



(86) Date de dépôt PCT/PCT Filing Date: 2002/04/25

(87) Date publication PCT/PCT Publication Date: 2002/11/21

(85) Entrée phase nationale/National Entry: 2003/10/27

(86) N° demande PCT/PCT Application No.: US 2002/012957

(87) N° publication PCT/PCT Publication No.: 2002/093560

(30) Priorité/Priority: 2001/05/10 (60/290,286) US

(51) Cl.Int.⁷/Int.Cl.⁷ G10L 19/14


(71) Demandeur/Applicant:
DOLBY LABORATORIES LICENSING CORPORATION,
US

(72) Inventeur/Inventor:
CROCKETT, BRETT G., US

(74) Agent: SMART & BIGGAR

(54) Titre : AMELIORATIONS APPORTEES A LA PERFORMANCE TRANSITOIRE DE SYSTEMES DE CODAGE AUDIO
A FAIBLE DEBIT BINAIRE PAR REDUCTION DU PRE-BRUIT

(54) Title: IMPROVING TRANSIENT PERFORMANCE OF LOW BIT RATE AUDIO CODING SYSTEMS BY REDUCING
PRE-NOISE



(57) Abrégé/Abstract:

Distortion artifacts preceding a signal transient in an audio signal stream processed by a transform-based low-bit-rate audio coding system employing coding blocks are reduced by detecting a transient in the audio signal stream and shifting the temporal relationship of the transient with respect to the coding blocks such that the time duration of the distortion artifacts is reduced. The audio data is time scaled in such a way that the transients are temporally repositioned prior to quantization in a transform-based low-bit-rate audio encoder so as to reduce the amount of pre-noise in the decoded audio signal. Alternatively, or in addition, in a transform-based low-bit-rate audio coding system, a transient in the audio signal stream is detected and a portion of the distortion artifacts are time compressed such that the time duration of the distortion artifacts is reduced.

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

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
21 November 2002 (21.11.2002)

PCT

(10) International Publication Number
WO 02/093560 A1

(51) International Patent Classification⁷: **G10L 19/14**

G. [US/US]; 100 Potrero Avenue, San Francisco, CA 94103 (US).

(21) International Application Number: PCT/US02/12957

(22) International Filing Date: 25 April 2002 (25.04.2002)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
60/290,286 10 May 2001 (10.05.2001) US

(71) Applicant (for all designated States except US): **DOLBY LABORATORIES LICENSING CORPORATION** [US/US]; 100 Potrero Avenue, San Francisco, CA 94103 (US).

(72) Inventor; and

(75) Inventor/Applicant (for US only): **CROCKETT, Brett,**


(74) Agents: **GALLAGHER, Thomas, A.** et al.; Gallagher & Lathrop, Suite 1111, 601 California Street, San Francisco, CA 94108-2805 (US).

(81) Designated States (*national*): AE, AG, AL, AM, AT (utility model), AT, AU, AZ, BA, BB, BG, BR, BY, BZ, CA, CH, CN, CO, CR, CU, CZ (utility model), CZ, DE (utility model), DE, DK (utility model), DK, DM, DZ, EC, EE (utility model), EE, ES, FI (utility model), FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, MZ, NO, NZ, OM, PH, PL, PT, RO, RU, SD, SE, SG, SI, SK (utility model), SK, SL, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VN, YU, ZA, ZM, ZW.

(84) Designated States (*regional*): ARIPO patent (GH, GM, KE, LS, MW, MZ, SD, SL, SZ, TZ, UG, ZM, ZW),

[Continued on next page]

(54) Title: IMPROVING TRANSIENT PERFORMANCE OF LOW BIT RATE AUDIO CODING SYSTEMS BY REDUCING PRE-NOISE



(57) Abstract: Distortion artifacts preceding a signal transient in an audio signal stream processed by a transform-based low-bit-rate audio coding system employing coding blocks are reduced by detecting a transient in the audio signal stream and shifting the temporal relationship of the transient with respect to the coding blocks such that the time duration of the distortion artifacts is reduced. The audio data is time scaled in such a way that the transients are temporally repositioned prior to quantization in a transform-based low-bit-rate audio encoder so as to reduce the amount of pre-noise in the decoded audio signal. Alternatively, or in addition, in a transform-based low-bit-rate audio coding system, a transient in the audio signal stream is detected and a portion of the distortion artifacts are time compressed such that the time duration of the distortion artifacts is reduced.

WO 02/093560 A1

WO 02/093560 A1



Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM),
European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR,
GB, GR, IE, IT, LU, MC, NL, PT, SE, TR), OAPI patent
(BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR,
NE, SN, TD, TG).

— *with amended claims*

For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

Published:

— *with international search report*

- 1 -

DESCRIPTION

Improving Transient Performance of Low Bit Rate Audio Coding
Systems by Reducing Pre-Noise

5

TECHNICAL FIELD

The invention relates generally to high-quality, low bit rate digital transform
10 encoding and decoding of information representing audio signals such as music or
voice signals. More particularly, the invention relates to the reduction of distortion
artifacts preceding a signal transient ("pre-noise") in an audio signal stream produced
by such an encoding and decoding system.

15

BACKGROUND ART**Time Scaling**

Time scaling refers to altering the time evolution or duration of an audio signal
while not altering its spectral content (perceived timbre) or perceived pitch (where
pitch is a characteristic associated with periodic audio signals). Pitch scaling refers
20 to modifying the spectral content or perceived pitch of an audio signal while not
affecting its time evolution or duration. Time scaling and pitch scaling are dual
methods of one another. For example, a digitized audio signal's pitch may be
increased 5% without affecting its time duration by time scaling it by 5% (*i.e.*,
increasing the time duration of the signal) and then reading out the samples at a 5%
25 higher sample rate (*e.g.*, by resampling), thereby maintaining its original time
duration. The resulting signal has the same time duration as the original signal but
with modified pitch or spectral characteristics. Resampling is not an essential step of
time scaling or pitch scaling unless it is desired to maintain a constant output
sampling rate or to maintain the input and output sampling rates the same.

30

In aspects of the present invention, time scaling processing of audio streams is
employed. However, as mentioned above, time scaling may also be performed using
pitch-scaling techniques, as they are duals of one another. Thus, while the term

- 2 -

"time scaling" is used herein, techniques that employ pitch scaling to achieve time scaling may also be employed.

Low Bit Rate Audio Coding

There is considerable interest among those in the field of signal processing to
5 minimize the amount of information required to represent a signal without perceptible loss in signal quality. By reducing information requirements, signals impose lower information capacity requirements upon communication channels and storage media. With respect to digital coding techniques, minimal informational requirements are synonymous with minimal binary bit requirements.

10 Some prior art techniques for coding audio signals intended for human hearing attempt to reduce information requirements without producing any audible degradation by exploiting psychoacoustic effects. The human ear displays frequency-analysis properties resembling those of highly asymmetrical tuned filters having variable center frequencies. The ability of the human ear to detect distinct
15 tones generally increases as the difference in frequency between the tones increases; however, the ear's resolving ability remains substantially constant for frequency differences less than the bandwidth of the above mentioned filters. Thus, the frequency-resolving ability of the human ear varies according to the bandwidth of these filters throughout the audio spectrum. The effective bandwidth of such an
20 auditory filter is referred to as a critical band. A dominant signal within a critical band is more likely to mask the audibility of other signals anywhere within that critical band than other signals at frequencies outside that critical band. A dominant signal may mask other signals occurring not only at the same time as the masking signal, but also occurring before and after the masking signal. The duration of pre-
25 and post-masking effects within a critical band depend upon the magnitude of the masking signal, but pre-masking effects are usually of much shorter duration than post-masking effects. See generally, the *Audio Engineering Handbook*, K. Blair Benson ed., McGraw-Hill, San Francisco, 1988, pages 1.40-1.42 and 4.8-4.10.

- 3 -

Signal recording and transmitting techniques that divide the useful signal bandwidth into frequency bands with bandwidths approximating the ear's critical bands can better exploit psychoacoustic effects than wider band techniques. Techniques that exploit psychoacoustic masking effects can encode and reproduce a
5 signal that is indistinguishable from the original input signal using a bit rate below that required by PCM coding.

Critical band techniques comprise dividing the signal bandwidth into frequency bands, processing the signal in each frequency band, and reconstructing a replica of the original signal from the processed signal in each frequency band. Two
10 such techniques are sub-band coding and transform coding. Sub-band and transform coders can reduce transmitted informational requirements in particular frequency bands where the resulting coding inaccuracy (noise) is psychoacoustically masked by neighboring spectral components without degrading the subjective quality of the encoded signal.

15 A bank of digital bandpass filters may implement sub-band coding. Transform coding may be implemented by any of several time-domain to frequency-domain discrete transforms that implement a bank of digital bandpass filters. The remaining discussion relates more particularly to transform coders, therefore the term "sub-band" is used here to refer to selected portions of the total signal bandwidth, whether
20 implemented by a sub-band coder or a transform coder. A sub-band as implemented by a transform coder is defined by a set of one or more adjacent transform coefficients; hence, the sub-band bandwidth is a multiple of the transform coefficient bandwidth. The bandwidth of a transform coefficient is proportional to the input signal sampling rate and inversely proportional to the number of coefficients
25 generated by the transform to represent the input signal.

Psychoacoustic masking may be more easily accomplished by transform coders if the sub-band bandwidth throughout the audible spectrum is about half the critical bandwidth of the human ear in the same portions of the spectrum. This is because the critical bands of the human ear have variable center frequencies that

- 4 -

adapt to auditory stimuli, whereas sub-band and transform coders typically have fixed sub-band center frequencies. To optimize the utilization of psychoacoustic-masking effects, any distortion artifacts resulting from the presence of a dominant signal should be limited to the sub-band containing the dominant signal. If the sub-band bandwidth is about half or less than half of the critical band and if filter
5 selectivity is sufficiently high, effective masking of the undesired distortion products is likely to occur even for signals whose frequency is near the edge of the sub-band passband bandwidth. If the sub-band bandwidth is more than half a critical band, there is a possibility that the dominant signal may cause the ear's critical band to be
10 offset from the coder's sub-band such that some of the undesired distortion products outside the ear's critical bandwidth are not masked. This effect is most objectionable at low frequencies where the ear's critical band is narrower.

The probability that a dominant signal may cause the ear's critical band to offset from a coder sub-band and thereby "uncover" other signals in the same coder
15 sub-band is generally greater at low frequencies where the ear's critical band is narrower. In transform coders, the narrowest possible sub-band is one transform coefficient, therefore psychoacoustic masking may be more easily accomplished if the transform coefficient bandwidth does not exceed one half the bandwidth of the ear's narrowest critical band. Increasing the length of the transform may decrease
20 the transform coefficient bandwidth. One disadvantage of increasing the length of the transform is an increase in the processing complexity to compute the transform and to encode larger numbers of narrower sub-bands. Other disadvantages are discussed below.

Of course, psychoacoustic masking may be achieved using wider sub-bands if
25 the center frequency of these sub-bands can be shifted to follow dominant signal components in much the same way the ear's critical band center frequency shifts.

The ability of a transform coder to exploit psychoacoustic masking effects also depends upon the selectivity of the filter bank implemented by the transform. Filter "selectivity," as that term is used here, refers to two characteristics of sub-band

- 5 -

bandpass filters. The first is the bandwidth of the regions between the filter pass-band and stopbands (the width of the transition bands). The second is the attenuation level in the stopbands. Thus, filter selectivity refers to the steepness of the filter response curve within the transition bands (steepness of transition band rolloff), and
5 the level of attenuation in the stopbands (depth of stopband rejection).

Filter selectivity is directly affected by numerous factors including the three factors discussed below: block length, window weighting functions, and transforms. In a very general sense, block length affects coder temporal and frequency resolution, and windows and transforms affect coding gain.

10 **Low Bit Rate Audio Coding / Block Length**

The input signal to be encoded is sampled and segmented into “signal sample blocks” prior to sub-band filtering. The number of samples in the signal sample block is the signal sample block length.

It is common for the number of coefficients generated by a transform filter
15 bank (the transform length) to be equal to the signal sample block length, but this is not necessary. An overlapping-block transform may be used and is sometimes described in the art as a transform of length N that transforms signal sample blocks with $2N$ samples. This transform can also be described as a transform of length $2N$ that generates only N unique coefficients. Because all the transforms discussed here
20 can be thought to have lengths equal to the signal sample block length, the two lengths are generally used here as synonyms for one another.

The signal sample block length affects the temporal and frequency resolution of a transform coder. Transform coders using shorter block lengths have poorer frequency resolution because the discrete transform coefficient bandwidth is wider
25 and filter selectivity is lower (decreased rate of transition band rolloff and a reduced level of stopband rejection). This degradation in filter performance causes the energy of a single spectral component to spread into neighboring transform coefficients. This undesirable spreading of spectral energy is the result of degraded filter performance called “sidelobe leakage.”

- 6 -

Transform coders using longer block lengths have poorer temporal resolution because quantization errors cause a transform encoder/decoder system to “smear” the frequency components of a sampled signal across the full length of the signal sample block. Distortion artifacts in the signal recovered from the inverse transform are most audible as a result of large changes in signal amplitude that occur during a time interval much shorter than the signal sample block length. Such amplitude changes are referred to here as “transients.” Such distortion manifests itself as noise in the form of an echo or ringing just before (pre-transient noise or “pre-noise”) and just after (post-transient noise) the transient. Pre-noise is of particular concern because it is highly audible and, unlike post-transient noise, only minimally masked (a transient provides only minimal temporal pre-masking). Pre-noise is produced when the high frequency components of transient audio material are temporally smeared through the length of the audio coder block in which it occurs. The present invention is concerned with minimizing pre-noise. Post-transient noise typically is substantially masked and is not the subject of the present invention.

Fixed block length transform coders use a compromise block length that trades off temporal resolution against frequency resolution. A short block length degrades sub-band filter selectivity, which may result in a nominal passband filter bandwidth that exceeds the ear's critical bandwidth at lower or at all, frequencies. Even if the nominal sub-band bandwidth is narrower than the ear's critical bandwidth, degraded filter characteristics manifested as a broad transition band and/or poor stopband rejection may result in significant signal artifacts outside the ear's critical bandwidth. On the other hand, a long block length may improve filter selectivity but reduces temporal resolution, which may result in audible signal distortion occurring outside the ear's temporal psychoacoustic masking interval.

Window Weighting Function

Discrete transforms do not produce a perfectly accurate set of frequency coefficients because they work with only a finite-length segment of the signal, the signal sample block. Strictly speaking, discrete transforms produce a time-frequency

- 7 -

representation of the input time-domain signal rather than a true frequency-domain representation which would require infinite signal sample block lengths. For convenience of discussion here, however, the output of discrete transforms is referred to as a frequency-domain representation. In effect, the discrete transform assumes
5 that the sampled signal only has frequency components whose periods are a submultiple of the signal sample block length. This is equivalent to an assumption that the finite-length signal is periodic. The assumption in general, of course, is not true. The assumed periodicity creates discontinuities at the edges of the signal sample block that cause the transform to create phantom spectral components.

10 One technique that minimizes this effect is to reduce the discontinuity prior to the transformation by weighting the signal samples such that samples near the edges of the signal sample block are zero or close to zero. Samples at the center of the signal sample block are generally passed unchanged, *i.e.*, weighted by a factor of one. This weighting function is called an “analysis window.” The shape of the
15 window directly affects filter selectivity.

As used here, the term “analysis window” refers only to the windowing function performed prior to application of the forward transform. The analysis window is a time-domain function. If no compensation for the window’s effects is provided, the recovered or “synthesized” signal is distorted according to the shape of
20 the analysis window. One compensation method known as overlap-add is well known in the art. This method requires the coder to transform overlapped blocks of input signal samples. By carefully designing the analysis window such that two adjacent windows add to unity across the overlap, the effects of the window are exactly compensated.

25 Window shape affects filter selectivity significantly. See generally, Harris, “On the Use of Windows for Harmonic Analysis with the Discrete Fourier Transform,” *Proc IEEE*, vol. 66, January, 1978, pp. 51—83. As a general rule, “smoother” shaped windows and larger overlap intervals provide better selectivity.

- 8 -

For example, a Kaiser-Bessel window generally provides for greater filter selectivity than a sine-tapered rectangular window.

When used with certain types of transforms such as the Discrete Fourier Transform (DFT), overlap-add increases the number of bits required to represent the
5 signal because the portion of the signal in the overlap interval must be transformed and transmitted twice, once for each of the two overlapped signal sample blocks. Signal analysis/synthesis for systems using such a transform with overlap-add is not critically sampled. The term "critically sampled" refers to a signal analysis/synthesis which over a period of time generates the same number of frequency coefficients as
10 the number of input signal samples it receives. Hence, for noncritically sampled systems, it is desirable to design the window with an overlap interval as small as possible in order to minimize the coded signal information requirements.

Some transforms also require that the synthesized output from the inverse transform be windowed. The synthesis window is used to shape each synthesized
15 signal block. Therefore, the synthesized signal is weighted by both an analysis and a synthesis window. This two-step weighting is mathematically similar to weighting the original signal once by a window whose shape is equal to a sample-by-sample product of the analysis and synthesis windows. Therefore, in order to utilize overlap-add to compensate for windowing distortion, both windows must be designed such
20 that the product of the two sums to unity across the overlap-add interval.

While there is no single criterion that may be used to assess a window's optimality, a window is generally considered "good" if the selectivity of the filter used with the window is considered "good." Therefore, a well designed analysis window (for transforms that use only an analysis window) or analysis/synthesis
25 window pair (for transforms that use both an analysis and a synthesis window) can reduce sidelobe leakage.

Block Switching

A common solution that addresses the compromise between temporal and frequency resolution in fixed block length transform coders is the use of transient

- 9 -

detection and block length switching. In this solution the presence and location of audio signal transients are detected using various transient detection methods. When transient audio signals are detected that are likely to introduce pre-noise when coded using a long audio coder block length, the low bit rate coder switches from the more efficient long block length to a less efficient shorter block length. While this reduces the frequency resolution and coding efficiency of the encoded audio signal it also reduces the length of transient pre-noise introduced by the coding process, improving the perceived quality of the audio upon low bit rate decoding. Techniques for block length switching are disclosed in U.S. Patents 5,394,473; 5,848,391; and 6,226,608 B1, each of which is hereby incorporated by reference in its entirety. Although the present invention reduces pre-noise without the complexity and disadvantages of block switching, it may be employed along with and in addition to block switching.

DISCLOSURE OF THE INVENTION

In accordance with a first aspect of the present invention, a method for reducing distortion artifacts preceding a signal transient in an audio signal stream processed by a transform-based low-bit-rate audio coding system employing coding blocks comprises detecting a transient in the audio signal stream, and shifting the temporal relationship of the transient with respect to the coding blocks such that the time duration of the distortion artifacts is reduced.

An audio signal is analyzed and the locations of transient signals are identified. The audio data is then time scaled in such a way that the transients are temporally repositioned prior to quantization in a transform-based low-bit-rate audio encoder so as to reduce the amount of pre-noise in the decoded audio signal. Such processing prior to encoding and decoding is referred to herein as “pre-processing.”

Thus, before quantization in the encoder, because the quantization process smears the transient throughout the encoding block creating the undesired pre-noise artifacts, the transient is shifted to a better position *vis-à-vis* block ends using time scaling (time compression or time expansion). Such pre-processing may also be

- 10 -

referred to as “transient time shifting”. Transient time shifting requires the identification of transients and also requires information as to their temporal location relative to block ends. In principle, transient time shifting may be accomplished in the time domain prior to application of the forward transform or in the frequency
5 domain following application of the forward transform but prior to quantization. In practice, transient time shifting may be more easily accomplished in the time domain prior to application of the forward transform, particularly when a compensating time scaling is performed as described below.

The results of transient time shifting may be audible because both the transient
10 and the audio stream are no longer in their original relative temporal positions — the time evolution of the audio stream is altered as a result of time compression or time expansion of the audio stream before the transient. A listener may perceive this as an alteration in the rhythm within a musical piece, for example.

There are several compensation techniques for reducing such an alteration in
15 the audio stream’s time evolution that form aspects of the present invention. These compensation techniques are optional because slight variations in the temporal evolution of an audio signal are not discernable to most listeners. Compensation techniques are discussed after the following discussion of a second aspect of the present invention.

20 In accordance with a second aspect of the present invention, in an encoder of a transform-based low-bit-rate audio coding system employing coding blocks, a method for reducing distortion artifacts preceding a signal transient in an audio signal stream subsequent to inverse transformation, comprises detecting a transient in the audio signal stream, and time compressing at least a portion of the distortion artifacts
25 such that the time duration of the distortion artifacts is reduced.

By such processing, referred to as “post-processing” herein, audio quality improvements to any audio signal that has undergone low bit rate audio encoding may be obtained whether or not pre-processing is employed and, if it is employed, whether or not the encoder transmits metadata useful for the post-processing. Any

- 11 -

audio signal that has undergone low bit rate audio encoding and decoding may be analyzed to identify the location of transient signals and to estimate the duration of transient pre-noise artifacts. Then, time scale post-processing may be performed on the audio so as to remove the transient signal pre-noise or reduce its duration.

5 As mentioned above, there are several compensation techniques for reducing alterations in the audio stream's time evolution. These time scaling compensation techniques also have the beneficial result of keeping the number of audio samples constant.

A first time scaling compensation technique, useful in connection with pre-
10 processing, is applied before the forward transform. It applies a compensating time scaling to the audio stream following the transient, the time scaling having a sense opposite to the sense of the time scaling employed to shift the transient position and, preferably, having substantially the same duration as the transient-shifting time scaling. For convenience in discussion, this type of compensation is referred to
15 herein as "sample number compensation" because it is capable of keeping the number of audio samples constant but is not capable of fully restoring the original temporal evolution of the audio signal stream (it leaves the transient and portions of the signal stream near the transient out of place temporally). Preferably, the time-scaling providing sample number compensation closely follows the transient such
20 that it is temporally post-masked by the transient.

Although sample number compensation leaves the transient shifted from its original temporal position, it does restore the audio stream following the compensating time scaling to its original relative temporal position. Thus, the likelihood of audibility of the transient time shifting is reduced, although it is not
25 eliminated, because the transient is still out of its original position. Nevertheless, this may provide a sufficient reduction in audibility and it has the advantage that it is done prior to low bit-rate audio encoding, allowing the use of a standard, unmodified decoder. As explained below, a full restoration of the audio signal stream's time evolution can only be accomplished by processing in the decoder or following the

- 12 -

decoder. In addition to reducing the possibility of audibility of the transient time shifting, time-scaling compensation before forward transformation has the advantage of keeping the number of audio samples constant, which may be important for processing and/or for the operation of hardware implementing the processing.

5 In order to provide optimum time-scaling compensation before forward transformation, information as to the location of the transient and the temporal length of the transient time shifting should be employed by the compensation process.

If transient time shifting is applied after blocking (but before applying the forward transform), it is necessary to employ sample number compensation within
10 the same block in which transient time shifting is done in order to keep the block length the same. Consequently, it is preferred to perform transient time shifting and sample number compensation before blocking.

Sample number compensation may also be employed after the inverse transform (either in the decoder or after decoding) in connection with post-
15 processing. In this case, information useful for performing compensation may be sent to the compensation process from the decoder (which information may have originated in the encoder and/or the decoder).

A more complete restoration of the audio signal stream's temporal evolution along with restoring the original number of audio samples may be accomplished after
20 the inverse transform (either in the decoder or following decoding), by apply a compensating time scaling to the audio stream before the transient in the sense opposite to the sense of the time scaling employed to shift the transient position and, preferably, of substantially the same duration as the transient-shifting time scaling. For convenience in discussion, this type of compensation is referred to herein as
25 "time evolution compensation." This time scaling compensation has the significant advantage of restoring the entire audio stream, including the transient, to its original relative temporal position. Thus, the likelihood of audibility of the time scaling processes is greatly reduced, although not eliminated, because the two time scaling processes themselves may cause audible artifacts.

- 13 -

In order to provide optimum time-evolution compensation, various information such as the location of the transient, the location of the block ends, the length of the transient time shifting, and the length of the pre-noise is useful. The length of the pre-noise is useful in assuring that the time-scaling of the time evolution compensation does not occur during the pre-noise, thus possibly expanding the temporal length of the pre-noise. The length of the transient time shifting is useful if it is desired to restore the audio stream to its original relative temporal position and to maintain the number of samples constant. The location of the transient is useful because the length of the pre-noise may be determined from the original location of the transient with respect to the ends of the coding blocks. The length of the pre-noise may be estimated by measuring a signal parameter, such as high-frequency content or a default value may be employed. If the compensation is performed in the decoder or after decoding, useful information may be sent by the encoder as metadata along with the encoded audio. When performed after decoding, metadata may be sent to the compensation process from the decoder (which information may have originated in the encoder and/or the decoder).

As mentioned above, post-processing to reduce the length of the pre-noise artifact may also be applied as an additional step to an audio coder that performs time scaling pre-processing and, optionally, provides metadata information. Such post-processing would act as an additional quality improvement scheme by reducing the pre-noise that may still remain after pre-processing.

Pre-processing may be preferred in coder systems employing professional encoders in which cost, complexity and time-delay are relatively immaterial in comparison to post-processing in connection with a decoder, which is typically a lower complexity consumer device.

The low bit rate audio coding system quality improvement technique of the present invention may be implemented using any suitable time-scaling technique, as well as any that may become available in the future. One suitable technique is described in International Patent Application PCT/US02/04317, filed February 12,

- 14 -

2002, entitled High Quality Time-Scaling and Pitch-Scaling of Audio Signals. Said application designates the United States and other entities. The application is hereby incorporated by reference in its entirety. As discussed above, since time scaling and pitch shifting are dual methods of one another, time scaling may also be implemented using any suitable pitch scaling technique, as well as any that may become available in the future. Pitch scaling following by reading out the audio samples at an appropriate rate that is different than the input sample rate results in a time scaled version of the audio with the same spectral content or pitch of the original audio and is applicable to the present invention.

10 As discussed in the low bit rate audio coding background summary, the selection of block length in an audio coding system is a trade-off between frequency and temporal resolution. In general, a longer block length is preferred as it provides increased efficiency of the coder (generally provides greater perceived audio quality with a reduced number of data bits) in comparison to a shorter block length.

15 However, transient signals and the pre-noise signals that they generate offset the quality gain of longer block lengths by introducing audible impairments. It is for this reason that block switching or fixed smaller block lengths are used in practical applications of low bit rate audio coders. However, applying time scaling pre-processing in accordance with the present invention to audio data that is to undergo

20 low bit rate audio coding and/or has undergone post-processing may reduce the duration of transient pre-noise. This allows longer audio coding block lengths to be used, thereby providing increased coding efficiency and improving perceived audio quality without adaptively switching block lengths. However, the reduction of pre-noise in accordance with the present invention may be also employed in coding

25 systems that employ block length switching. In such systems, some pre-noise may exist even for the smallest window size. The larger the window, the longer and, consequently, more audible the pre-noise is. Typical transients provide approximately 5 msec of premasking, which translates to 240 samples at a 48 kHz

- 15 -

sampling rate. If a window is larger than 256 samples, which is common in a block switching arrangement, the invention provides some benefit.

Audio Coding Transient Pre-Noise Artifacts

FIGS. 1a-1e show examples of transient pre-noise artifacts generated by a
5 fixed block length audio coder system. FIG. 1a shows six, 50% overlapped, audio coding windowed blocks of fixed length 1 through 6. In this figure and all other figures herein, each window is contiguous with an audio coding block and is referred to as a "windowed block," "window," or "block." In this figure and certain other figures herein, the windows are shown generally in the shape of a Kaiser-Bessel
10 window. Other figures show windows in the shape of semi-circles for simplicity in presentation. Window shape is not critical to the present invention. While the length of the windowed blocks in FIG. 1a and other figures is not critical to the invention, fixed length windowed blocks typically are in the range of 256 to 2048 samples in length. The four audio signal examples in FIGS. 1b through 1e illustrate,
15 respectively, the effects of temporal relationships between the audio coding windowed blocks and the transient pre-noise artifacts.

FIG. 1b illustrates the relationship between the location of a transient signal in an input audio stream to be coded and the borders of the 50% overlapping windowed blocks. While a 50% overlapping fixed block length is shown, the invention is
20 applicable to both fixed and variable block length coding systems and to blocks having other than a 50% overlap, including no overlap as is discussed below in connection with FIGS. 2a through 5b.

FIG. 1c shows the audio signal stream output of the audio coding system for the case of an audio signal stream input as shown in FIG. 1b. As shown in FIGS. 1b and 1c, the transient is located between the end of windowed block 3 and the end of windowed block 4. FIG. 1c illustrates the location and length of transient pre-noise introduced by the low bit rate audio coding process in relation to the location of the transient and the end of windowed block 2. Note that the pre-noise is prior to the transient and is limited to windowed blocks 4 and 5, the sample blocks in which the

- 16 -

transient lies. Thus, the pre-noise extends back to the beginning of windowed block 4.

Similarly to FIGS. 1b and 1c, FIGS. 1d and 1e show, respectively, the relationship between an input audio signal stream that contains a transient located
5 between the end of windowed block 2 and the end of windowed block 3 and the pre-noise introduced in the output audio signal stream by the audio coding system. Because the pre-noise is limited to windowed blocks 3 and 4, within which the transient lies, the pre-noise extends back to the beginning of windowed block 3. In this case, the pre-noise has a longer duration because the transient is nearer the end of
10 windowed block 3 than the transient of FIGS. 1b and 1c is to the end of windowed block 4. The ideal transient location is closely following the last block end so that the pre-noise extends back only to the next prior block end (about half of the block length in the case of this 50% block overlap example).

It should be noted that the examples in FIGS. 1a-1e do not explicitly take into
15 account the effects of cross fading at the coding window boundaries. In general, as the audio coding windows taper off, the pre-noise artifacts are scaled accordingly and their audibility is reduced. For simplicity in presentation, scaling of the pre-noise artifacts is not shown in the idealized waveforms of the figures herein.

As suggested in FIGS. 1a-1e and shown in more detail in FIGS. 2A, 2B, 3A,
20 3B, 4A, 4B, 5A and 5B, an audio coder's transient pre-noise artifacts may be minimized if the location of transient signals is judiciously positioned prior to audio encoding.

Examples of repositioning the location of a transient in order to reduce pre-noise are shown in FIGS. 2a, 2b, 3a, 3b, 4a, 4b, 5a and 5b for the cases of non-
25 overlapping blocks (FIGS. 2a and 2b), less than 50% block overlap (FIGS. 3a and 3b), 50% block overlap (FIGS. 4a and 4b), and greater than 50% block overlap (FIGS. 5a and 5b). In each case, unless the original position of the transient is equidistant between two successive block ends (in which case there is no preference), it is preferred to shift the transient to a position closely following the nearest block

- 17 -

end. Whether the shift is to the prior block end or to the next block end, whether or not the nearest block end, the resulting pre-noise is substantially the same. However, by temporally shifting the transient to a location closely following the nearest block end, disruption to the time evolution of the audio stream is minimized, thereby
5 minimizing the possible audibility of shifting the transient. Nevertheless, in some cases, shifting to the more distant block end may also be inaudible. Moreover, even if a shifting to the more distant block end is audible, time evolution compensation, as described below, may be employed to reduce or eliminate such audibility.

FIGS. 2a and 2b show a series of idealized non-overlapping windowed blocks.
10 In FIG. 2a, a transient's initial location is, as shown by the solid-lined arrow, closer to the last window end than it is to the next window end. The pre-noise for the transient's initial location extends back in time to the end of the beginning of the window, as shown. If it is desired to minimize the degree of temporal shift of the transient, it should be shifted "left" (backward in time) to a location closely following
15 the end of the last windowed block, as shown. Although the resulting pre-noise still extends back to the beginning of the windowed block, this length is very short compared to the pre-noise resulting from the initial transient location. In this and other figures, the distance of the shifted transient from the windowed block end is exaggerated for clarity in presentation. In FIG. 2b, the initial position of the transient
20 is closer to the next window end than to the previous window end. Thus, if it is desired to minimize the degree of temporal shift of the transient, it should be shifted "right" (later in time) to a location closely following the end of the next windowed block, as shown. It will be noted that the improvement in pre-noise reduction increases as the initial transient position becomes later in the windowed block.

25 FIGS. 3a and 3b show a series of idealized windowed blocks that overlap by less than 50%. In FIG. 3a, a transient's initial location is, as shown by the solid-lined arrow, closer to the last window end than it is to the next window end. The pre-noise for the transient's initial location extends back in time to the end of the beginning of the window, as shown. If it is desired to minimize the degree of temporal shift of the

- 18 -

transient, it should be shifted "left" to a location closely following the end of the last windowed block, as shown. The resulting pre-noise still extends back to the beginning of the windowed block, but this length is short compared to the pre-noise resulting from the initial transient location. In FIG. 3b, the initial position of the transient is closer to the next window end than to the previous window end. Thus, if it is desired to minimize the degree of temporal shift of the transient, it should be shifted "right" to a location closely following the end of the next windowed block, as shown. It will be noted that the improvement in pre-noise reduction increases as the initial transient position is later in the interval between successive windowed blocks.

FIGS. 4a and 4b show a series of idealized windowed blocks that overlap by 50%. In FIG. 4a, a transient's initial location is, as shown by the solid-lined arrow, closer to the last window end than it is to the next window end. The pre-noise for the transient's initial location extends back in time to the end of the beginning of the window, as shown. If it is desired to minimize the degree of temporal shift of the transient, it should be shifted "left" to a location closely following the end of the last windowed block, as shown. The resulting pre-noise still extends back to the beginning of the windowed block, but this length is shorter than the pre-noise resulting from the initial transient location. In FIG. 4b, the initial position of the transient is closer to the next window end than to the previous window end. Thus, if it is desired to minimize the degree of temporal shift of the transient, it should be shifted "right" to a location closely following the end of the next windowed block, as shown. It will be noted that the improvement in pre-noise reduction increases as the initial transient position is later in the interval between successive windowed block ends, as in the case of less than 50% overlapped blocks.

FIGS. 5a and 5b show a series of idealized windowed blocks that overlap by greater than 50%. In FIG. 5a, a transient's initial location is, as shown by the solid-lined arrow, closer to the last window end than it is to the next window end. The pre-noise for the transient's initial location extends back in time to the end of the beginning of the window, as shown. If it is desired to minimize the degree of

- 19 -

temporal shift of the transient, it should be shifted "left" to a location closely following the end of the last windowed block, as shown. The resulting pre-noise still extends back to the beginning of the windowed block, but this length is still somewhat shorter than the pre-noise resulting from the initial transient location. In FIG. 5b, the initial position of the transient is closer to the next window end than to the previous window end. Thus, if it is desired to minimize the degree of temporal shift of the transient, it should be shifted "right" to a location closely following the end of the next windowed block, as shown. It will be noted that the improvement in pre-noise reduction increases as the initial transient position is later in the interval between successive windowed block ends, as in the case of 50% overlapped blocks.

It will be noted that the improvement in pre-noise reduction is greatest for non-overlapping blocks and decreases as the degree of block overlap increases.

DESCRIPTION OF THE DRAWINGS

FIGS. 1a-1e are a series of idealized waveforms illustrating examples of transient pre-noise artifacts generated by a fixed block length audio coder system for two cases of input signal conditions.

FIGS. 2a and 2b show a series of idealized non-overlapping windowed blocks illustrating initial and shifted transient temporal locations, along with the pre-noise for such locations, for the case of an initial position closer to the last window end than to the next window end and for the case of an initial position closer to the next window end than to the previous window end, respectively.

FIGS. 3a and 3b show a series of idealized less than 50% overlapping windowed blocks illustrating initial and shifted transient temporal locations, along with the pre-noise for such locations, for the case of an initial position closer to the last window end than to the next window end and for the case of an initial position closer to the next window end than to the previous window end, respectively.

FIGS. 4a and 4b show a series of idealized 50% overlapping windowed blocks illustrating initial and shifted transient temporal locations, along with the pre-noise

- 20 -

for such locations, for the case of an initial position closer to the last window end than to the next window end and for the case of an initial position closer to the next window end than to the previous window end, respectively.

FIGS. 5a and 5b show a series of idealized greater than 50% overlapping windowed blocks illustrating initial and shifted transient temporal locations, along with the pre-noise for such locations, for the case of an initial position closer to the last window end than to the next window end and for the case of an initial position closer to the next window end than to the previous window end, respectively.

FIG. 6 is a flow chart showing steps to reduce transient pre-noise artifacts by time scaling prior to low bit rate encoding.

FIG. 7 is a conceptual representation of an input data buffer used for transient detection.

FIGS. 8a-8e are a series of idealized waveforms illustrating an example of audio time scaling pre-processing in accordance with aspects of the present invention when a transient exists in an audio coding block and is located closer to the last windowed block end than to the next windowed block end.

FIGS. 9a-9e are a series of idealized waveforms illustrating an example of audio time scaling processing when a transient exists in a windowed audio coding block and is located approximately T samples before a block end.

FIGS. 10a-10d are a series of idealized waveforms illustrating time scaling for the case of multiple transients.

FIGS. 11a-11f are a series of idealized waveforms illustrating intelligent time evolution compensation of time scaling using metadata conveyed in the audio stream.

FIG. 12 is a flow chart of time scaling post-processing in conjunction with a low bit rate audio decoder.

FIGS. 13a-13c are a series of idealized waveforms illustrating an example of post-processing for a single transient to reduce the pre-noise artifacts present after decoding.

- 21 -

FIG. 14 is a flow chart of a post-processing process for improving the perceived quality of audio that has undergone low bit rate coding without time scaling pre-processing.

FIGS. 15a-15c are a series of idealized waveforms demonstrating the technique of using a default value to time-scale the audio before each transient to reduce pre-noise without performing sample number compensation.

FIGS. 16a-16c are a series of idealized waveforms demonstrating the technique of using a computed pre-noise duration to time-scale the audio before each transient to reduce pre-noise duration with sample number and time evolution compensation.

BEST MODE FOR CARRYING OUT THE INVENTION

Time Scaling Pre-Processing Overview

FIG. 6 is a flow chart illustrating a method for time-scaling audio prior to low bit rate audio encoding to reduce the amount of transient pre-noise (*i.e.*, “pre-processing”). This method processes the input audio in N sample blocks, where N may correspond to a number greater than or equal to the number of audio samples used in the audio coding block. Processing sizes with N greater than the size of the audio coding block may be desirable to provide additional audio data outside of the audio coding block for use in time scaling processing. This additional data may be used, for example, to sample number compensate for time scaling processing performed to improve the location of a transient.

The first step 202 in the process of FIG. 6 checks for the availability of N audio data samples for time scaling processing. These audio data samples may be from, for example, a file on a PC-based hard disk or a data buffer in a hardware device. The audio data may also be provided by a low bit rate audio coding process that invokes the time scaling processor prior to audio encoding. If N audio data samples are available they are passed (step 204) to and then used by the time scaling pre-processing process in the following steps.

- 22 -

The third step 206 in the pre-processing process is detecting the location of audio data transient signals that are likely to introduce pre-noise artifacts. Many different processes are available to perform this function and the specific implementation is not critical as long as it provides accurate detection of transient signals that are likely to introduce pre-noise artifacts. Many audio coding processes perform audio signal transient detection and this step may be skipped if the audio coding process provides the transient information to the subsequent time scaling processing block 210 along with the input audio data.

Transient Detection

One suitable method for performing audio signal transient detection is as follows. The first step in the transient detection analysis is to filter the input data (treating the data samples as a time function). The input data may, for example, be filtered with a 2nd order IIR high-pass filter with a 3 dB cutoff frequency of approximately 8 kHz. The filter characteristics are not critical. This filtered data is then used in the transient analysis. Filtering the input data isolates the high frequency transients and makes them easier to identify. Next, the filtered input data are processed in sixty-four sub-blocks (in the case of a 4096 sample signal sample block) of approximately 1.5 msec (or 64 samples at 44.1 kHz) as shown in FIG. 7. While the actual size of the processing sub-block is not constrained to 1.5 msec and may vary, this size provides a good trade-off between real-time processing requirements (as larger block sizes require less processing overhead) and resolution of transient location (smaller blocks provide more detailed information on the location of transients). The use of 4096 sample signal sample blocks and the use of 64 sample sub-blocks is merely an example and is not critical to the invention.

The next step of transient detection processing is to perform a low-pass filtering of the maximum absolute data values contained in each 64-sample sub-block. This processing is performed to smooth the maximum absolute data and provide a general indication of the average peak values in the input buffer to which

- 23 -

the actual sub-buffer peak value can be compared. The method described below is one method of doing the smoothing.

To smooth the data, each 64-sample sub-block is scanned for the maximum absolute data signal value. The maximum absolute data signal value is then used to
 5 compute a smoothed, moving average peak value. The filtered, high frequency moving averages for each k^{th} sub-buffer, $hi_mavg(k)$ respectively, are computed using Equations 1 and 2.

for buffer $k = 1:1:64$

$hi_mavg(k) = hi_mavg(k-1) + ((hi\ freq\ peak\ val\ in\ buffer\ k) - hi_mavg(k-1)) * AVG_WHT$ (1)
 10 end

where $hi_mavg(0)$ is set equal to $hi_mavg(64)$ from the previous input buffer for continuous processing. In the current implementation the parameter AVG_WHT is set equal to 0.25. This value was decided upon following experimental analysis
 15 using a wide range of common audio material.

Next, the transient detection processing compares the peak in each sub-block to the array of smoothed, moving average peak values to determine whether a transient exists. While a number of methods exist to compare these two measures the approach outlined below was taken because it allows tuning of the comparison by use
 20 of a scaling factor that has been set to perform optimally as determined by analyzing a wide range of audio signals.

The peak value in the k^{th} sub-block, for the filtered data, is multiplied by the high frequency scaling value HI_FREQ_SCALE , and compared to the computed smoothed, moving average peak value of each k . If a sub-block's scaled peak value
 25 is greater than the moving average value a transient is flagged as being present. These comparisons are outlined below in Equations 3 and 4.

for buffer $k = 1:1:64$

if $((hi\ freq\ peak\ value\ in\ buffer\ k) * HI_FREQ_SCALE) > hi_mavg(k)$
 (2)

- 24 -

flag high frequency transient in sub-block k = TRUE

end

end

Following transient detection, several corrective checks are made to determine
5 whether the transient flag for a 64-sample sub-block should be cancelled (reset from
TRUE to FALSE). These checks are performed to reduce false transient detections.
First, if the high frequency peak values fall below a minimum peak value then the
transient is cancelled (to address low level transients). Second, if the peak in a sub-
block triggers a transient but is not significantly larger than the previous sub-block,
10 which also would have triggered a transient flag, then the transient in the current sub-
block is cancelled. This reduces a smearing of the information on the location of a
transient.

Referring again to FIG. 6, the next step 208 in processing is to
determine whether transients exist in the current N sample input data array. If no
15 transients exist the input data may be output (or passed back to a low-bit rate audio
coder) with no time scaling processing performed. If transients do exist, the number
of transients that exist in the current N samples of audio data and their location(s) are
passed to the audio time scaling processing portion 210 of the process for temporal
modification of the input audio data. The result of suitable time-scale processing is
20 described in connection with the description of FIGS. 8a-8e. Note that the process
requires information from the encoder as to, for example, the location of the
windowed sample blocks with respect to the audio data stream. If, optionally, time
scaling metadata information is output (as shown in FIG. 6), for the case of no
transients it would indicate that no pre-processing was performed. Time scaling
25 metadata may include, for example, time scaling parameters such as the location and
amount of time scaling performed and, if cross fading of spliced audio segments is
employed by the time scaling technique, the cross fade length. Metadata in the
encoded audio bit stream may also include information about transients, including

- 25 -

their location after and/or before and after temporal shifting. Audio data is output in step 212.

Audio Pre-Processing

FIGS. 8a-8e illustrate an example of audio time scaling pre-processing in accordance with aspects of the present invention when a transient exists in an audio coding block and is located closer to the last windowed block end than to the next windowed block end. For this example, a 50% block overlap is assumed, in the manner of FIGS. 1a-1e and FIGS. 4a and 4b. As discussed previously, to reduce the amount of transient pre-noise introduced by low bit rate audio coding, it is desired to adjust the time evolution of the input audio signal such that the audio signal transient is located closely following the last windowed block end. Such a shift in the transient location is preferred because it minimizes the disruption to the time evolution of the signal stream while optimally limiting the length of the transient pre-noise. However, as discussed above, a shift to a location closely following the next windowed block end also optimally limits the length of the transient pre-noise but does not minimize disruption to the signal stream's time evolution. In some cases the difference in disruption may be of little or no audible significance, particularly if time evolution compensation is also employed. Thus, a shift to either of the closest block ends is contemplated by the present invention in the present example and in other examples herein. As mentioned above, the transient time shifting time scaling need not be accomplished within a single block unless the processing is performed after the audio signal stream is divided into blocks by the encoder.

FIG. 8a shows three consecutive 50% overlapped windowed coding blocks. FIG. 8b shows the relationship between the original input audio data stream, containing a single transient and the windowed audio coding blocks. The onset of the transient is T samples after the preceding block end. Because the transient is closer to the preceding block end than the next block end, it is preferred to shift the transient to the left to a location closely following the preceding block end by applying time compression that has the effect of deleting T samples prior to the

- 26 -

transient. FIG. 8c shows two regions in the audio stream where audio time scaling may be performed. The first region corresponds to the audio samples before the transient where reducing the duration of the audio by T samples “slides” or shifts the position of the transient left to the desired location closely following the end of the preceding block by providing time compression. As in FIGS. 2A through 5B and other figures to be described, the spacing of the transient from the block end in FIGS. 8d and 8e is exaggerated in the figure for clarity of presentation. The second region shows the region where time scaling optionally may be performed after the transient to increase the duration of the audio by T samples by providing time expansion so that the overall length of the audio data remains at N samples. Although the deletion of T samples and the optional sample number compensating addition of T samples are both shown as occurring within a windowed audio coding sample block, this is not essential — the compensating time-scaling processing need not occur within a single audio coding block unless the transient time shifting is performed after the audio signal stream is divided into blocks by the encoder. The optimum location for such time-scaling processing may be determined by the time-scaling process employed. Because the transient may provide useful post-masking, sample number compensating time scaling preferably is done close to the transient.

FIG. 8d demonstrates the resulting signal stream if time scaling processing is performed on the input audio data stream by reducing the time duration of the audio input data stream by T samples in the area before the transient and no sample number compensating time scale expansion is performed after the transient signal. As discussed previously, slight variations in the temporal evolution of an audio signal are not discernable to most listeners. Therefore, if it is not required for the number of time scaled audio data stream samples to equal the number of input samples, N; it may be sufficient only to process the audio stream before the transient. FIG. 8e illustrates the case when the audio data stream before the transient is reduced in duration by T samples and the audio data stream following the transient is increased by T samples, thereby maintaining N audio samples in and out of the time scaling

- 27 -

processing block and restoring the time evolution of the audio signal stream except for the transient and portions of the signal stream close to the transient. The variations in lengths of the signal waveforms in FIGS. 8b-8e are intended to show schematically that the number of samples in the audio data stream varies for the
5 described conditions. When the number of audio samples is reduced, as in FIG. 8d, additional samples may need to be acquired before additional audio coding can be performed. This may mean reading more samples from a file or waiting for more audio to be buffered in a real-time system.

FIGS. 9a-9e illustrate an example of audio time scaling processing when a
10 transient exists in a windowed audio coding block and is located approximately T samples before a block end. To reduce the amount of transient pre-noise introduced by low bit rate audio coding while minimizing the transient shift, it is preferred to temporally adjust the input audio signal such that the audio signal transient closely follows the next block end. In the case of 50% overlapped blocks, a shift to the end
15 of the next block end (or the previous block end) limits the transient pre-noise to the first half of an audio coding block, instead of spreading the transient pre-noise throughout that block and the previous audio block.

FIG. 9a shows three consecutive 50% overlapped windowed coding blocks. FIG. 9b shows the relationship between the original input audio data,
20 containing a single transient and the audio blocks. The onset of the transient is T samples before the next block end. Because the transient is closer to the next block end than the previous block end, it is preferred to shift the transient to the right to a location closely following the next block end by applying time expansion that has the effect of adding T samples prior to the transient. FIG. 9c shows two regions where
25 audio time scaling may be performed. The first region corresponds to the audio samples before the transient where increasing the duration of the audio by T samples slides the position of the transient to the desired location closely after the next block end. FIG. 9c also shows the region where time scaling may be performed after the transient to reduce the duration of the audio by T samples so that the overall length of

- 28 -

the audio data stream, N samples, remains constant. FIG. 9d demonstrates the result if time scaling processing is performed on the input audio data stream by increasing the time duration of the audio input data stream by T samples in the time region before the transient but without performing sample number compensating time scale expansion after the transient signal. As discussed previously, slight variations in the temporal evolution of an audio signal are not discernable to most listeners.

Therefore, if it is not required for the number of audio stream samples after time scaling to equal the input, N. It may be sufficient only to process the audio before the transient.

FIG. 9e illustrates the case when the audio prior to the transient is increased in duration by T samples and the audio following the transient is reduced by T samples, thereby maintaining a constant number of audio samples before and after time scaling. As in other figures, the spacing of the transient from the block end in FIGS. 9d and 9e is exaggerated in the figures for clarity of presentation.

Audio Time Scaling Processing for Multiple Transients

Depending upon the length of the audio coding block size and the content of the audio data being coded, it is possible for an input audio data stream being processed to contain, within the N samples being processed, more than one transient signal that may introduce pre-noise artifacts. As mentioned above, the N samples being processed may include more than an audio coding block.

FIGS. 10a-10d illustrate processing solutions when two transients occur in an audio coding block. In general, two or more transients may be handled in the same manner as a single transient, with the earliest transient in the audio data stream being treated as the transient of interest.

FIG. 10a shows three consecutive 50% overlapped windowed coding blocks. FIG. 10b shows the case where two transients in the input audio straddle the end of an audio coding block. For this case, the earlier transient introduces the most perceptible pre-noise because a portion of the pre-noise resulting from the second transient is post-masked by the first transient. To minimize the pre-noise artifacts,

- 29 -

the input audio signal may be time scaled to shift the first transient to the right such that the audio before the first transient is time scale expanded by T samples, where T is the number of samples which places the first transient to a position closely following the next block end.

5 In order to sample number compensate for the time scale expansion processing before the first transient in FIG. 10b and to optimize post-masking of the pre-noise resulting from the second transient by moving the transients more closely together in time, the audio following the first transient and before the second transient preferably is time scaled to be reduced in duration by T samples. As illustrated in FIG. 10b,
10 there is sufficient audio processing data between the first and second transients to perform time scale processing. However, in some cases it may be that the second transient is so close to the first transient that there is not enough audio data to perform time scale processing between them. The amount of audio data required between transients is dependent upon the time scaling process used for the
15 processing. If insufficient audio data exists between the two transients, it may be necessary to time scale expand the audio data following the second transient in order to provide sample number compensation. In order to accomplish expansion of the audio data after the second transient, may be necessary for the time scaling process to have access to a larger segment of audio data than the number of samples in a block
20 used in the audio coding process, as mentioned above.

FIG. 10c illustrates the case when the first transient is closer to the last block end than the next block end and all of the transients (in this case two) are sufficiently close together that the pre-noise resulting from the first transient is substantially post-masked by the first transient. Thus, the audio stream prior to the first transient
25 preferably is time scale compressed by T samples so that the first transient is shifted to a location just after the prior block end. Sample number compensation to restore the original number of samples, in the form of time scale expansion, may be performed in the audio data stream following the second transient.

- 30 -

FIG. 10d illustrates the case when the first transient is closer to the next block end than the last block end and all of the transients (in this case, two) are sufficiently close together that the pre-noise resulting from the second is substantially post-masked by the first transient. Thus, the audio stream prior to the first transient is
5 time scale expanded by T samples so that the first transient is shifted to a location just after the next block end. Sample number compensation, in the form of time scale compression, optionally may be performed in the audio data stream following the second transient.

For the multiple transient case, if it is desired to time evolution compensate for
10 pre-processing in a near perfect manner, metadata information may be conveyed with each coded audio block in a manner similar to the single transient case described above.

Metadata Controlled Time Evolution Compensation of Time Scaling Pre-Processing

15 As mentioned above, it may be desirable to apply, subsequent to inverse transformation by the decoder, a compensating time scaling to the audio signal stream after the transient such that the time evolution of the processed audio signal stream is substantially the same as that of the original audio signal stream, thus restoring the original time evolution of the signal stream. However, experimental
20 studies have shown that slight temporal modifications of audio are not perceptible to most listeners and therefore time evolution compensation may not be necessary. Also, on average, transients are advanced and retarded equally and, thus, over a sufficiently long time period, the cumulative effect without time evolution compensation may be negligible. Another issue to be considered is that depending
25 upon the type of time scaling used for pre-processing, the additional time evolution compensating processing may introduce audible artifacts in the audio. Such artifacts may arise because time scaling processing, in many cases, is not a perfectly reversible process. In other words, reducing audio by a fixed amount using a time

- 31 -

scaling process and then time expanding the same audio later may introduce audible artifacts.

One benefit of processing audio that contains transient material by time scaling is that time scaling artifacts may be masked by the temporal masking properties of transient signals. An audio transient provides both forward and backward temporal masking. Transient audio material “masks” audible material both before and after the transient such that the audio directly preceding and following is not perceptible to a listener. Pre-masking has been measured and is relatively short and lasts only a few milliseconds while post-masking may last longer than 100 msec. Therefore, time-scaling time evolution compensation processing may be inaudible due to temporal post-masking effects. Thus, if performed, it is advantageous to perform time evolution compensation time-scaling within temporally masked regions.

FIGS. 11a-11f shows an example where intelligent time evolution compensation is performed following inverse transformation in the decoder using metadata information. The metadata greatly reduces the amount of analysis required to perform time evolution compensation because it indicates where time scaling processing should be performed and the duration of time scaling required. As explained above, time evolution compensating processing is intended to return the decoded audio signal to its original temporal evolution in which the signal stream, including the transient, has its original location in the audio stream. FIG. 11a shows three consecutive 50% overlapped windowed coding blocks. FIG. 11b shows an input audio stream prior to pre-processing having a transient T samples after a block end. FIG. 11c shows that the input audio stream is processed by deleting T samples prior to the transient to shift the transient to an earlier location. T samples are added after the transient in order to leave the number of audio data sample unchanged (sample number compensation). FIG. 11d shows the modified audio stream in which the transient is shifted to an earlier location and audio following the transient is shifted back to its original location. FIG. 11e shows the required time evolution compensating time scaling regions in which the deletion of T samples (time

- 32 -

compression) is compensated by adding T samples (time expansion) and the addition of T samples (time expansion) is compensated by deleting T samples (time compression). The result, shown in FIG. 11f is a compensated “near perfect” output signal having the same time evolution as the input signal of FIG. 11a (subject mainly to imperfections in the time scaling processes).

Time Scaling Post-Processing to Reduce Transient Pre-noise

As demonstrated in a number of previous examples, even with optimal placement of a transient in an audio coding block, some pre-noise is still introduced by the low bit rate audio coding system process. As was stated above, longer audio coding blocks are preferred over shorter coding blocks because they provide greater frequency resolution and increased coding gain. However, even if transients are optimally placed by time scaling prior to audio encoding (pre-processing), as the length of the audio coding block increases, the pre-noise also increases. Pre-masking of transient temporal pre-noise is on the order of 5 msec (milliseconds), which corresponds to 240 samples for audio sampled at 48 kHz. This implies that for coders with block sizes greater than approximately 512 samples, transient pre-noise begins to be audible even with optimal placement (only half is masked in the case of 50% overlapped block). (This does not take into account the reduction of transient pre-noise caused by windowing edge effects in the coder's blocks.)

While transient pre-noise may not be removed entirely from a low bit rate coding system, it is possible to perform time scaling post-processing (by itself or in addition to pre-processing) on audio data that has undergone inverse transformation in a transform-based low bit rate audio decoder to reduce the amount of transient pre-noise whether or not pre-processing is also applied. Time scaling post-processing may be performed either in conjunction with a low bit rate audio decoder (*i.e.*, as part of the decoder and/or by receiving metadata from the decoder and/or from the encoder via the decoder) or as a stand-alone post-process. Using metadata is preferred because useful information such as the location of transients in relation to audio coding blocks as well as the audio coding block length(s) are readily available

- 33 -

and may be passed to the post-processing process via the metadata. However, post-processing may be used without interaction with a low bit rate audio decoder. Both methods are discussed below.

5 **Time Scaling Post-Processing in Conjunction with
a Low Bit Rate Audio Decoder (Receiving Metadata)**

FIG. 12 is a flowchart of a process for performing time scaling post-processing in conjunction with a low bit rate audio decoder to reduce transient pre-noise artifacts. The process illustrated in FIG. 12 assumes that the input data is low bit rate encoded audio data (step 802). Following decoding of the compressed data into
10 audio (step 804), the audio corresponding to a block (or blocks) is sent to the time scaler 806 along with metadata information useful in reducing the transient pre-noise duration. This information may include, for example, the location of transients, the length of the audio coder block(s), the relation of the coder block boundaries to the audio data, and a desired length of the transient pre-noise. If the location of the
15 transients in relation to the audio coder's block borders is available, the length and location of the pre-noise artifact may be estimated and accurately reduced by post-processing. Since transients do provide some temporal pre-masking, it may not be necessary to completely remove the transient pre-noise. By giving the time scaling post-processing process a desired pre-noise length, some control over the amount of
20 pre-noise that is left in the output audio output by step 808 may be achieved. The results of suitable time-scale processing for step 806 is described below in connection with the description of FIGS. 13a-13c.

Note that post-processing may be useful whether or not pre-processing has been applied prior to encoding. Regardless of where the transient is located with
25 respect to block ends, some transient pre-noise exists. For example, at a minimum it is half the length of the audio coding window for the case of 50% overlap. Large window sizes still may introduce audible artifacts. By performing post processing, it is possible to reduce the length of the pre-noise even more than it was reduced by

- 34 -

optimally placing the transient with respect to block ends prior to quantization by the encoder.

FIGS. 13a-13c illustrate an example of post-processing for a single transient to reduce the pre-noise artifact present after inverse transformation. As shown in FIG. 13a, a single transient introduces a pre-noise artifact.. Depending on the coding block length, the pre-noise, even after pre-processing, if any, may have a longer time than may be masked by transient temporal pre-masking effects. However, as shown in FIG. 13b, by using the transient location metadata information from the decoder, one may identify a region of audio containing the pre-noise in which the pre-noise may be reduced in length by time scaling the audio to reduce the pre-noise by T samples. The number T may be chosen such that the pre-noise length is minimized to take advantage of pre-masking or may be chosen so as to remove the pre-noise completely or nearly completely. If it is desired to maintain the same number of samples as in the original signal, the audio following the transient may be time scale expanded by +T samples. Alternatively, as shown in connection with the example of FIG. 16A, such sample number compensation may be applied prior to the pre-noise, which has the advantage of also providing time evolution compensation.

It should be noted that if post-processing is performed in conjunction with time scaling pre-processing, one may minimize the amount of further disruption to the output audio stream's time evolution. Since the time scaling pre-processing discussed earlier reduces the length of the pre-noise to $N/2$ samples for the case of 50% block overlap (where N is the length of the audio coding block) one is guaranteed to introduce less than $N/2$ samples of further time evolution disruption in the output audio as compared to the original input audio. In the absence of pre-processing, the pre-noise can be up to N samples, the coding block length, for the case of 50% block overlap.

In some low bit rate audio coding systems, the location of the signal transients may not be readily available if the encoder does not convey the location information. If that is the case, the decoder or the time scaling process may, using any number of

transient detection processes or the efficient method described previously, perform transient detection.

For multiple transients, the same issues apply as for pre-processing, as discussed above.

5

As mentioned above, in some cases it may be desired to improve the perceived quality of audio that has undergone low bit rate coding using compression systems that do not implement transient pre-noise time scaling processing (pre-processing).

10

15

20

25

- 36 -

The fifth and sixth steps 1410 in processing involve estimating the location and duration of the transient pre-noise artifacts and reducing their length with time scaling processing 1412. Since, by definition, the pre-noise artifacts are limited to the regions preceding transients in the audio data, the search area is limited by the information provided by the transient detection processing. As shown in FIG. 1, the length of the pre-noise is limited from a minimum of $N/2$ to a maximum of N samples where N is the number of audio samples in a 50% overlapped audio coding block. Thus, when N is 1024 samples and audio is sampled at 48 kHz, transient pre-noise may range from 10.7 msec to 21.3 msec before the onset of the transient, depending on the transient location in the audio stream, which significantly exceeds any temporal masking that may be expected from transient signals. Alternatively, instead of estimating the length of the pre-noise artifacts preceding a transient, step 1410 may apply assume that the pre-noise artifacts have a default length.

Two approaches for transient pre-noise reduction may be implemented. The first assumes that all transients contain pre-noise and therefore the audio before every transient may be time scaled (time compressed) by a predetermined (default) amount that is based on an expected amount of pre-noise per transient. If this technique is used, time scale expansion of the audio prior to the temporal pre-noise may be done to provide both sample number compensation for the time compression time scaling processing employed to reduce the length of the pre-noise and to provide time evolution compensation (time expansion prior to the pre-noise that compensates for time compression within the pre-noise leaves the transient in or nearly in its original temporal location). However, if the exact location of the start of the pre-noise is not known, such sample number compensation processing may unintentionally increase the duration of parts of the pre-noise component.

FIGS. 15a-15c demonstrate a technique that uses a default value to time-scale the audio before each transient to reduce pre-noise duration but does not perform sample number compensation. As shown in FIG. 15a, an audio signal stream from a low bit rate audio decoder has a transient preceded by pre-noise. FIG. 15b shows a

- 37 -

default processing length used as the amount of time compression to be performed by the time scaling processing. FIG. 15c shows the resulting audio signal stream having reduced pre-noise. In this example, time evolution compensation is not performed to return the transient to its original location in the audio data stream. However, in a manner similar to previous processing examples, if a constant number of input to output samples are desired, time scale expansion processing following the transient may be performed, similar to the example of FIG. 13b or, possibly, before the pre-noise as described below in connection with the example of FIGS. 16a-16c.

However, when applying a default processing length, providing such compensation prior to the pre-noise runs the risk of performing the time scale expansion processing within the pre-noise (thus, undesirably increasing the pre-noise length) if the actual length of the pre-noise exceeds the default length. Moreover, in some cases, the post-processing may not have access to the audio stream prior to the pre-noise – the audio may already be output in order to reduce latency.

A second post-processing pre-noise reduction technique, illustrated in FIGS. 16a-16c, involves performing analysis of the pre-noise resulting from a transient to determine its length and processing the audio so that only the pre-noise segment is processed. As noted above, transient pre-noise is produced when the high frequency components of transient audio material are temporally smeared throughout a block as a result of the quantizing process in the encoder. Therefore one straight-forward method of detection is to high pass filter the audio prior to a transient and measure the high frequency energy. The start of the transient pre-noise is identified when the noise-like, high frequency pre-noise related to and caused by the transient exceeds a predetermined threshold. When the size and location of the transient pre-noise is known, compensating time scale expansion of the audio may be performed prior to time scale reduction of the pre-noise to return the audio to its original temporal evolution and to restore the time evolution of the audio stream substantially to its original condition. The invention is not limited to employing high frequency

- 38 -

detection. Other techniques for detecting or estimating the length of the pre-noise may be employed.

In FIG. 16a, an audio signal stream from a low bit rate audio decoder has a transient preceded by pre-noise. FIG. 16b shows a time compression processing length used as the amount of time scale reduction to be performed by the time scaling processing based on an estimated pre-noise length as measured by the high frequency audio content in the block. FIG. 16b also shows the use of time expansion by T samples in order to restore the original time evolution of the signal stream and also to restore the original number of samples. FIG. 16c shows the resulting audio signal stream having reduced pre-noise along with the original time evolution and the same number of samples as the original signal stream.

The present invention and its various aspects may be implemented as software functions performed in digital signal processors, programmed general-purpose digital computers, and/or special purpose digital computers. Interfaces between analog and digital signal streams may be performed in appropriate hardware and/or as functions in software and/or firmware.

AMENDED CLAIMS

[received by the International Bureau on 04 October 2002 (04.10.02);
original claims 1-38 replaced by amended claims 1-37 (7 pages)]

1. A method for reducing distortion artifacts preceding a signal transient in an audio signal stream processed by a transform-based low-bit-rate audio coding system
5 employing coding blocks, comprising
detecting a transient in the audio signal stream prior to processing by said coding system, and
shifting the temporal relationship of said transient with respect to said coding blocks by time scaling a segment of said audio signal stream preceding said signal
10 transient such that the time duration of said distortion artifacts is reduced.
2. The method of claim 1 wherein said shifting shifts the temporal relationship of said transient with respect to said coding blocks prior to forward transforming in the encoder of said coding system.
15
3. The method of claim 2 wherein said transient is shifted to a temporal position closely following the next block end or closely following the last block end.
4. The method of claim 3 wherein said transient is shifted to a temporal
20 position closely following the next block end or closely following the last block end which results in the shorter shift of temporal position.
5. A method according to any one of claims 1-4 further comprising removing at least a portion of remaining distortion artifacts after inverse transformation in the
25 decoder of said coding system.
6. The method of claim 5 wherein the portion of remaining distortion artifacts is determined at least in part by metadata information carried in said coding system.

7. The method of claim 5 wherein the portion of remaining distortion artifacts is determined at least in part by a default parameter.

8. The method of claim 5 wherein the portion of remaining distortion artifacts
5 is determined at least in part by a measure of high frequency audio components in said audio signal stream.

9. The method of claim 1 further comprising applying a compensating time scaling to the audio signal stream subsequent to inverse transformation in the decoder
10 of said coding system such that the time evolution of the processed audio signal stream is substantially the same as that of the audio signal stream prior to said shifting.

10. The method of claim 9 wherein said compensating time scaling is applied
15 to a segment of said audio signal stream preceding said signal transient.

11. The method of claim 9 wherein said coding system includes an encoder and a decoder, said encoder transmitting metadata to said decoder along with an encoded version of said audio signal stream, said metadata including information
20 useful for applying said compensating time scaling.

12. The method of claim 1 wherein said time scaling is performed on a segment of said audio stream closely preceding said transient.

25 13. The method of claim 12 wherein said time scaling is performed on a segment of said audio stream that is at least partially temporally pre-masked by transient.

14. The method of claim 1 wherein said time scaling has the effect of deleting signal components from or adding signal components to the audio signal stream applied to the coding system.

5 15. The method of claim 14 wherein a further time scaling is applied following said signal transient, said further time scaling acting in the opposite sense to the said first-recited time scaling.

16. The method of claim 15 wherein said further time scaling is applied prior
10 to forward transforming in the encoder of said coding system.

17. The method of claim 15 wherein said further time scaling is applied subsequent to inverse transformation in the decoder of said coding system.

15 18. The method of claim 15 wherein the time duration of the signal components added or deleted by said further time scaling is substantially the same as the time duration of signal components deleted or added by said first-recited time scaling, respectively, whereby the time duration of said audio signal stream is substantially unchanged.

20

19. The method of claim 14 further comprising applying compensating time scaling to the audio signal stream preceding said distortion artifacts, which precede said transient, and subsequent to inverse transformation in the decoder of said coding system such that the time evolution of the processed audio signal stream is
25 substantially the same as that of the audio signal stream prior to said shifting and the time duration of said audio signal stream is substantially unchanged.

20. The method of claim 19 wherein said coding system includes an encoder and a decoder, said encoder transmitting metadata to said decoder, said metadata including information useful for applying said compensating time scalings.

5 21. The method of claim 1 wherein said audio signal stream applied to the coding system is a digital signal stream in which the audio information is represented by samples, the order of said samples representing time, and wherein said time scaling has the effect of deleting samples from or adding samples to the digital signal stream applied to the coding system.

10 22. The method of claim 1 wherein a further time scaling is applied following said signal transient, said further time scaling acting in the opposite sense to the said first-recited time scaling.

15 23. The method of claim 22 wherein said further time scaling is performed on a segment of said audio stream closely following said transient.

20 24. The method of claim 23 wherein said time scaling is performed on a segment of said audio stream that is at least partially temporally post-masked by transient.

25 25. The method of claim 22 wherein said first-recited time scaling has the effect of deleting signal components from or adding signal components to the audio signal stream applied to the coding system and said further time scaling has the effect of adding signal components to the audio signal stream when said first-recited time scaling deletes signal components and said further time scaling has the effect of deleting signal components to the audio signal stream when said first-recited time scaling adds signal components.

26. The method of claim 25 wherein the time duration of the signal components added or deleted by said further time scaling is substantially the same as the time duration of signal components deleted or added by said first-recited time scaling, respectively, whereby the time duration of said audio signal stream is substantially unchanged.

27. The method of claim 22 wherein said audio signal stream applied to the coding system is a digital signal stream in which the audio information is represented by samples, the order of said samples representing time, and wherein said first-recited time scaling has the effect of deleting samples from or adding samples to the digital signal stream applied to the coding system and said further time scaling has the effect of adding samples to the digital signal stream when said first-recited time sampling deletes samples from the digital signal stream and said further time scaling has the effect of deleting samples from the digital signal stream when said first-recited time sampling adds samples to the digital signal stream.

28. The method of claim 1 wherein said detecting detects multiple transients and said shifting shifts the temporal location of the first of said transients to reduce distortion artifacts prior to the first of said transients.

29. The method of claim 28 wherein the temporal location of the first of said transients with respect to said coding blocks is shifted by time scaling said audio signal stream preceding the first of said signal transients.

30. The method of claim 29 wherein a further time scaling is applied following the first of said transients and before one or more other of said multiple

transients, said further time scaling acting in the opposite sense to the said first-recited time scaling.

31. The method of claim 29 wherein a further time scaling is applied
5 following said transients, said further time scaling acting in the opposite sense to the said first-recited time scaling.

32. In a decoder of a transform-based low-bit-rate audio coding system
employing coding blocks, a method for reducing distortion artifacts preceding a
10 signal transient in an audio signal stream subsequent to inverse transformation,
comprising
detecting a transient in the audio signal stream, and
time compressing at least a portion of said distortion artifacts such that the
time duration of said distortion artifacts is reduced.

15

33. The method of claim 32 wherein the portion of the distortion artifacts is
determined at least in part by the location of the detected transient and a default
parameter.

20







34. The method of claim 32 the portion of the distortion artifacts is
determined at least in part by the location of the detected transient and signal
characteristics preceding said transient.

35. The method of claim 34 wherein said signal characteristics include a
25 measure of high-frequency components of the audio signal stream.





36. The method of claim 33 or 34 further comprising time expanding prior to said time compression such that the time evolution and length of the audio signal stream is substantially unchanged.

5 37. The method of claim 33 or 34 further comprising time expanding subsequent to said time compression such that the length of the audio signal stream is substantially unchanged.




1 / 8

FIG._1A**FIG._1B****FIG._1C****FIG._1D****FIG._1E****FIG._2A****FIG._2B**




2 / 8





3 / 8

**FIG. 5A****FIG. 5B****FIG. 6**



4 / 8

**FIG._7****FIG._8A****FIG._8B****FIG._8C**







-T samples +T samples

**FIG._8D****FIG._8E**







5 / 8

FIG._9A**FIG._9B****FIG._9C**


+T samples -T samples

FIG._9D**FIG. 9E****FIG._10A****FIG._10B****FIG._10C****FIG._10D**



6 / 8

FIG._11A**FIG._11B****FIG._11C****FIG._11D****FIG._11E****FIG._11F****FIG._12**

7 / 8

FIG. 13A**FIG. 13B****FIG. 13C**

Reduced pre-noise

**FIG. 14**

8 / 8

