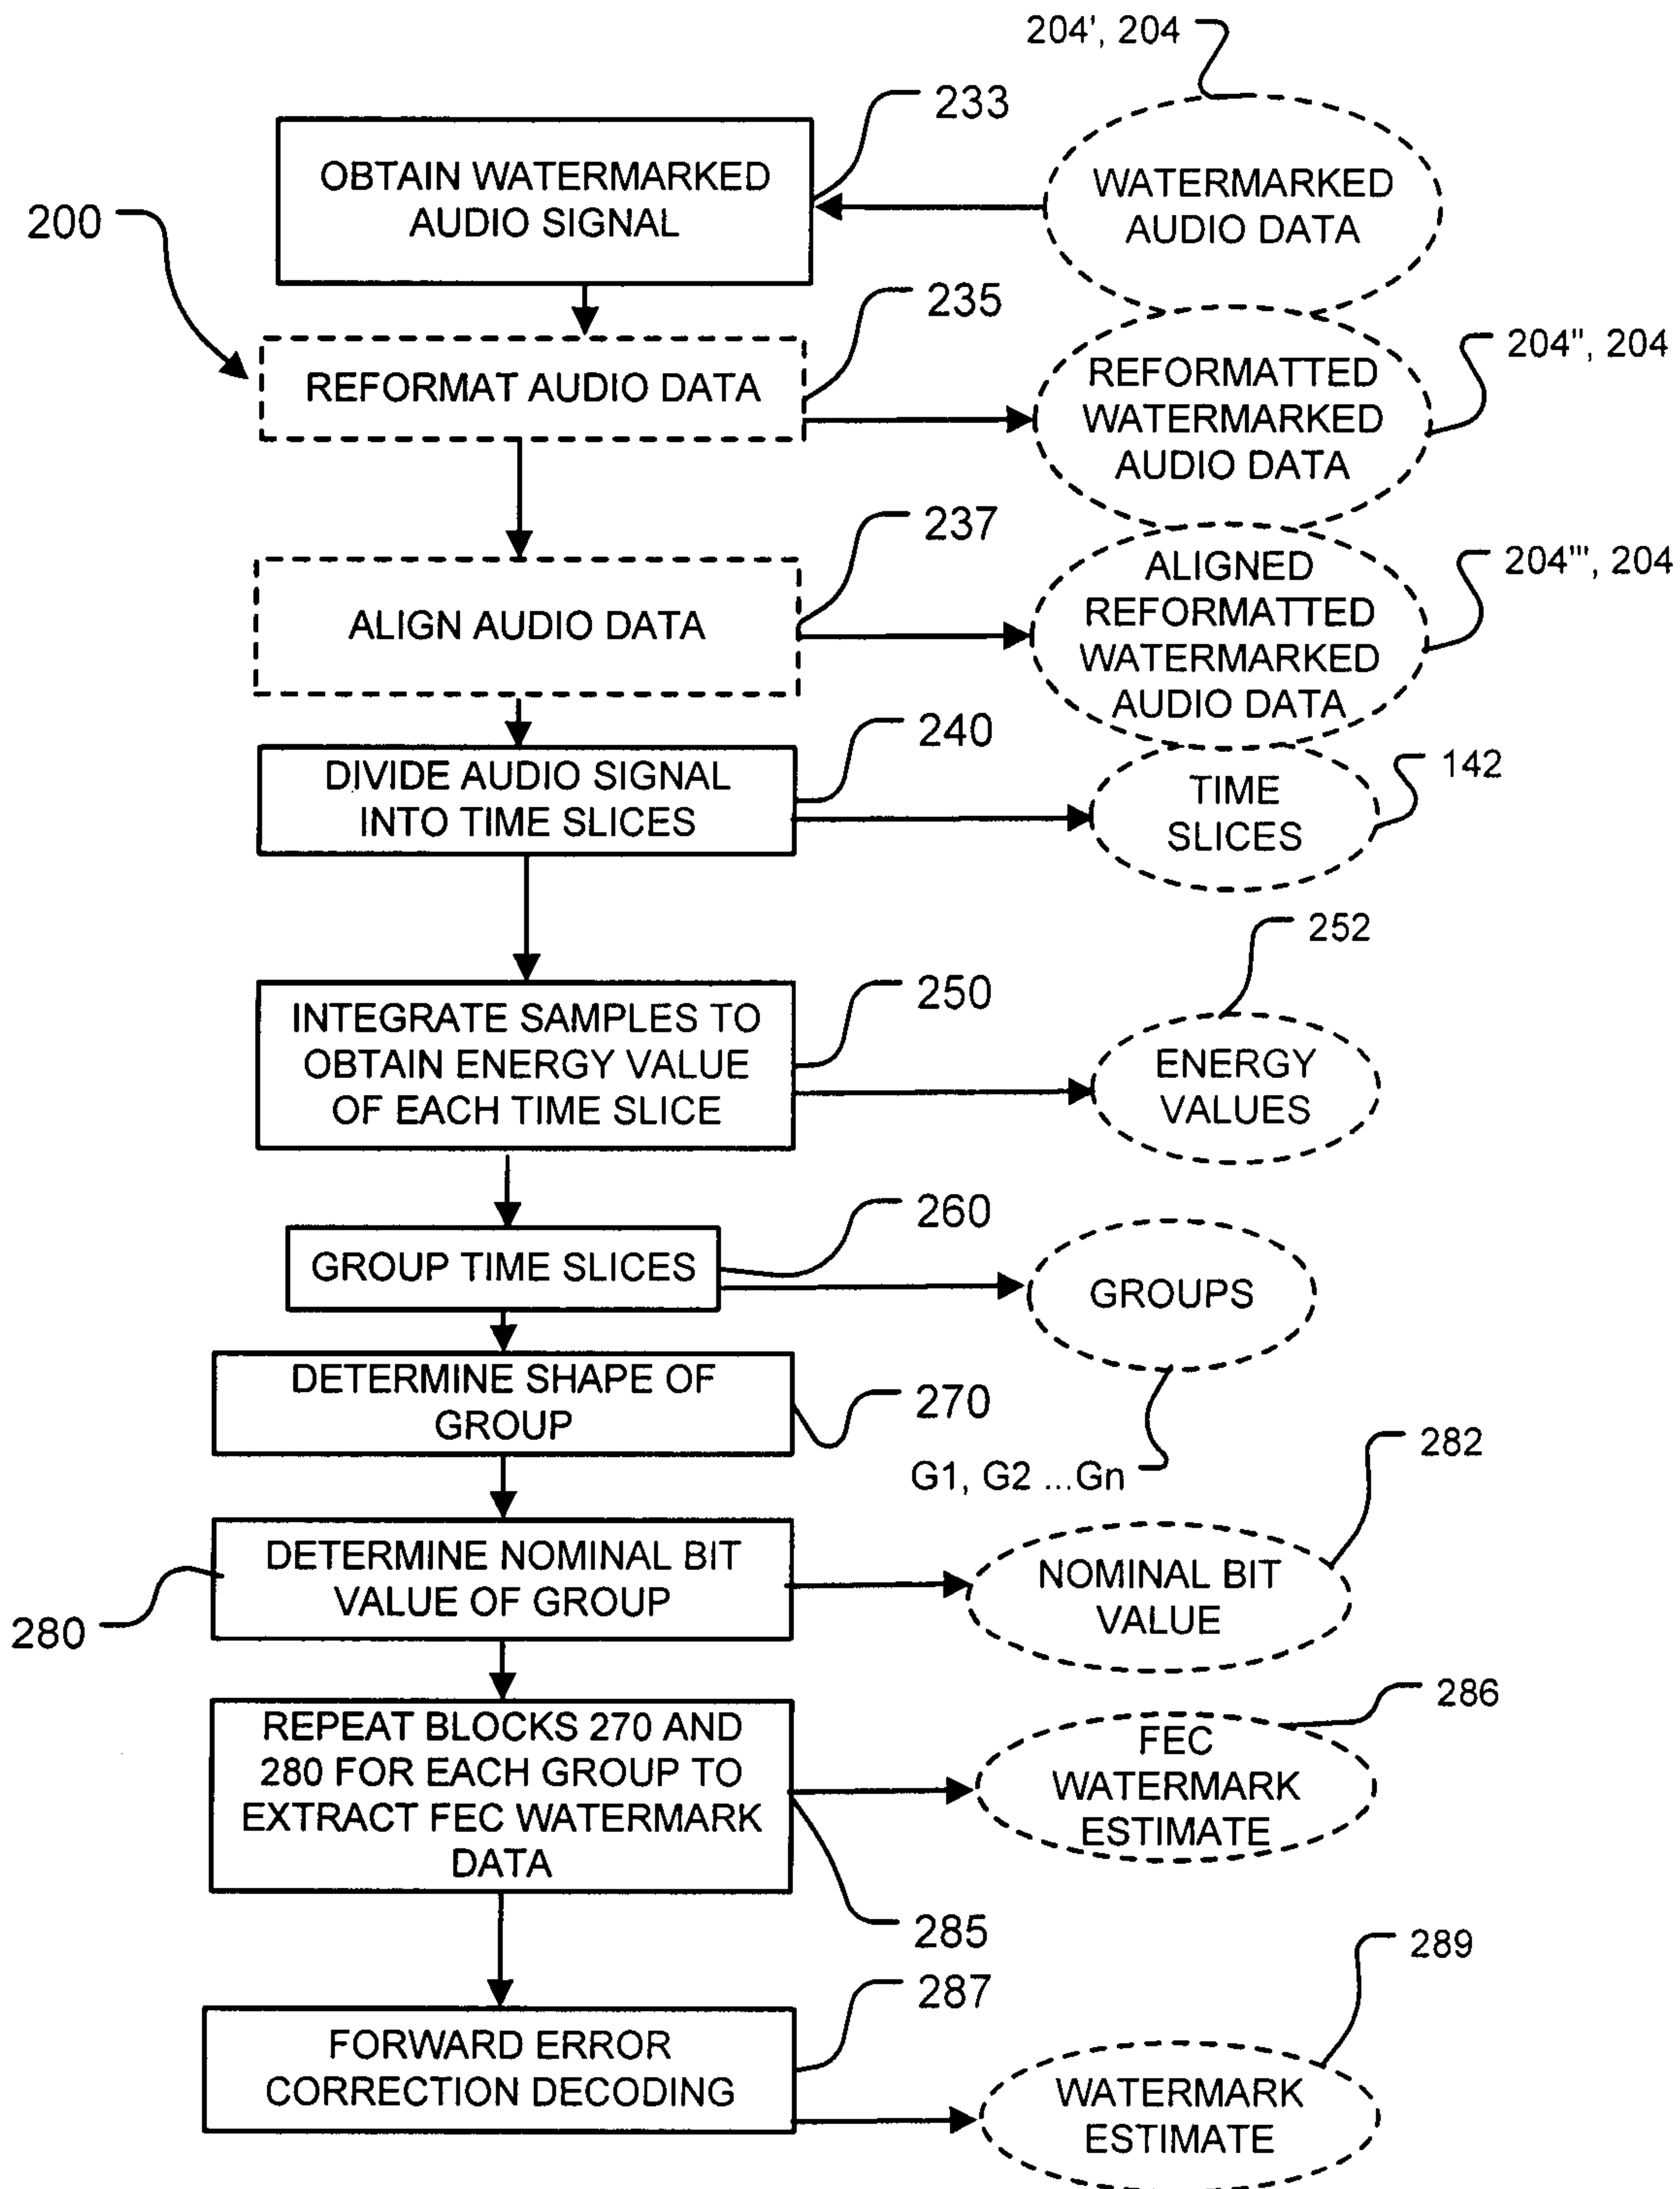


FIGURE 6

6/7




FIGURE 7

7/7

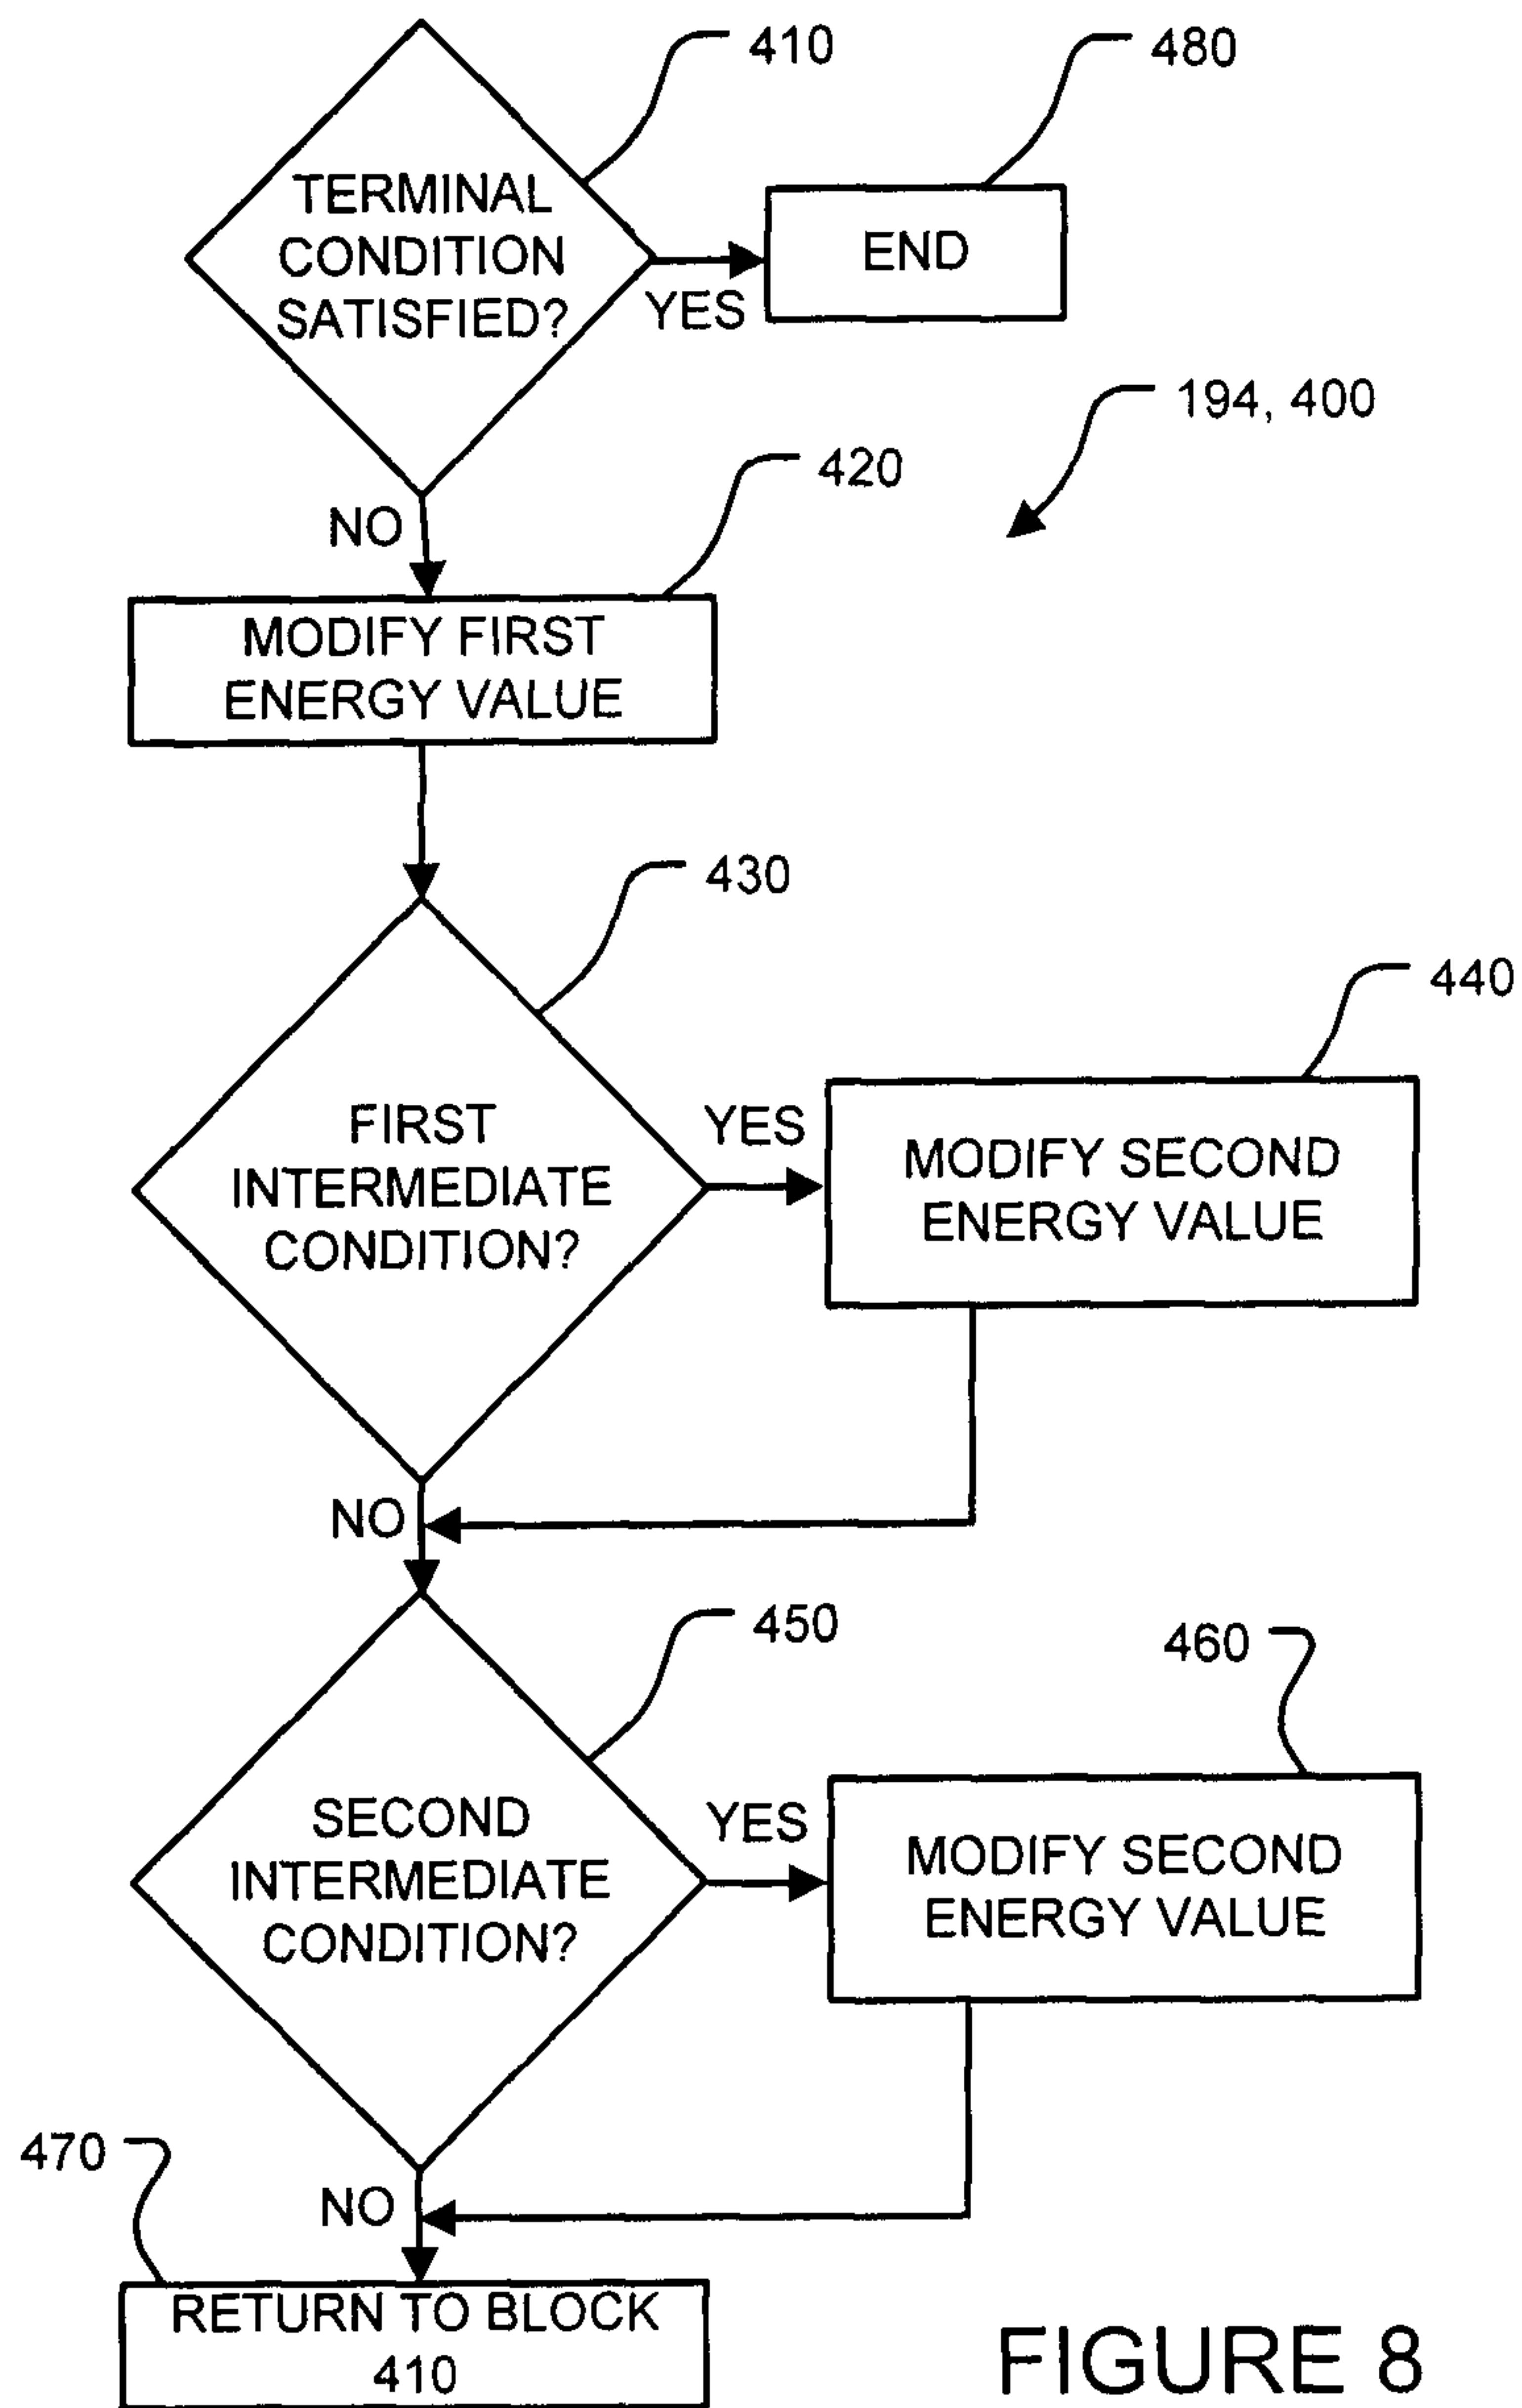


FIGURE 8

