



(22) Date de dépôt/Filing Date: 2003/05/30

(41) Mise à la disp. pub./Open to Public Insp.: 2003/12/24

(45) Date de délivrance/Issue Date: 2015/02/24

(62) Demande originale/Original Application: 2 489 441

(30) Priorité/Priority: 2002/06/17 (US10/174,493)

(51) Cl.Int./Int.Cl. *G10L 19/032* (2013.01)

(72) Inventeurs/Inventors:

TRUMAN, MICHAEL MEAD, US;
DAVIDSON, GRANT ALLEN, US;
FELLERS, MATTHEW CONRAD, US;
VINTON, MARK STUART, US;
WATSON, MATTHEW AUBREY, US;
ROBINSON, CHARLES QUITO, US


(73) Propriétaire/Owner:

DOLBY LABORATORIES LICENSING CORPORATION,
US

(74) Agent: SMART & BIGGAR

(54) Titre : SYSTÈME DE CODAGE AUDIO UTILISANT UN REMPLISSAGE DE TROUS SPECTRAUX

(54) Title: AUDIO CODING SYSTEM USING SPECTRAL HOLE FILLING



(57) Abrégé/Abstract:

Audio coding processes like quantization can cause spectral components of an encoded audio signal to be set to zero, creating spectral holes in the signal. These spectral holes can degrade the perceived quality of audio signals that are reproduced by audio coding systems. An improved decoder avoids or reduces the degradation by filling the spectral holes with synthesized spectral components. An improved encoder may also be used to realize further improvements in the decoder.

73221-74F

ABSTRACT

Audio coding processes like quantization can cause spectral components of an encoded audio signal to be set to zero, creating spectral holes in the signal. These spectral holes can degrade the perceived quality of audio signals that are reproduced by audio coding systems. An improved decoder avoids or reduces the degradation by filling the spectral holes with synthesized spectral components. An improved encoder may also be used to realize further improvements in the decoder.

73221-74F

- 1 -

Audio Coding System Using Spectral Hole Filling

5 This is a divisional application of Canadian National Phase Application
No. 2,489,441, filed on 30th May, 2003.

TECHNICAL FIELD

10 The present invention is related generally to audio coding systems, and is
related more specifically to improving the perceived quality of the audio signals
obtained from audio coding systems.

BACKGROUND ART

15 Audio coding systems are used to encode an audio signal into an encoded
signal that is suitable for transmission or storage, and then subsequently receive or
retrieve the encoded signal and decode it to obtain a version of the original audio
signal for playback. Perceptual audio coding systems attempt to encode an audio
signal into an encoded signal that has lower information capacity requirements than
20 the original audio signal, and then subsequently decode the encoded signal to provide
an output that is perceptually indistinguishable from the original audio signal. One
example of a perceptual audio coding system is described in the Advanced Television
Standards Committee (ATSC) A52 document (1994), which is referred to as Dolby[®]
AC-3. Another example is described in Bosi et al., "ISO/IEC MPEG-2 Advanced
25 Audio Coding." J. AES, vol. 45, no. 10, October 1997, pp. 789-814, which is referred
to as Advanced Audio Coding (AAC). These two coding systems, as well as many
other perceptual coding systems, apply an analysis filterbank to an audio signal to
obtain spectral components that are arranged in groups or frequency bands. The band
widths typically vary and are usually commensurate with widths of the so called
30 critical bands of the human auditory system.

 Perceptual coding systems can be used to reduce the information capacity
requirements of an audio signal while preserving a subjective or perceived measure of
audio quality so that an encoded representation of the audio signal can be conveyed
through a communication channel using less bandwidth or stored on a recording
35 medium using less space. Information capacity requirements are reduced by

- 2 -

quantizing the spectral components. Quantization injects noise into the quantized signal, but perceptual audio coding systems generally use psychoacoustic models in an attempt to control the amplitude of quantization noise so that it is masked or rendered inaudible by spectral components in the signal.

5 The spectral components within a given band are often quantized to the same quantizing resolution and a psychoacoustic model is used to determine the largest minimum quantizing resolution, or the smallest signal-to-noise ratio (SNR), that is possible without injecting an audible level of quantization noise. This technique works fairly well for narrow bands but does not work as well for wider bands when
10 information capacity requirements constrain the coding system to use a relatively coarse quantizing resolution. The larger-valued spectral components in a wide band are usually quantized to a non-zero value having the desired resolution but smaller-valued spectral components in the band are quantized to zero if they have a magnitude that is less than the minimum quantizing level. The number of spectral components in
15 a band that are quantized to zero generally increases as the band width increases, as the difference between the largest and smallest spectral component values within the band increases, and as the minimum quantizing level increases.

Unfortunately, the existence of many quantized-to-zero (QTZ) spectral components in an encoded signal can degrade the perceived quality of the audio
20 signal even if the resulting quantization noise is kept low enough to be deemed inaudible or psychoacoustically masked by spectral components in the signal. This degradation has at least three causes. The first cause is the fact that the quantization noise may not be inaudible because the level of psychoacoustic masking is less than what is predicted by the psychoacoustic model used to determine the quantizing
25 resolution. A second cause is the fact that the creation of many QTZ spectral components can audibly reduce the energy or power of the decoded audio signal as compared to the energy or power of the original audio signal. A third cause is relevant to coding processes that uses distortion-cancellation filterbanks such as the Quadrature Mirror Filter (QMF) or a particular modified Discrete Cosine Transform
30 (DCT) and modified Inverse Discrete Cosine Transform (IDCT) known as Time-Domain Aliasing Cancellation (TDAC) transforms, which are described in Princen et al., "Subband/Transform Coding Using Filter Bank Designs Based on Time Domain Aliasing Cancellation," *ICASSP 1987 Conf. Proc.*, May 1987, pp. 2161-64.

73221-74F

- 3 -

Coding systems that use distortion-cancellation filterbanks such as the QMF or the TDAC transforms use an analysis filterbank in the encoding process that introduces distortion or spurious components into the encoded signal, but use a synthesis filterbank in the decoding process that can, in theory at least, cancel the distortion. In practice, however, the ability of the synthesis filterbank to cancel the distortion can be impaired significantly if the values of one or more spectral components are changed significantly in the encoding process. For this reason, QTZ spectral components may degrade the perceived quality of a decoded audio signal even if the quantization noise is inaudible because changes in spectral component values may impair the ability of the synthesis filterbank to cancel distortion introduced by the analysis filterbank.

Techniques used in known coding systems have provided partial solutions to these problems. Dolby[®] AC-3 and AAC transform coding systems, for example, have some ability to generate an output signal from an encoded signal that retains the signal level of the original audio signal by substituting noise for certain QTZ spectral components in the decoder. In both of these systems, the encoder provides in the encoded signal an indication of power for a frequency band and the decoder uses this indication of power to substitute an appropriate level of noise for the QTZ spectral components in the frequency band. A Dolby[®] AC-3 encoder provides a coarse estimate of the short-term power spectrum that can be used to generate an appropriate level of noise. When all spectral components in a band are set to zero, the decoder fills the band with noise having approximately the same power as that indicated in the coarse estimate of the short-term power spectrum. The AAC coding system uses a technique called Perceptual Noise Substitution (PNS) that explicitly transmits the power for a given band. The decoder uses this information to add noise to match this power. Both systems add noise only in those bands that have no non-zero spectral components.

Unfortunately, these systems do not help preserve power levels in bands that contain a mixture of QTZ and non-zero spectral components. Table 1 shows a hypothetical band of spectral components for an original audio signal, a 3-bit quantized representation of each spectral component that is assembled into an encoded signal, and the corresponding spectral components obtained by a decoder from the encoded signal. The quantized band in the encoded signal has a combination of QTZ and non-zero spectral components.

73221-74F

- 4 -

Original Signal Components	Quantized Components	Dequantized Components
10101010	101	10100000
00000100	000	00000000
00000010	000	00000000
00000001	000	00000000
00011111	000	00000000
00010101	000	00000000
00001111	000	00000000
01010101	010	01000000
11110000	111	11100000

Table 1

5 The first column of the table shows a set of unsigned binary numbers representing spectral components in the original audio signal that are grouped into a single band. The second column shows a representation of the spectral components quantized to three bits. For this example, the portion of each spectral component below the 3-bit resolution has been removed by truncation. The quantized spectral components are transmitted to the decoder and subsequently dequantized by

10 appending zero bits to restore the original spectral component length. The dequantized spectral components are shown in the third column. Because a majority of the spectral components have been quantized to zero, the band of dequantized spectral components contains less energy than the band of original spectral components and

15 that energy is concentrated in a few non-zero spectral components. This reduction in energy can degrade the perceived quality of the decoded signal as explained above.

DISCLOSURE OF INVENTION

20 It is an object of some embodiments of the present invention to improve the perceived quality of audio signals obtained from audio coding systems by avoiding or reducing degradation related to zero-valued quantized spectral components.

73221-74F

- 4a -

According to an aspect of the present invention, there is provided a method for generating audio information, wherein the method comprises: receiving an input signal that conveys an encoded representation of quantized subband signals, wherein spectral components that had a magnitude less than a threshold were quantized to a zero value;

- 5 decoding the encoded representation and identifying a particular subband signal in which one or more spectral components have non-zero values and a plurality of spectral components have a zero value; establishing a scaling envelope that is less than or equal to the threshold using different ways adapted or selected as a function of frequency; generating synthesized spectral components that correspond to the zero-valued spectral components that are scaled
- 10 according to the scaling envelope; generating a modified set of subband signals by substituting the synthesized spectral components for corresponding zero-valued spectral components in the particular subband signal; and generating the audio information by applying a synthesis filterbank to the modified set of subband signals.

- According to another aspect of the present invention, there is provided an
- 15 apparatus for generating audio information, wherein the apparatus comprises means for performing all of the steps in the above described method.

According to another aspect of the present invention, there is provided a medium storing a program of instructions and that is readable by a device for executing the program of instructions to perform all of the steps in the above described method.

- 20 According to another aspect of the present invention, there is provided a method for generating audio information, wherein the method comprises: receiving an input signal that conveys an encoded representation of quantized subband signals, wherein spectral components that had a magnitude less than a threshold were quantized to a zero value;
- 25 decoding the encoded representation and identifying a particular subband signal in which one or more spectral components have non-zero values and a plurality of spectral components have a zero value; deriving a scaling envelope from output of a frequency-domain filter applied to the spectral components, wherein the scaling envelope is less than or equal to the threshold; generating synthesized spectral components that correspond to the zero-valued spectral components that are scaled according to the scaling envelope; generating a modified

73221-74F

- 4b -

set of subband signals by substituting the synthesized spectral components for corresponding zero-valued spectral components in the particular subband signal; and generating the audio information by applying a synthesis filterbank to the modified set of subband signals.

5 According to another aspect of the present invention, there is provided an apparatus for generating audio information, wherein the apparatus comprises means for performing all of the steps in the above described method.

According to another aspect of the present invention, there is provided a medium storing a program of instructions and that is readable by a device for executing the program of instructions to perform all of the steps in the above described method.

10 In another aspect, audio information is provided by receiving an input signal and obtaining therefrom a set of subband signals each having one or more spectral components representing spectral content of an audio signal; identifying within the set of subband signals a particular subband signal in which one or more spectral components have a non-zero value and are quantized by a quantizer having a minimum quantizing level that
15 corresponds to a threshold, and in which a

73221-74F

- 5 -

plurality of spectral components have a zero value; generating synthesized spectral components that correspond to respective zero-valued spectral components in the particular subband signal and that are scaled according to a scaling envelope less than or equal to the threshold; generating a modified set of subband signals by substituting
5 the synthesized spectral components for corresponding zero-valued spectral components in the particular subband signal; and generating the audio information by applying a synthesis filterbank to the modified set of subband signals.

In another aspect, an output signal, preferably an encoded output signal, is provided by generating a set of subband signals each having one or more
10 spectral components representing spectral content of an audio signal by quantizing information that is obtained by applying an analysis filterbank to audio information; identifying within the set of subband signals a particular subband signal in which one or more spectral components have a non-zero value and are quantized by a quantizer having a minimum quantizing level that corresponds to a threshold, and in
15 which a plurality of spectral components have a zero value; deriving scaling control information from the spectral content of the audio signal, wherein the scaling control information controls scaling of synthesized spectral components to be synthesized and substituted for the spectral components having a zero value in a receiver that generates audio information in response to the output signal; and generating the
20 output signal by assembling the scaling control information and information representing the set of subband signals.

The various features of the present invention and its preferred embodiments may be better understood by referring to the following discussion and the accompanying drawings in which like reference numerals refer to like elements in the
25 several figures. The contents of the following discussion and the drawings are set forth as examples only and should not be understood to represent limitations upon the scope of the present invention.

BRIEF DESCRIPTION OF DRAWINGS

30 Fig. 1a is a schematic block diagram of an audio encoder.
Fig. 1b is a schematic block diagram of an audio decoder.
Figs. 2a-2c are graphical illustrations of quantization functions.

Fig. 3 is a graphical schematic illustration of the spectrum of a hypothetical audio signal.

Fig. 4 is a graphical schematic illustration of the spectrum of a hypothetical audio signal with some spectral components set to zero.

5 Fig. 5 is a graphical schematic illustration of the spectrum of a hypothetical audio signal with synthesized spectral components substituted for zero-valued spectral components.

Fig. 6 is a graphical schematic illustration of a hypothetical frequency response for a filter in an analysis filterbank.

10 Fig. 7 is a graphical schematic illustration of a scaling envelope that approximates the roll off of spectral leakage shown in Fig. 6.

Fig. 8 is a graphical schematic illustration of scaling envelopes derived from the output of an adaptable filter.

15 Fig. 9 is a graphical schematic illustration of the spectrum of a hypothetical audio signal with synthesized spectral components weighted by a scaling envelope that approximates the roll off of spectral leakage shown in Fig. 6.

Fig. 10 is a graphical schematic illustration of hypothetical psychoacoustic masking thresholds.

20 Fig. 11 is a graphical schematic illustration of the spectrum of a hypothetical audio signal with synthesized spectral components weighted by a scaling envelope that approximates psychoacoustic masking thresholds.

Fig. 12 is a graphical schematic illustration of a hypothetical subband signal.

Fig. 13 is a graphical schematic illustration of a hypothetical subband signal with some spectral components set to zero.

25 Fig. 14 is a graphical schematic illustration of a hypothetical temporal psychoacoustic masking threshold.

Fig. 15 is a graphical schematic illustration of a hypothetical subband signal with synthesized spectral components weighted by a scaling envelope that approximates temporal psychoacoustic masking thresholds.

30 Fig. 16 is a graphical schematic illustration of the spectrum of a hypothetical audio signal with synthesized spectral components generated by spectral replication.

Fig. 17 is a schematic block diagram of an apparatus that may be used to implement various aspects of the present invention in an encoder or a decoder.

MODES FOR CARRYING OUT THE INVENTION

A. Overview

Various aspects of the present invention may be incorporated into a wide
5 variety of signal processing methods and devices including devices like those
illustrated in Figs. 1a and 1b. Some aspects may be carried out by processing
performed in only a decoding method or device. Other aspects require cooperative
processing performed in both encoding as well as decoding methods or devices. A
description of processes that may be used to carry out these various aspects of the
10 present invention is provided below following an overview of typical devices that
may be used to perform these processes.

1. Encoder

Fig 1a illustrates one implementation of a split-band audio encoder in which
the analysis filterbank 12 receives from the path 11 audio information representing an
15 audio signal and, in response, provides digital information that represents frequency
subbands of the audio signal. The digital information in each of the frequency
subbands is quantized by a respective quantizer 14, 15, 16 and passed to the encoder
17. The encoder 17 generates an encoded representation of the quantized information,
which is passed to the formatter 18. In the particular implementation shown in the
20 figure, the quantization functions in quantizers 14, 15, 16 are adapted in response to
quantizing control information received from the model 13, which generates the
quantizing control information in response to the audio information received from the
path 11. The formatter 18 assembles the encoded representation of the quantized
information and the quantizing control information into an output signal suitable for
25 transmission or storage, and passes the output signal along the path 19.

Many audio applications use uniform linear quantization functions $q(x)$ such
as the 3-bit mid-tread asymmetric quantization function illustrated in Fig. 2a;
however, no particular form of quantization is important to the present invention.
Examples of two other functions $q(x)$ that may be used are shown in Figs. 2b and 2c.
30 In each of these examples, the quantization function $q(x)$ provides an output value
equal to zero for any input value x in the interval from the value at point 30 to the
value at point 31. In many applications, the two values at points 30, 31 are equal in
magnitude and opposite in sign; however, this is not necessary as shown in Fig. 2b.

73221-74F

- 8 -

For ease of discussion, a value x that is within the interval of input values quantized to zero (QTZ) by a particular quantization function $q(x)$ is referred to as being less than the minimum quantizing level of that quantization function.

In this disclosure, terms like "encoder" and "encoding" are not intended to
5 imply any particular type of information processing. For example, encoding is often used to reduce information capacity requirements; however, these terms in this disclosure do not necessarily refer to this type of processing. The encoder 17 may perform essentially any type of processing that is desired. In one implementation, quantized information is encoded into groups of scaled numbers having a common
10 scaling factor. In the Dolby[®] AC-3 coding system, for example, quantized spectral components are arranged into groups or bands of floating-point numbers where the numbers in each band share a floating-point exponent. In the AAC coding system, entropy coding such as Huffman coding is used. In another implementation, the encoder 17 is eliminated and the quantized information is assembled directly into the
15 output signal. No particular type of encoding is important to the present invention.

The model 13 may perform essentially any type processing that may be desired. One example is a process that applies a psychoacoustic model to audio information to estimate the psychoacoustic masking effects of different spectral components in the audio signal. Many variations are possible. For example, the model
20 13 may generate the quantizing control information in response to the frequency subband information available at the output of the analysis filterbank 12 instead of, or in addition to, the audio information available at the input of the filterbank. As another example, the model 13 may be eliminated and quantizers 14, 15, 16 use quantization functions that are not adapted. No particular modeling process is
25 important to the present invention.

2. Decoder

Fig 1b illustrates one implementation of a split-band audio decoder in which the deformatter 22 receives from the path 21 an input signal conveying an encoded representation of quantized digital information representing frequency subbands of an
30 audio signal. The deformatter 22 obtains the encoded representation from the input signal and passes it to the decoder 23. The decoder 23 decodes the encoded representation into frequency subbands of quantized information. The quantized digital information in each of the frequency subbands is dequantized by a respective

- 9 -

dequantizer 25, 26, 27 and passed to the synthesis filterbank 28, which generates along the path 29 audio information representing an audio signal. In the particular implementation shown in the figure, the dequantization functions in the dequantizers 25, 26, 27 are adapted in response to quantizing control information received from the model 24, which generates the quantizing control information in response to control information obtained by the deformatter 22 from the input signal.

In this disclosure, terms like "decoder" and "decoding" are not intended to imply any particular type of information processing. The decoder 23 may perform essentially any type of processing that is needed or desired. In one implementation that is inverse to an encoding process described above, quantized information in groups of floating-point numbers having shared exponents are decoded into individual quantized components that do not share exponents. In another implementation, entropy decoding such as Huffman decoding is used. In another implementation, the decoder 23 is eliminated and the quantized information is obtained directly by the deformatter 22. No particular type of decoding is important to the present invention.

The model 24 may perform essentially any type of processing that may be desired. One example is a process that applies a psychoacoustic model to information obtained from the input signal to estimate the psychoacoustic masking effects of different spectral components in an audio signal. As another example, the model 24 is eliminated and dequantizers 25, 26, 27 may either use quantization functions that are not adapted or they may use quantization functions that are adapted in response to quantizing control information obtained directly from the input signal by the deformatter 22. No particular process is important to the present invention.

3. Filterbanks

The devices illustrated in Figs. 1a and 1b show components for three frequency subbands. Many more subbands are used in a typical application but only three are shown for illustrative clarity. No particular number is important in principle to the present invention.

The analysis and synthesis filterbanks may be implemented in essentially any way that is desired including a wide range of digital filter technologies, block transforms and wavelet transforms. In one audio coding system having an encoder and a decoder like those discussed above, the analysis filterbank 12 is implemented by the TDAC modified DCT and the synthesis filterbank 28 is implemented by the

- 10 -

TDAC modified IDCT mentioned above; however, no particular implementation is important in principle.

Analysis filterbanks that are implemented by block transforms split a block or interval of an input signal into a set of transform coefficients that represent the spectral content of that interval of signal. A group of one or more adjacent transform coefficients represents the spectral content within a particular frequency subband having a bandwidth commensurate with the number of coefficients in the group.

Analysis filterbanks that are implemented by some type of digital filter such as a polyphase filter, rather than a block transform, split an input signal into a set of subband signals. Each subband signal is a time-based representation of the spectral content of the input signal within a particular frequency subband. Preferably, the subband signal is decimated so that each subband signal has a bandwidth that is commensurate with the number of samples in the subband signal for a unit interval of time.

The following discussion refers more particularly to implementations that use block transforms like the TDAC transform mentioned above. In this discussion, the term "subband signal" refers to groups of one or more adjacent transform coefficients and the term "spectral components" refers to the transform coefficients. Principles of the present invention may be applied to other types of implementations, however, so the term "subband signal" generally may be understood to refer also to a time-based signal representing spectral content of a particular frequency subband of a signal, and the term "spectral components" generally may be understood to refer to samples of a time-based subband signal.

4. Implementation

Various aspects of the present invention may be implemented in a wide variety of ways including software in a general-purpose computer system or in some other apparatus that includes more specialized components such as digital signal processor (DSP) circuitry coupled to components similar to those found in a general-purpose computer system. Fig. 17 is a block diagram of device 70 that may be used to implement various aspects of the present invention in an audio encoder or audio decoder. DSP 72 provides computing resources. RAM 73 is system random access memory (RAM) used by DSP 72 for signal processing. ROM 74 represents some form of persistent storage such as read only memory (ROM) for storing programs needed to

WO 03/107328

PCT/US03/17078

- 11 -

operate device 70 and to carry out various aspects of the present invention. I/O control 75 represents interface circuitry to receive and transmit signals by way of communication channels 76, 77. Analog-to-digital converters and digital-to-analog converters may be included in I/O control 75 as desired to receive and/or transmit analog audio signals. In the embodiment shown, all major system components connect to bus 71, which may represent more than one physical bus; however, a bus architecture is not required to implement the present invention.

In embodiments implemented in a general purpose computer system, additional components may be included for interfacing to devices such as a keyboard or mouse and a display, and for controlling a storage device having a storage medium such as magnetic tape or disk, or an optical medium. The storage medium may be used to record programs of instructions for operating systems, utilities and applications, and may include embodiments of programs that implement various aspects of the present invention.

The functions required to practice various aspects of the present invention can be performed by components that are implemented in a wide variety of ways including discrete logic components, one or more ASICs and/or program-controlled processors. The manner in which these components are implemented is not important to the present invention.

Software implementations of the present invention may be conveyed by a variety of machine readable media such as baseband or modulated communication paths throughout the spectrum including from supersonic to ultraviolet frequencies, or storage media including those that convey information using essentially any magnetic or optical recording technology including magnetic tape, magnetic disk, and optical disc. Various aspects can also be implemented in various components of computer system 70 by processing circuitry such as ASICs, general-purpose integrated circuits, microprocessors controlled by programs embodied in various forms of ROM or RAM, and other techniques.

B. Decoder

Various aspects of the present invention may be carried out in a decoder that do not require any special processing or information from an encoder. These aspects are described in this section of the disclosure. Other aspects that do require special processing or information from an encoder are described in the following section.

1. Spectral Holes

Fig. 3 is a graphical illustration of the spectrum of an interval of a hypothetical audio signal that is to be encoded by a transform coding system. The spectrum 41 represents an envelope of the magnitude of transform coefficients or spectral components. During the encoding process, all spectral components having a magnitude less than the threshold 40 are quantized to zero. If a quantization function such as the function $q(x)$ shown in Fig. 2a is used, the threshold 40 corresponds to the minimum quantizing levels 30, 31. The threshold 40 is shown with a uniform value across the entire frequency range for illustrative convenience. This is not typical in many coding systems. In perceptual audio coding systems that uniformly quantize spectral components within each subband signal, for example, the threshold 40 is uniform within each frequency subband but it varies from subband to subband. In other implementations, the threshold 40 may also vary within a given frequency subband.

Fig. 4 is a graphical illustration of the spectrum of the hypothetical audio signal that is represented by quantized spectral components. The spectrum 42 represents an envelope of the magnitude of spectral components that have been quantized. The spectrum shown in this figure as well as in other figures does not show the effects of quantizing the spectral components having magnitudes greater than or equal to the threshold 40. The difference between the QTZ spectral components in the quantized signal and the corresponding spectral components in the original signal are shown with hatching. These hatched areas represent "spectral holes" in the quantized representation that are to be filled with synthesized spectral components.

In one implementation of the present invention, a decoder receives an input signal that conveys an encoded representation of quantized subband signals such as that shown in Fig. 4. The decoder decodes the encoded representation and identifies those subband signals in which one or more spectral components have non-zero values and a plurality of spectral components have a zero value. Preferably, the frequency extents of all subband signals are either known *a priori* to the decoder or they are defined by control information in the input signal. The decoder generates synthesized spectral components that correspond to the zero-valued spectral components using a process such as those described below. The synthesized components are scaled according to a scaling envelope that is less than or equal to the

- 13 -

threshold 40, and the scaled synthesized spectral components are substituted for the zero-valued spectral components in the subband signal. The decoder does not require any information from the encoder that explicitly indicates the level of the threshold 40 if the minimum quantizing levels 30, 31 of the quantization function $q(x)$ used to
5 quantize the spectral components is known.

2. Scaling

The scaling envelope may be established in a wide variety of ways. A few ways are described below. More than one way may be used. For example, a composite scaling envelope may be derived that is equal to the maximum of all
10 envelopes obtained from multiple ways, or by using different ways to establish upper and/or lower bounds for the scaling envelope. The ways may be adapted or selected in response to characteristics of the encoded signal, and they can be adapted or selected as a function of frequency.

a) Uniform Envelope

15 One way is suitable for decoders in audio transform coding systems and in systems that use other filterbank implementations. This way establishes a uniform scaling envelope by setting it equal to the threshold 40. An example of such a scaling envelope is shown in Fig. 5, which uses hatched areas to illustrate the spectral holes that are filled with synthesized spectral components. The spectrum 43 represents an
20 envelope of the spectral components of an audio signal with spectral holes filled by synthesized spectral components. The upper bounds of the hatched areas shown in this figure as well as in later figures do not represent the actual levels of the synthesized spectral components themselves but merely represents a scaling envelope for the synthesized components. The synthesized components that are used to fill
25 spectral holes have spectral levels that do not exceed the scaling envelope.

b) Spectral Leakage

A second way for establishing a scaling envelope is well suited for decoders in audio coding systems that use block transforms, but it is based on principles that may be applied to other types of filterbank implementations. This way provides a non-
30 uniform scaling envelope that varies according to spectral leakage characteristics of the prototype filter frequency response in a block transform.

The response 50 shown in Fig. 6 is a graphical illustration of a hypothetical frequency response for a transform prototype filter showing spectral leakage between

coefficients. The response includes a main lobe, usually referred to as the passband of the prototype filter, and a number of side lobes adjacent to the main lobe that diminish in level for frequencies farther away from the center of the passband. The side lobes represent spectral energy that leaks from the passband into adjacent frequency bands.

5 The rate at which the level of these side lobes decrease is referred to as the rate of roll off of the spectral leakage.

The spectral leakage characteristics of a filter impose constraints on the spectral isolation between adjacent frequency subbands. If a filter has a large amount of spectral leakage, spectral levels in adjacent subbands cannot differ as much as they
10 can for filters with lower amounts of spectral leakage. The envelope 51 shown in Fig. 7 approximates the roll off of spectral leakage shown in Fig. 6. Synthesized spectral components may be scaled to such an envelope or, alternatively, this envelope may be used as a lower bound for a scaling envelope that is derived by other techniques.

15 The spectrum 44 in Fig. 9 is a graphical illustration of the spectrum of a hypothetical audio signal with synthesized spectral components that are scaled according to an envelope that approximates spectral leakage roll off. The scaling envelope for spectral holes that are bounded on each side by spectral energy is a composite of two individual envelopes, one for each side. The composite is formed by
20 taking the larger of the two individual envelopes.

c) Filter

A third way for establishing a scaling envelope is also well suited for decoders in audio coding systems that use block transforms, but it is also based on principles that may be applied to other types of filterbank implementations. This way provides a
25 non-uniform scaling envelope that is derived from the output of a frequency-domain filter that is applied to transform coefficients in the frequency domain. The filter may be a prediction filter, a low pass filter, or essentially any other type of filter that provides the desired scaling envelope. This way usually requires more computational resources than are required for the two ways described above, but it allows the scaling
30 envelope to vary as a function of frequency.

Fig. 8 is a graphical illustration of two scaling envelopes derived from the output of an adaptable frequency-domain filter. For example, the scaling envelope 52 could be used for filling spectral holes in signals or portions of signals that are

- 15 -

deemed to be more tone like, and the scaling envelope 53 could be used for filling spectral holes in signals or portions of signals that are deemed to be more noise like. Tone and noise properties of a signal can be assessed in a variety of ways. Some of these ways are discussed below. Alternatively, the scaling envelope 52 could be used
5 for filling spectral holes at lower frequencies where audio signals are often more tone like and the scaling envelope 53 could be used for filling spectral holes at higher frequencies where audio signal are often more noise like.

d) Perceptual Masking

A fourth way for establishing a scaling envelope is applicable to decoders in
10 audio coding systems that implement filterbanks with block transforms and other types of filters. This way provides a non-uniform scaling envelope that varies according to estimated psychoacoustic masking effects.

Fig. 10 illustrates two hypothetical psychoacoustic masking thresholds. The threshold 61 represents the psychoacoustic masking effects of a lower-frequency
15 spectral component 60 and the threshold 64 represents the psychoacoustic masking effects of a higher-frequency spectral component 63. Masking thresholds such as these may be used to derive the shape of the scaling envelope.

The spectrum 45 in Fig. 11 is a graphical illustration of the spectrum of a hypothetical audio signal with substitute synthesized spectral components that are
20 scaled according to envelopes that are based on psychoacoustic masking. In the example shown, the scaling envelope in the lowest-frequency spectral hole is derived from the lower portion of the masking threshold 61. The scaling envelope in the central spectral hole is a composite of the upper portion of the masking threshold 61 and the lower portion of the masking threshold 64. The scaling envelope in the
25 highest-frequency spectral hole is derived from the upper portion of the masking threshold 64.

e) Tonality

A fifth way for establishing a scaling envelope is based on an assessment of the tonality of the entire audio signal or some portion of the signal such as for one or
30 more subband signals. Tonality can be assessed in a number of ways including the calculation of a Spectral Flatness Measure, which is a normalized quotient of the arithmetic mean of signal samples divided by the geometric mean of the signal samples. A value close to one indicates a signal is very noise like, and a value close to

- 16 -

zero indicates a signal is very tone like. SFM can be used directly to adapt the scaling envelope. When the SFM is equal to zero, no synthesized components are used to fill a spectral hole. When the SFM is equal to one, the maximum permitted level of synthesized components is used to fill a spectral hole. In general, however, an encoder
5 is able to calculate a better SFM because it has access to the entire original audio signal prior to encoding. It is likely that a decoder will not calculate an accurate SFM because of the presence of QTZ spectral components.

A decoder can also assess tonality by analyzing the arrangement or distribution of the non-zero-valued and the zero-valued spectral components. In one
10 implementation, a signal is deemed to be more tone like rather than noise like if long runs of zero-valued spectral components are distributed between a few large non-zero-valued components because this arrangement implies a structure of spectral peaks.

In yet another implementation, a decoder applies a prediction filter to one or
15 more subband signals and determines the prediction gain. A signal is deemed to be more tone like as the prediction gain increases.

f) Temporal Scaling

Fig. 12 is a graphical illustration of a hypothetical subband signal that is to be encoded. The line 46 represents a temporal envelope of the magnitude of spectral
20 components. This subband signal may be composed of a common spectral component or transform coefficient in a sequence of blocks obtained from an analysis filterbank implemented by a block transform, or it may be a subband signal obtained from another type of analysis filterbank implemented by a digital filter other than a block transform such as a QMF. During the encoding process, all spectral components
25 having a magnitude less than the threshold 40 are quantized to zero. The threshold 40 is shown with a uniform value across the entire time interval for illustrative convenience. This is not typical in many coding systems that use filterbanks implemented by block transforms.

Fig. 13 is a graphical illustration of the hypothetical subband signal that is
30 represented by quantized spectral components. The line 47 represents a temporal envelope of the magnitude of spectral components that have been quantized. The line shown in this figure as well as in other figures does not show the effects of quantizing the spectral components having magnitudes greater than or equal to the threshold 40.

- 17 -

The difference between the QTZ spectral components in the quantized signal and the corresponding spectral components in the original signal are shown with hatching. The hatched area represents a spectral hole within an interval of time that are is to be filled with synthesized spectral components.

- 5 In one implementation of the present invention, a decoder receives an input signal that conveys an encoded representation of quantized subband signals such as that shown in Fig. 13. The decoder decodes the encoded representation and identifies those subband signals in which a plurality of spectral components have a zero value and are preceded and/or followed by spectral components having non-zero values.
- 10 The decoder generates synthesized spectral components that correspond to the zero-valued spectral components using a process such as those described below. The synthesized components are scaled according to a scaling envelope. Preferably, the scaling envelope accounts for the temporal masking characteristics of the human auditory system.
- 15 Fig. 14 illustrates a hypothetical temporal psychoacoustic masking threshold. The threshold 68 represents the temporal psychoacoustic masking effects of a spectral component 67. The portion of the threshold to the left of the spectral component 67 represents pre-temporal masking characteristics, or masking that precedes the occurrence of the spectral component. The portion of the threshold to the right of the
- 20 spectral component 67 represents post-temporal masking characteristics, or masking that follows the occurrence of the spectral component. Post-masking effects generally have a duration that is much longer than the duration of pre-masking effects. A temporal masking threshold such as this may be used to derive a temporal shape of the scaling envelope.
- 25 The line 48 in Fig. 15 is a graphical illustration of a hypothetical subband signal with substitute synthesized spectral components that are scaled according to envelopes that are based on temporal psychoacoustic masking effects. In the example shown, the scaling envelope is a composite of two individual envelopes. The individual envelope for the lower-frequency part of the spectral hole is derived from
- 30 the post-masking portion of the threshold 68. The individual envelope for the higher-frequency part of the spectral hole is derived from the pre-masking part of the threshold 68.

3. Generation of Synthesized Components

The synthesized spectral components may be generated in a variety of ways. Two ways are described below. Multiple ways may be used. For example, different ways may be selected in response to characteristics of the encoded signal or as a function of frequency.

A first way generates a noise-like signal. Essentially any of a wide variety of ways for generating pseudo-noise signals may be used.

A second way uses a technique called spectral translation or spectral replication that copies spectral components from one or more frequency subbands. Lower-frequency spectral components are usually copied to fill spectral holes at higher frequencies because higher frequency components are often related in some manner to lower frequency components. In principle, however, spectral components may be copied to higher or lower frequencies.

The spectrum 49 in Fig. 16 is a graphical illustration of the spectrum of a hypothetical audio signal with synthesized spectral components generated by spectral replication. A portion of the spectral peak is replicated down and up in frequency multiple times to fill the spectral holes at the low and middle frequencies, respectively. A portion of the spectral components near the high end of the spectrum are replicated up in frequency to fill the spectral hole at the high end of the spectrum. In the example shown, the replicated components are scaled by a uniform scaling envelope; however, essentially any form of scaling envelope may be used.

C. Encoder

The aspects of the present invention that are described above can be carried out in a decoder without requiring any modification to existing encoders. These aspects can be enhanced if the encoder is modified to provide additional control information that otherwise would not be available to the decoder. The additional control information can be used to adapt the way in which synthesized spectral components are generated and scaled in the decoder.

1. Control Information

An encoder can provide a variety of scaling control information, which a decoder can use to adapt the scaling envelope for synthesized spectral components. Each of the examples discussed below can be provided for an entire signal and/or for frequency subbands of the signal.

If a subband contains spectral components that are significantly below the minimum quantizing level, the encoder can provide information to the decoder that indicates this condition. The information may be a type of index that a decoder can use to select from two or more scaling levels, or the information may convey some measure of spectral level such as average or root-mean-square (RMS) power. The decoder can adapt the scaling envelope in response to this information.

As explained above, a decoder can adapt the scaling envelope in response to psychoacoustic masking effects estimated from the encoded signal itself; however, it is possible for the encoder to provide a better estimate of these masking effects when the encoder has access to features of the signal that are lost by an encoding process. This can be done by having the model 13 provide psychoacoustic information to the formatter 18 that is otherwise not available from the encoded signal. Using this type of information, the decoder is able to adapt the scaling envelope to shape the synthesized spectral components according to one or more psychoacoustic criteria.

The scaling envelope can also be adapted in response to some assessment of the noise-like or tone-like qualities of a signal or subband signal. This assessment can be done in several ways by either the encoder or the decoder; however, an encoder is usually able to make a better assessment. The results of this assessment can be assembled with the encoded signal. One assessment is the SFM described above.

An indication of SFM can also be used by a decoder to select which process to use for generating synthesized spectral components. If the SFM is close to one, the noise-generation technique can be used. If the SFM is close to zero, the spectral replication technique can be used.

An encoder can provide some indication of power for the non-zero and the QTZ spectral components such as a ratio of these two powers. The decoder can calculate the power of the non-zero spectral components and then use this ratio or other indication to adapt the scaling envelope appropriately.

2. Zero Spectral Coefficients

The previous discussion has sometimes referred to zero-valued spectral components as QTZ (quantized-to-zero) components because quantization is a common source of zero-valued components in an encoded signal. This is not essential. The value of spectral components in an encoded signal may be set to zero by essentially any process. For example, an encoder may identify the largest one or two

WO 03/107328

PCT/US03/17078

- 20 -

spectral components in each subband signal above a particular frequency and set all other spectral components in those subband signals to zero. Alternatively, an encoder may set to zero all spectral components in certain subbands that are less than some threshold. A decoder that incorporates various aspects of the present invention as

5 described above is able to fill spectral holes regardless of the process that is responsible for creating them.

73221-74F

- 21 -

CLAIMS:

1. A method for generating audio information, wherein the method comprises:

receiving an input signal that conveys an encoded representation of quantized subband signals, wherein spectral components that had a magnitude less than a threshold were
5 quantized to a zero value;

decoding the encoded representation and identifying a particular subband signal in which one or more spectral components have non-zero values and a plurality of spectral components have a zero value;

10 establishing a scaling envelope that is less than or equal to the threshold using different ways adapted or selected as a function of frequency;

generating synthesized spectral components that correspond to the zero-valued spectral components that are scaled according to the scaling envelope;

15 generating a modified set of subband signals by substituting the synthesized spectral components for corresponding zero-valued spectral components in the particular subband signal; and

generating the audio information by applying a synthesis filterbank to the modified set of subband signals.

2. The method according to claim 1 wherein the synthesis filterbank is implemented by a block transform that has spectral leakage between adjacent spectral
20 components and the scaling envelope varies at a rate substantially equal to a rate of roll off of the spectral leakage of the block transform.

3. The method according to claim 1 or 2 wherein the synthesis filterbank is implemented by a block transform and the method comprises:

73221-74F

- 22 -

applying a frequency-domain filter to one or more spectral components in a set of the quantized subband signals; and

deriving the scaling envelope from an output of the frequency-domain filter.

4. The method according to claim 3 that comprises varying the response of the
5 frequency-domain filter as a function of frequency.

5. The method according to any one of claims 1 through 4 that comprises:

obtaining a measure of tonality of an audio signal represented by a set of the quantized subband signals; and

adapting the scaling envelope in response to the measure of tonality.

10 6. The method according to claim 5 that obtains the measure of tonality from the input signal.

7. The method according to claim 5 that comprises deriving the measure of tonality from the way in which the zero-valued spectral components are arranged in the particular subband signal.

15 8. The method according to any one of claims 1 through 7 wherein the synthesis filterbank is implemented by a block transform and the method comprises:

obtaining a sequence of sets of subband signals from the input signal;

identifying a common subband signal in the sequence of sets of subband signals where, for each set in the sequence, one or more spectral components have a non-zero
20 value and a plurality of spectral components have a zero value;

identifying a common spectral component within the common subband signal that has a zero value in a plurality of adjacent sets in the sequence that are either preceded or followed by a set with the common spectral components having a non-zero value;

73221-74F

- 23 -

scaling the synthesized spectral components that correspond to the zero-valued common spectral components according to the scaling envelope that varies from set to set in the sequence according to temporal masking characteristics of the human auditory system;

5 generating a sequence of modified sets of subband signals by substituting the synthesized spectral components for the corresponding zero-valued common spectral components in the sets; and

generating the audio information by applying the synthesis filterbank to the sequence of modified sets of subband signals.

9. The method according to any one of claims 1 through 8 wherein the synthesis
10 filterbank is implemented by a block transform and the method generates the synthesized spectral components by spectral translation of other spectral components in a set of the quantized subband signals.

10. The method according to any one of claims 1 through 9 wherein the scaling envelope varies according to temporal masking characteristics of the human auditory system.

15 11. The method according to any one of claims 1 through 10 that establishes the scaling envelope to vary according to estimated psychoacoustic masking effects.

12. An apparatus for generating audio information, wherein the apparatus comprises means for performing all of the steps in the method of any one of claims 1 through 11.

13. A medium storing a program of instructions and that is readable by a device for
20 executing the program of instructions to perform all of the steps in the method of any one of claims 1 through 11.

73221-74F

- 24 -

14. A method for generating audio information, wherein the method comprises:

receiving an input signal that conveys an encoded representation of quantized subband signals, wherein spectral components that had a magnitude less than a threshold were quantized to a zero value;

5 decoding the encoded representation and identifying a particular subband signal in which one or more spectral components have non-zero values and a plurality of spectral components have a zero value;

deriving a scaling envelope from output of a frequency-domain filter applied to the spectral components, wherein the scaling envelope is less than or equal to the threshold;

10 generating synthesized spectral components that correspond to the zero-valued spectral components that are scaled according to the scaling envelope;

generating a modified set of subband signals by substituting the synthesized spectral components for corresponding zero-valued spectral components in the particular subband signal; and

15 generating the audio information by applying a synthesis filterbank to the modified set of subband signals.

15. The method according to claim 14 that comprises varying the response of the frequency-domain filter as a function of frequency.

16. The method according to claim 14 or 15 that comprises generating the
20 synthesized spectral components by generating a noise-like signal.



17. The method according to any one of claims 14 through 16 that comprises generating the synthesized spectral components by multiple ways selected as a function of frequency.




73221-74F




- 25 -




18. An apparatus for generating audio information, wherein the apparatus comprises means for performing all steps in the method according to any one of claims 14 through 17.




19. A medium storing a program of instructions and that is readable by an apparatus
5 for executing the program of instructions to perform all steps in the method according to any one of claims 14 through 17.




**Fig. 1a****Fig. 1b**






**Fig. 3****Fig. 4****Fig. 5**

**Fig. 6****Fig. 7****Fig. 8**

**Fig. 9****Fig. 10****Fig. 11**

**Fig. 12****Fig. 13****Fig. 14**

**Fig. 15****Fig. 16**

70 ↗

Fig. 17

