

**COMPANDING APPARATUS AND METHOD TO REDUCE QUANTIZATION
NOISE USING ADVANCED SPECTRAL EXTENSION**

CROSS-REFERENCE TO RELATED APPLICATIONS

- 5 **[0001]** This application claims priority to United States Provisional Patent Application Nos. 61/809,028 filed 5 April 2013 and 61/877,167 filed 12 September 2013, which are hereby incorporated by reference in its entirety.

FIELD OF THE INVENTION

- 10 **[0002]** One or more embodiments relate generally to audio signal processing, and more specifically to reducing coding noise in audio codecs using compression/expansion (companding) techniques.

BACKGROUND

- 15 **[0003]** Many popular digital sound formats utilize lossy data compression techniques that discard some of the data to reduce storage or data rate requirements. The application of lossy data compression not only reduces the fidelity of source content (e.g., audio content), but it can also introduce noticeable distortion in the form of compression artifacts. In the context of audio coding systems, these sound artifacts are called coding noise or quantization noise.

- 20 **[0004]** Digital audio systems employ codecs (coder-decoder components) to compress and decompress audio data according to a defined audio file format or streaming media audio format. Codecs implement algorithms that attempt to represent the audio signal with a minimum number of bits while retaining as high a fidelity as possible. The lossy compression techniques typically used in audio codecs work on a psychoacoustic model of human hearing perception. The audio formats usually involve the use of a time/frequency domain transform (e.g., a modified discrete cosine transform - MDCT), and use masking effects, such as frequency masking or temporal masking so that certain sounds, including any apparent quantization noise is hidden or masked by actual content.

- 25 **[0005]** Most audio coding systems are frame based. Within a frame, audio codecs normally shape the coding noise in the frequency domain so that it becomes least audible.
- 30 Several present digital audio formats utilize frames of such long durations that a frame may contain sounds of several different levels or intensities. Since the coding noise is usually stationary in level over the evolution of a frame, coding noise may be most audible during low intensity parts of the frame. Such an effect may be manifested as pre-echo distortion in which silence (or low-level signal) preceding a high intensity segment is swamped by noise

in the decoded audio signal. Such an effect may be most noticeable in transient sounds or impulses from percussion instruments, such as castanets or other sharp percussive sound sources. Such distortion is typically caused by the quantization noise introduced in the frequency domain being spread over the entire transform window of the codec in the time domain.

[0006] Present measures to avoid or minimize pre-echo artifacts include the use of filters. Such filters, however, introduce phase distortion and temporal smearing. Another possible solution includes the use of smaller transform windows, however this approach can significantly reduce frequency resolution.

[0007] The subject matter discussed in the background section should not be assumed to be prior art merely as a result of its mention in the background section. Similarly, a problem mentioned in the background section or associated with the subject matter of the background section should not be assumed to have been previously recognized in the prior art. The subject matter in the background section merely represents different approaches, which in and of themselves may also be inventions.

BRIEF SUMMARY OF EMBODIMENTS

[0008] Embodiments are directed to a method of processing a received audio signal by expanding the audio signal to an expanded dynamic range through a process that includes dividing the received audio signal into a plurality of time segments using a defined window shape, calculating a wideband gain for each time segment in the frequency domain using a non-energy based average of a frequency domain representation of the audio signal, and applying the gain value to each time segment to obtain the expanded audio signal. The gain values of the wideband gain applied to each time segment are selected to have the effect of amplifying segments of relatively high intensity and attenuating segments of relatively low intensity. For this method, the received audio signal comprises an original audio signal that was compressed from an original dynamic range through a compression process including dividing the original audio signal into a plurality of time segments using a defined window shape, calculating a wideband gain in the frequency domain using a non-energy based average of frequency domain samples of the initial audio signal, and applying the wideband gain to the original audio signal. In the compression process, the gain values of the wideband gain applied to each time segment are selected to have the effect of amplifying segments of relatively low intensity and attenuating segments of relatively high intensity. The expansion process is configured to substantially restore the dynamic range of the initial audio signal,

and the wideband gain of the expansion process may be substantially the inverse of the wideband gain of the compression process.

[0009] In a system that implements a method of processing a received audio signal by an expansion process, a filterbank component may be used to analyze the audio signal to obtain its frequency domain representation, and the defined window shape for segmentation into the plurality of time segments may be the same as the prototype filter for the filterbank.

Likewise, in a system that implements a method of processing a received audio signal by a compression process, a filterbank component may be used to analyze the original audio signal to obtain its frequency domain representation, and the defined window shape for

segmentation into the plurality of time segments may be the same as the prototype filter for the filterbank. The filterbank in either case may be one of a QMF bank or a short-time Fourier transform. In this system, a received signal for the expansion process is obtained after modification of the compressed signal by an audio encoder that generates a bitstream, and a decoder that decodes the bitstream. The encoder and decoder may comprise at least part of a transform-based audio codec. The system may further comprise components that process control information that is received through the bitstream and determines an activation state of the expansion process.

BRIEF DESCRIPTION OF THE DRAWINGS

[0010] In the following drawings like reference numbers are used to refer to like elements. Although the following figures depict various examples, the one or more implementations are not limited to the examples depicted in the figures.

[0011] FIG. 1 illustrates a system for compressing and expanding an audio signal in a transform-based audio codec, under an embodiment.

[0012] FIG. 2A illustrates an audio signal divided into a plurality of short time segments, under an embodiment.

[0013] FIG. 2B illustrates the audio signal of FIG. 2A after the application of wideband gain over each of the short time segments, under an embodiment.

[0014] FIG. 3A is a flowchart illustrating a method of compressing an audio signal, under an embodiment.

[0015] FIG. 3B is a flowchart illustrating a method of expanding an audio signal, under an embodiment.

[0016] FIG. 4 is a block diagram illustrating a system for compressing an audio signal, under an embodiment.

[0017] FIG. 5 is a block diagram illustrating a system for expanding an audio signal, under an embodiment.

[0018] FIG. 6 illustrates the division of an audio signal into a plurality of short time segments, under an embodiment.

5

DETAILED DESCRIPTION

[0019] Systems and methods are described for the use of companding techniques to achieve temporal noise shaping of quantization noise in an audio codec. Such embodiments include the use of a companding algorithm implemented in the QMF-domain to achieve
10 temporal shaping of quantization noise. Processes include encoder control of the desired decoder companding level, and extension beyond monophonic applications to stereo and multi-channel companding.

[0020] Aspects of the one or more embodiments described herein may be implemented in an audio system that processes audio signals for transmission across a network that includes
15 one or more computers or processing devices executing software instructions. Any of the described embodiments may be used alone or together with one another in any combination. Although various embodiments may have been motivated by various deficiencies with the prior art, which may be discussed or alluded to in one or more places in the specification, the embodiments do not necessarily address any of these deficiencies. In other words, different
20 embodiments may address different deficiencies that may be discussed in the specification. Some embodiments may only partially address some deficiencies or just one deficiency that may be discussed in the specification, and some embodiments may not address any of these deficiencies.

[0021] FIG. 1 illustrates a companding system for reducing quantization noise in a codec-based audio processing system, under an embodiment. FIG. 1 illustrates an audio signal
25 processing system that is built around an audio codec comprising encoder (or “core encoder”) 106 and decoder (or “core decoder”) 112. The encoder 106 encodes audio content into data stream or signal for transmission over network 110 where it is decoded by decoder 112 for playback or further processing. In an embodiment, the encoder 106 and decoder 112 of the
30 codec implement a lossy compression method to reduce the storage and/or data rate requirements of the digital audio data, and such a codec may be implemented as a MP3, Vorbis, Dolby Digital (AC-3), AAC, or similar codec. The lossy compression method of the codec creates coding noise that generally is stationary in level over the evolution of a frame defined by the codec. Such coding noise is often most audible during low intensity parts of a

frame. System 100 includes components that reduce the perceived coding noise in existing coding systems by providing a compression pre-step component 104 prior to the core encoder 106 of the codec and an expansion post-step component 114 operating on the core decoder 112 output. The compression component 104 is configured to divide the original audio input signal 102 into a plurality of time segments using a defined window shape, calculate and apply a wideband gain in the frequency domain using a non-energy based average of frequency domain samples of the initial audio signal, wherein the gain values applied to each time segment amplify segments of relatively low intensity and attenuate segments of relatively high intensity. This gain modification has the effect of compressing or significantly reducing the original dynamic range of the input audio signal 102. The compressed audio signal is then coded in encoder 106, transmitted over network 110 and decoded in decoder 112. The decoded compressed signal is input to expansion component 114, which is configured to perform the inverse operation of the compression pre-step 104 by applying inverse gain values to each time segment to expand the dynamic range of the compressed audio signal back to the dynamic range of the original input audio signal 102. Thus, the audio output signal 116 comprises an audio signal having the original dynamic range, with the coding noise removed through the pre- and post-step companding process.

[0022] As shown in FIG. 1, a compression component or compression pre-step 104 is configured to reduce the dynamic range of the audio signal 102 input to the core encoder 106. The input audio signal is divided into a number of short segments. The size or length of each short segment is a fraction of the frame size used by the core encoder 106. For example, a typical frame size of the core coder may be on the order of 40 to 80 milliseconds. In this case, each short segment may be on the order of 1 to 3 milliseconds. The compression component 104 calculates an appropriate wideband gain value to compress the input audio signal on a per-segment basis. This is achieved by modifying short segments of the signal by an appropriate gain value for each segment. Relatively large gain values are selected to amplify segments of relatively low intensity, and small gain values are selected to attenuate segments of high intensity.

[0023] FIG. 2A illustrates an audio signal divided into a plurality of short time segments, under an embodiment, and FIG. 2B illustrates the same audio signal after application of wideband gain by a compression component. As shown in FIG. 2A, audio signal 202 represents a transient or sound impulse such as may be produced by a percussive instrument (e.g., castanets). The signal features a spike in amplitude as shown in the plot of voltage, V , versus time, t . In general, the amplitude of the signal is related to the acoustic energy or

intensity of the sound and represents a measure of the sound's power at any point in time.

When the audio signal 202 is processed through a frame-based audio codec, portions of the signal are processed within transform (e.g., MDCT) frames 204. Typical present digital

audio systems utilize frames of relatively long duration, so that for sharp transient or short

5 impulse sounds, a single frame may include sounds of low intensity as well high intensity.

Thus, as shown in FIG. 1, the single MDCT frame 204 includes the impulse portion (peak) of

the audio signal as well as a relatively large amount of low intensity signal before and after

the peak. In an embodiment, a compression component 104 divides the signal into a number

of short time segments, 206, and applies a wideband gain to each segment in order to

10 compress the dynamic range of the signal 202. The number and size of each short segment

may be selected based on application needs and system constraints. Relative to the size of an

individual MDCT frame, the number of short segments may range from 12 to 64 segments,

and may typically comprise 32 segments, but embodiments are not so limited.

[0024] FIG. 2B illustrates the audio signal of FIG. 2A after the application of wideband

15 gain over each of the short time segments, under an embodiment. As shown in FIG. 2B,

audio signal 212 has the same relative shape as the original signal 202, however, the

amplitude of the low intensity segments has been increased by application of amplifying gain

values, and the amplitude of the high intensity segments has been decreased by application of
attenuating gain values.

20 **[0025]** The output of the core decoder 112 is the input audio signal with reduced dynamic
range (e.g., signal 212) plus quantization noise introduced by the core encoder 106. This

quantization noise features an almost uniform level across time within each frame. The

expansion component 114 acts on the decoded signal to restore the dynamic range of the
original signal. It uses the same short time resolution based on the short segment size 206

25 and inverts the gains applied in the compression component 104. Thus, the expansion

component 114 applies a small gain (attenuation) on segments that in the original signal had

low intensity, and had been amplified by the compressor, and applies a large gain

(amplification) on segments that in the original signal had high intensity and had been

attenuated by the compressor. The quantization noise added by the core coder, that had a

30 uniform time envelope, is thus concurrently shaped by the post-processor gain to

approximately follow the temporal envelope of the original signal. This processing

effectively renders the quantization noise less audible during quiet passages. Although the

noise may be amplified during passages of high intensity, it remains less audible due to the

masking effect of the loud signal of the audio content itself.

[0026] As shown in FIG. 2A, the companding process modifies discrete segments of the audio signal individually with respective gain values. In certain cases, this can result in discontinuities at the output of the compression component that can cause problems in the core encoder 106. Likewise, discontinuities in gain at the expansion component 114 could result in discontinuities in the envelope of the shaped noise, which could result in audible clicks in the audio output 116. Another issue related to application of individual gain values to short segments of the audio signal is based on the fact that typical audio signals are a mixture of many individual sources. Some of these sources may be stationary across time, and some may be transients. A stationary signal is generally constant in their statistical parameters over time, whereas transient signals are generally not constant. Given the broadband nature of transients, their fingerprint in such a mixture is usually more visible at the higher frequencies. A gain calculation that is based on the short-term energy (RMS) of the signal tends to be biased towards the stronger low frequencies and hence is dominated by the stationary sources, and exhibits little variation across time. Thus, this energy-based approach is generally ineffective in shaping the noise introduced by the core encoder.

[0027] In an embodiment, system 100 calculates and applies the gain at the compression and expansion components in a filter-bank with a short prototype filter in order to resolve the potential issues associated with the application of individual gain values. The signal to be modified (the original signal at the compression component 104, and the output of the core decoder 112 in the expansion component 114) is first analyzed by the filter-bank and the wideband gain is applied directly in the frequency domain. The corresponding effect in the time domain is to naturally smooth the gain application according to the shape of the prototype filter. This resolves the issues of discontinuities described above. The modified frequency domain signal is then converted back to the time domain via a corresponding synthesis filter-bank. Analyzing the signal with a filterbank provides access to its spectral content, and allows the calculation of a gain that preferentially boosts the contribution due to the high frequencies (or to boost the contribution due to any spectral content that is weak), providing gain values that are not dominated by the strongest components in the signal. This resolves the problem associated with audio sources that comprise a mixture of different sources, as described above. In an embodiment, the system calculates the gain using a p-norm of the spectral magnitudes where p is typically less than 2 ($p < 2$). This enables more emphasis to the weak spectral content, as compared to when it is based on energy ($p=2$).

[0028] As stated above, the system includes a prototype filter to smooth the gain application. In general, a prototype filter is the basic window shape in a filterbank, which is

modulated by sinusoidal waveforms to get the impulse responses for the different subband filters in the filterbanks. For instance, a short-time Fourier transform (STFT) is a filterbank, and each frequency line of this transform is a subband of the filterbank. The short-time Fourier transform is implemented by multiplying a signal with a window shape (an N-sample window), which could be rectangular, Hann, Kaiser-Bessel derived (KBD), or some other shape. The windowed signal is then subject to a discrete Fourier transform (DFT) operation, to obtain the STFT. The window shape in this case is the prototype filter. The DFT is composed of sinusoidal basis functions, each of a different frequency. The window shape multiplied by a sinusoidal function then provides the filter for the subband corresponding to that frequency. Since the window shape is the same at all frequencies, it is referred to as a “prototype”.

[0029] In an embodiment, the system utilizes a QMF (Quadrature Modulated Filter) bank for the filterbank. In a particular implementation, the QMF bank may have a 64-pt window, which forms the prototype. This window modulated by cosine and sine functions (corresponding to 64 equally spaced frequencies) forms the subband filters for the QMF bank. After each application of the QMF function, the window is moved over by 64 samples, i.e., the overlap between time segments in this case is $640 - 64 = 576$ samples. However although the window shape spans ten time segments in this case ($640 = 10 \times 64$), the main lobe of the window (where its sample values are very significant) is about 128 samples long.

Thus, the effective length of the window is still relatively short.

[0030] In an embodiment, the expansion component 114 ideally inverts the gains applied by the compression component 104. Although it is possible to transmit the gains applied by the compression component through the bitstream to the decoder, such an approach would typically consume a significant bit-rate. In an embodiment, system 100 instead estimates the gains required by the expansion component 114 directly from the signal available to it, i.e., the output of the decoder 112, which effectively requires no additional bits. The filterbank at the compression and expansion components are selected to be identical in order to calculate gains that are inverses of each other. In addition, these filterbanks are time synchronized so that any effective delays between the output of the compression component 104 and the input to the expansion component 114 are multiple of the stride of the filterbank. If the core encoder-decoder were lossless, and the filterbank provides perfect reconstruction, the gains at the compression and expansion components would be exact inverses of each other, thus allowing for exact reconstruction of the original signal. In practice, however, the gain

applied by the expansion component 114 is only a close approximation of the inverse of the gain applied by the compression component 104.

[0031] In an embodiment, the filterbank used in the compression and expansion components is a QMF bank. In a typical use application, a core audio frame could be 4096 samples long with an overlap of 2048 with the neighboring frame. At 48kHz such a frame would be 85.3 milliseconds long. In contrast, a QMF bank that is used may have a stride of 64 samples (which is 1.3 ms long), which provides a fine temporal resolution for the gains. Further, the QMF has a smooth prototype filter that is 640 samples long ensuring that the gain application varies smoothly across time. Analysis with this QMF filterbank provides a time-frequency tiled representation of the signal. Each QMF time-slot is equal to a stride and in each QMF time-slot there are 64 uniformly spaced subbands. Alternatively, other filterbanks could be employed, such as a short term Fourier transform (STFT), and such a time-frequency tiled representation could still be obtained.

[0032] In an embodiment, the compression component 104 performs a pre-processing step that scales the codec input. For this embodiment, $S_t(k)$ is a complex valued filter bank sample at time slot t and frequency bin k . FIG. 6 illustrates the division of an audio signal into a number of time slots for a range of frequencies, under an embodiment. For the embodiment of diagram 600, there are 64 frequency bins k , and 32 time slots t that produce a plurality of time-frequency tiles as shown (though not necessarily drawn to scale). The compression pre-steps scales the codec input to become $S'_t(k) = S_t(k)/g_t$. In this equation, $g_t = (\bar{S}_t/S_0)^\gamma$ is a normalized slot mean.

[0033] In the above equation, the expression, $\bar{S}_t = \frac{1}{K} \sum_{k=1}^K |S_t(k)|$ is the *mean absolute level/1-norm* and S_0 is a suitable constant. A generic p-norm is defined in this context as follows:

$$\bar{S}_t = \left(\frac{1}{K} \sum_{k=1}^K |S_t(k)|^p \right)^{1/p}$$

[0034] It has been shown that the 1-norm may give significantly better results than using the energy (rms/2-norm). The value of the exponent term γ is typically in the range of between 0 and 1, and may be chosen to be 1/3. The constant S_0 ensures reasonable gain values independent of the implementation platform. For instance, it may be 1 when implemented in a platform where all the $S_t(k)$ values might be limited in absolute value to 1. It could potentially be different in a platform where $S_t(k)$ may have a different maximum absolute value. It could also be used to make sure that mean gain value across a large set of

signals is close to 1. That is, it could be an intermediate signal value between a maximum signal value and a minimum signal value determined from large corpora of content.

[0035] In the post-step process performed by the expansion component 114, the codec output is expanded by an inverse gain applied by the compression component 104. This requires an exact or near-exact replica of the filter bank of the compression component. In this case, $\tilde{S}_t(k)$ represents a complex valued sample of this second filter bank. The expansion component 114 scales the codec output to become $\tilde{S}'_t(k) = \tilde{S}_t(k) \cdot \tilde{g}_t$.

[0036] In the above equation $\tilde{g}_t$ is a normalized slot mean given as:

$$\tilde{g}_t = (\bar{\tilde{S}}_t/S_0)^{\gamma/(1-\gamma)}$$

10 and

$$\bar{\tilde{S}}_t = (\frac{1}{K} \sum_{k=1}^K |\tilde{S}_t(k)|^p)^{1/p}$$

[0037] In general, the expansion component 114 will use the same p-norm as used in the compression component 104. Thus if the mean absolute level is used to define $\bar{\tilde{S}}_t$ in the

15 compression component 104, $\bar{\tilde{S}}_t$ is also defined using the 1-norm (p=1) in the above equation.

[0038] When a complex filterbank (comprising of both cosine and sine basis functions), such as the STFT or the complex-QMF is used in the compression and expansion components, the calculation of the magnitude, $|\tilde{S}_t(k)|$ or $|S_t(k)|$ of a complex subband sample requires a computationally intensive square-root operation. This can be circumvented by approximating the magnitude of the complex subband sample in a variety of ways, for instance, by summing up the magnitude of its real and imaginary parts.

[0039] In the above equations, the value K is equal to the number of subbands in the filterbank, or lower. In general, the p-norm could be calculated using any subset of the subbands in the filterbank. However, the same subset should be employed at both the encoder 106 and decoder 112. In an embodiment, the high frequency portions (e.g., audio components above 6 kHz) of the audio signal could be coded with an advanced spectral extension (A-SPX) tool. Additionally it may be desirable to use only the signal above 1 kHz (or a similar frequency) to guide the noise-shaping. In such a case only those subbands in the range 1 kHz to 6 kHz may be used to calculate p-norm, and hence the gain value.

30 Furthermore, although the gain is calculated from one subset of subbands it could still be applied to a different, and possibly larger, subset of subbands.

[0040] As shown in FIG. 1, the companding function to shape quantization noise introduced by the core encoder 106 of an audio codec is performed two separate components 104 and 114 performing certain pre-encoder compression functions and post-decoder expansion functions. FIG. 3A is a flowchart illustrating a method of compressing an audio signal in a pre-encoder compression component, under an embodiment, and FIG. 3B is a flowchart illustrating a method of expanding an audio signal in a post-decoder expansion component, under an embodiment.

[0041] As shown in FIG. 3A, process 300 begins with the compression component receiving the input audio signal (302). This component then divides the audio signal into short time-segments (304) and compresses the audio signal to a reduced dynamic range by applying wideband gain values to each of the short segments (306). The compression component also implements certain prototype filtering and QMF filterbank components to reduce or eliminate any discontinuities caused by applying different gain values to contiguous segments, as described above (308). In certain cases, such as based on the type of audio content or certain characteristics of the audio content, compression and expansion of the audio signal before and after the encode/decode stages of the audio codec may degrade rather than enhance the output audio quality. In such instances, the companding process may be turned off, or modified to return different companding (compression/expansion) levels. Thus, the compression component determines the appropriateness of the companding function and/or the optimum level of companding required for the specific signal input and audio playback environment, among other variables (310). This determination step 310 may occur at any practical point of process 300, such as prior to the division of the audio signal 304 or the compression of the audio signal 306. If companding is deemed to be appropriate, the gains are applied (306), and the encoder then encodes the signal for transmission to the decoder in accordance with the data format of the codec (312). Certain companding control data, such as activation data, synchronization data, companding level data, and other similar control data may be transmitted as part of the bitstream for processing by the expansion component.

[0042] FIG. 3B is a flowchart illustrating a method of expanding an audio signal in a post-decoder expansion component, under an embodiment. As shown in process 350, the decoder stage of the codec receives the bitstream encoding the audio signal from the encoder stage (352). The decoder then decodes the encoded signal in accordance with the codec data format (353). The expansion component then processes the bitstream and applies any encoded control data to switch off expansion or modify the expansion parameters based on

the control data (354). The expansion component divides the audio signal into time segments using a suitable window shape (356). In an embodiment, the time segments correspond to the same time segments used by the compression component. The expansion component then calculates the appropriate gain values for each segment in the frequency domain (358) and applies the gain values to each time segment to expand the dynamic range of the audio signal back to the original dynamic range, or any other appropriate dynamic range (360).

Companding Control

[0043] The compression and expansion components comprising the compander of system 100 may be configured to apply the pre and post-processing steps only at certain time during audio signal processing, or only for certain types of audio content. For example, companding may exhibit benefits for speech and musical transient signals. However, for other signals, such as stationary signals companding may degrade the signal quality. Thus, as shown in FIG. 3A, a companding control mechanism is provided as block 310, and control data is transmitted from the compression component 104 to the expansion component 114 to coordinate the companding operation. The simplest form of such a control mechanism is to switch off the companding function for the blocks of audio samples where application of the companding is degrading the audio quality. In an embodiment, the companding on/off decision is detected in the encoder and transmitted as bitstream element to the decoder so that compressor and expander are able to be switched on/off at the same QMF time slot.

[0044] The switching between the two states will usually lead to a discontinuity in the applied gain, resulting in audible switching artifacts or clicks. Embodiments include mechanisms to reduce or eliminate these artifacts. In a first embodiment, the system allows switching of the companding function off and on only at frames where the gain is close to 1. In this case, there is only a small discontinuity between switching the companding function on/off. In a second embodiment, a third *weak* companding mode, that is in between on and off mode is applied in an audio frame between on and off frames, and is signaled in the bitstream. The weak companding mode slowly transitions the exponent term γ from its default value during companding to 0, which is the equivalent of no companding. As an alternative to the intermediate weak companding mode, the system may implement start-frames and stop-frames that over a block of audio samples smoothly fades into an out-of-companding mode instead of abruptly switching off the companding function. In a further embodiment, the system is configured not to simply switch off the companding but rather apply an average gain. In certain cases, the audio quality of tonal-stationary signals can be

increased if a constant gain factor is applied to an audio frame that more greatly resembles the gain factors of adjacent companding-on frames than a constant gain factor of 1.0 in a companding off situation. Such a gain factor can be calculated by averaging all companding gains over one frame. A frame containing constant average companding gain is thus signaled in the bitstream.

[0045] Although embodiments are described in the context of a monophonic audio channel, it should be noted that in a straightforward extension multiple channels could be handled by repeating the approach individually on each channel. However, audio signals that comprise two or more channels present certain additional complexities that are addressed by embodiments of the companding system of FIG. 1. The companding strategy should depend on the similarity between channels.

[0046] For example, in the case of stereo-panned transient signals it has been observed that independent companding of the individual channels may result in audible image artifacts. In an embodiment, the system determines a single gain value for each time-segment from the subband samples of both channels and uses the same gain value to compress/expand the two signals. This approach is generally suitable whenever the two channels have very similar signals, wherein similarity is defined using cross correlation, for instance. A detector calculates the similarity between channels and switches between using individual companding of the channels or jointly companding the channels. Extensions to more channels would divide the channels into groups of channels using similarity criteria and apply joint companding on the groups. This grouping information can then be transmitted through the bitstream.

System Implementation

[0047] FIG. 4 is a block diagram illustrating a system for compressing an audio signal in conjunction with an encoder stage of a codec, under an embodiment. FIG. 4 illustrates a hardware circuit or system that implements at least a portion of the compression method for use in a codec-based system shown in FIG. 3A. As shown in system 400, an input audio signal 401 in the time domain is input to a QMF filterbank 402. This filterbank performs an analysis operation that separates the input signal into multiple components in which each bandpass filter carries a frequency sub-band of the original signal. Reconstruction of the signal is performed in a synthesis operation performed by QMF filterbank 410. In the example embodiment of FIG. 4, both the analysis and synthesis filterbanks handle 64 bands. The core encoder 412 receives the audio signal from the synthesis filterbank 410 and

generates a bitstream 414 by encoding the audio signal in the appropriate digital format (e.g., MP3, AAC, etc.).

[0048] System 400 includes a compressor 406 that applies gain values to each of the short segments that the audio signal has been divided into. This produces a compressed dynamic range audio signal, such as shown in FIG. 2B. A companding control unit 404 analyzes the audio signal to determine whether or how much compression should be applied based on the type of signal (e.g., speech), or the characteristics of the signal (e.g. stationary versus transient), or other relevant parameters. The control unit 404 may include a detection mechanism to detect the temporal peakness characteristic of the audio signal. Based on the detected characteristic of the audio signal and certain pre-defined criteria, the control unit 404 sends appropriate control signals to the compressor 406 to either turn off the compression function or modify the gain values applied to the short segments.

[0049] In addition to companding, many other coding tools could also operate in the QMF domain. One such tool is A-SPX (advanced spectral extension), which is shown in block 408 of FIG. 4. A-SPX is a technique that is used to allow perceptually less important frequencies to be coded with a coarser coding scheme than more important frequencies. For example, in A-SPX at the decoder end, the QMF subband samples from the lower frequency may be replicated at the higher frequencies, and the spectral envelope in the high frequency band is then shaped using side information transmitted from the encoder to the decoder.

[0050] In a system where both companding and A-SPX are performed in the QMF domain, at the encoder, the A-SPX envelope data for the higher frequencies may be extracted from the yet uncompressed subband samples as shown in FIG. 4, and compression may be applied only to the lower frequency QMF samples that correspond to the frequency range of the signal encoded by the core encoder 412. At the decoder 502 of FIG. 5, after QMF analysis 504 of the decoded signal, the expansion process 506 is applied first, and the A-SPX operation 508 subsequently reproduces the higher subband samples from the expanded signal in the lower frequencies.

[0051] In this example implementation, the QMF synthesis filterbank 410 at the encoder and the QMF analysis filterbank at the decoder 504 together introduce $640 - 64 + 1$ sample delay (~ 9 QMF slots). The core codec delay in this example is 3200 samples (50 QMF slots), so the total delay is 59 slots. This delay is accounted for by embedding control data into the bitstream and using it at the decoder, so that both the encoder compressor and the decoder expander operations are in sync.

[0052] Alternatively, at the encoder, compression may be applied on the entire bandwidth of the original signal. The A-SPX envelope data may be subsequently extracted from the compressed subband samples. In such a case, the decoder, after QMF analysis, first runs the A-SPX tool to first reconstruct the full bandwidth compressed signal. The expansion stage is then applied to recover the signal with its original dynamic range.

[0053] Yet another tool that could operate in the QMF domain may be an advanced coupling (AC) tool (not shown) in FIG. 4. In an advanced coupling system, two channels are encoded as a mono downmix with additional parametric spatial information that can be applied in the QMF domain at the decoder to reconstruct a stereo output. When AC and companding are used in conjunction with each other, the AC tool could either be placed after the compression stage 406 at the encoder, in which case it would be applied before the expansion stage 506 at the decoder. Alternatively, the AC side-information could be extracted from the uncompressed stereo signal in which case the AC tool would operate after the expansion stage 506 at the decoder. A hybrid AC mode may also be supported in which AC is used above a certain frequency and discrete stereo is used below this frequency; or alternatively, discrete stereo is used above the certain frequency and AC is used below this frequency.

[0054] As shown in FIG. 3A and 3B, the bitstream transmitted between the encoder stage and the decoder stage of the codec includes certain control data. Such control data constitutes side-information that allows the system to switch between different companding modes. The switching control data (for switching companding on/off) plus potentially some intermediate states may add on the order of 1 or 2 bits per channel. Other control data can include a signal to determine if all the channels of a discrete stereo or multichannel configuration will use common companding gain factors or if they should be calculated independently for each channel. Such data may only require a single extra bit per channel. Other similar control data elements and their appropriate bit weights may be used depending on system requirements and constraints.

Detection Mechanism

[0055] In an embodiment, a companding control mechanism is included as part of the compression component 104 to provide control of the companding in the QMF-domain. Companding control can be configured based on a number of factors, such as audio signal type. For example, in most applications, companding should be turned on for speech signals and transient signals or any other signals within the class of temporally peaky signals. The

system includes a detection mechanism to detect the peakness of a signal in order to help generate an appropriate control signal for the compander function.

[0056] In an embodiment, a measure for temporal peakness $TP(k)_{frame}$ is computed over frequency bin k for a given core codec, and is calculated using the following formula:

$$TP(k)_{frame} = \frac{\sqrt{\frac{1}{T} \sum_{t=0}^{T-1} S_t(k)^4}}{\sqrt{\frac{1}{T} \sum_{t=0}^{T-1} S_t(k)^2}},$$

[0057] In the above equation, $S_t(k)$ is the sub-band signal, and T is the number of QMF slots corresponding to one core encoder frame. In an example implementation, the value of T may be 32. The temporal peakness computed per band can be used to classify the sound content into general two categories: stationary music signals, and musical transient signals or speech signals. If the value of $TP(k)_{frame}$ is less than a defined value (e.g., 1.2), the signal in that subband of the frame is likely to be a stationary music signal. If the value of $TP(k)_{frame}$ is greater than this value, then the signal is likely to be musical transient signals or speech signals. If the value is greater than an even higher threshold value (e.g., 1.6), the signal is very likely to be a pure musical transient signal, e.g. castanets. Furthermore, it has been observed that for naturally occurring signals the values of temporal peakness obtained in different bands was more or less similar, and this characteristic could be employed to reduce the number of subbands for which temporal peakness value is to be calculated. Based on this observation, the system may implement one of the following two.

[0058] In a first embodiment, the detector executes the following process. As a first step it computes the number of bands that have a temporal peakness greater than 1.6. As a second step it then computes the mean of temporal peakness values of bands where it is less than 1.6. If the number of bands found in the first step is greater than 51, or if the mean value determined in the second step is greater than 1.45, the signal is determined to be a musical transient signal and hence companding should be switched on. Otherwise, it is determined to be a signal for which companding should not be switched on. Such a detector will switch off most of the time for speech signals. In some embodiments, speech signals will usually be coded by a separate speech coder, and so this is not generally a problem. However, in certain cases, it may be desired to switch on the companding function also for speech. In this case, a second type of detector may be preferable.

[0059] In an embodiment, this second type of detector executes the following process. As a first step, it computes the number of bands that have a temporal peakness greater than 1.2. In a second step it then computes mean of temporal peakness values of bands where it is less than 1.2. It then applies the following rule: if the result of the first step is greater than 55: turn companding on, if the result of the first step is less than 15: turn companding off; if the result of the first step lies between 15 and 55 and the result of the second step is greater than 1.16: turn companding on; and if the result of the first step lies between 15 and 55 and the result of the second step is less than 1.16: turn companding off. It should be noted that the two types of detectors described only two examples of many possible solutions for a detector algorithm, and other similar algorithms may also or alternatively be used.

[0060] The companding control function provided by element 404 of FIG. 4 may be implemented in any appropriate manner to allow companding to be used or not used based on certain operational modes. For example, companding is generally not used on the LFE (low frequency effects) channel of a surround sound system, and is also not used when there is no A-SPX (i.e., no QMF) functionality implemented. In an embodiment, the companding control function may be provided by a program executed by a circuit or processor-based elements, such as companding control element 404. Following is some example syntax of a program segment that can implement companding control, under an embodiment:

```

20   Companding_control(nCh)
    {
        sync_flag=0;
        if (nCh>1){
            sync_flag
25        }
        b_needAvg=0
        ch_count=sync_flag?1:nCh
        for (ch=0; ch<ch_count; ch++){
            b_compand_on[ch]
30            if (!b_compand_on[ch]){
                b_needAvg=1;
            }
        }
        if (b_needAvg){
35            b_compand_avg;
        }
    }

```

40 The sync_flag, b_compand_on[ch], and b_compand_avg flags or program elements may be on the order of 1-bit long, or any other length depending on system constraints and

requirements. It should be noted that the program code illustrated above is an example of one way of implementing a companding control function, and other programs or hardware components may be used to implement companding control according to some embodiments.

[0061] Although embodiments described so far include the companding process for reducing quantization noise introduced by an encoder in a codec, it should be noted that aspects of such a companding process may also be applied in signal processing systems that do not include encoder and decoder (codec) stages. Furthermore, in the event that the companding process is used in conjunction with a codec, the codec may be transform-based or non transform-based.

[0062] Aspects of the systems described herein may be implemented in an appropriate computer-based sound processing network environment for processing digital or digitized audio files. Portions of the adaptive audio system may include one or more networks that comprise any desired number of individual machines, including one or more routers (not shown) that serve to buffer and route the data transmitted among the computers. Such a network may be built on various different network protocols, and may be the Internet, a Wide Area Network (WAN), a Local Area Network (LAN), or any combination thereof.

[0063] One or more of the components, blocks, processes or other functional components may be implemented through a computer program that controls execution of a processor-based computing device of the system. It should also be noted that the various functions disclosed herein may be described using any number of combinations of hardware, firmware, and/or as data and/or instructions embodied in various machine-readable or computer-readable media, in terms of their behavioral, register transfer, logic component, and/or other characteristics. Computer-readable media in which such formatted data and/or instructions may be embodied include, but are not limited to, physical (non-transitory), non-volatile storage media in various forms, such as optical, magnetic or semiconductor storage media.

[0064] Unless the context clearly requires otherwise, throughout the description and the claims, the words “comprise,” “comprising,” and the like are to be construed in an inclusive sense as opposed to an exclusive or exhaustive sense; that is to say, in a sense of “including, but not limited to.” Words using the singular or plural number also include the plural or singular number respectively. Additionally, the words “herein,” “hereunder,” “above,” “below,” and words of similar import refer to this application as a whole and not to any particular portions of this application. When the word “or” is used in reference to a list of two or more items, that word covers all of the following interpretations of the word: any of the items in the list, all of the items in the list and any combination of the items in the list.

[0065] While one or more implementations have been described by way of example and in terms of the specific embodiments, it is to be understood that one or more implementations are not limited to the disclosed embodiments. To the contrary, it is intended to cover various modifications and similar arrangements as would be apparent to those skilled in the art.

- 5 Therefore, the scope of the appended claims should be accorded the broadest interpretation so as to encompass all such modifications and similar arrangements.

CLAIMS

1. A method of expanding an audio signal comprising:
receiving an audio signal; and
5 expanding the audio signal to an expanded dynamic range through an expansion process comprising: dividing the received audio signal into a plurality of time-segments using a defined window shape, calculating a wideband gain for each time-segment in the frequency domain using a non-energy based average of a frequency domain representation of the audio signal, and applying individual gain values to each time-segment to obtain the expanded
10 dynamic range audio signal, wherein the application of the individual gain values amplifies segments of relatively high intensity and attenuates segments of relatively low intensity.
2. The method of claim 1 wherein the segments are overlapping.
- 15 3. The method of claim 2 wherein a first filterbank is used to analyze the audio signal to obtain a frequency domain representation, and the defined window shape corresponds to a prototype filter for the first filterbank.
4. The method of claim 3 wherein the first filterbank is one of a quadrature modulated
20 filter (QMF) bank or a short-time Fourier transform.
5. The method of claim 3 wherein the wideband gain for each time segment is calculated using the subband samples in a subset of subbands in the respective time segment.
- 25 6. The method of claim 5 wherein the subset of subbands corresponds to the entire frequency range spanned by the first filterbank.
7. The method of claim 5 wherein the gain for each respective time segment is derived from the p-norm of the subband samples in each time segment where p is a positive real
30 number not equal to two.
8. The method of claim 5 wherein the wideband gain is applied in the domain of the first filterbank.

9. The method of claim 8 wherein each wideband gain value is calculated from a first subset of subbands of the first filterbank and applied to a second subset of subbands of the first filterbank, wherein the second set of subbands includes the first subset of subbands.

5

10. The method of claim 9 wherein the first and second subset of subbands are identical and correspond to a low frequency region of the audio signal.

11. The method of claim 9 wherein the first subset of subbands corresponds to a low frequency region of the audio signal and the second subset of subbands corresponds to an entire frequency range spanned by the first filterbank.

12. The method of claim 3 wherein the received audio signal has been previously compressed by a process comprising:

15 receiving an initial audio signal; and

compressing the initial audio signal to substantially reduce an original dynamic range of the initial audio signal through a compression process comprising dividing the initial audio signal into a plurality of segments using the defined window shape, calculating a wideband gain for each segment using a non-energy based average of frequency domain samples of the initial audio signal, and applying the gain values calculated from the initial audio signal to each segment of the plurality of segments to amplify segments of relatively low intensity and attenuate segments of relatively high intensity.

13. The method of claim 12 wherein the wideband gain calculated by the expansion process is substantially the inverse of the wideband gain calculated for the corresponding time-segment by the compression process.

14. The method of claim 12 wherein a second filterbank is used in the compression process to analyze the initial audio signal to obtain its frequency domain representation, and the defined window shape for segmentation is the same as the prototype filter for the filterbank, and further wherein the second filterbank is the same as the first filterbank.

30

15. The method of claim 12 wherein the received signal for the expansion process is obtained after modification of the compressed signal by an audio encoder that generates a bitstream and decoder that decodes the bitstream.

16. The method of claim 15, wherein the audio encoder and the decoder are both transform-based; and wherein the time segments of the audio signals in the compression process and the expansion process are substantially shorter than a single window length of the transform in the audio encoder and decoder.

17. The method of claim 15 further comprising:

generating control information determining an activation state of the expansion process; and

transmitting the control information in a bitstream that is transmitted from the encoder to the decoder.

18. The method of claim 17 wherein the audio signal in the bitstream is segmented into frames with each frame corresponding to a plurality of time-segments of the expansion process, and wherein the activation state is selected from the group consisting of: applying the expansion process on each time-segment in a frame; not applying the expansion for any time-segment in a frame; applying the expansion process on each time-segment in a frame with a modified gain calculation, wherein the gain applied in each of the time-segments is the average gain of all time-segments in the frame; applying the expansion process on each time-segment in a frame with a modified gain calculation resulting in a gain value that is intermediate to not applying the expansion process at all; using a stop frame to fade out from a frame where the expansion process is applied and fade in to a frame where the expansion process is not applied; using a start frame to fade out from a frame where the expansion process is not applied and fade in to a frame where the expansion process is applied; and applying the expansion process completely.

19. The method of claim 18 wherein, the control information for the expansion process is determined by the compressing step based on one or more characteristics of the initial audio signal, including at least one of a content type of the audio signal and a stationary versus transient nature of the audio signal.

20. The method of claim 19 wherein the control information is determined such that switching between activation states minimizes an occurrence of signal discontinuities.

21. The method of claim 20 wherein, the control information also controls the compression process and has the effect of turning off the compression process when the expansion process will also be switched off; turning on the compression process when the expansion process will also be switched on; and enabling a modified gain calculation for compression when a modified gain calculation for expansion will be effected; using a stop frame when a stop frame will be used at the expander; and using a start frame when a start frame will be used at the expander.

22. The method of claim 15 wherein:

the compressed audio signal and the audio signal received by the expander have a number, N, channels, where N is greater than 1;

the channels are grouped into one or more discrete subsets;

the grouping at the compressor and the expander being the same;

the channels in each group compressed sharing the same gains at the compressor and expanded sharing the same gains at the expander.

23. The method of 22 wherein the grouping is pre-determined and known at the compressor and expander.

24. The method of 23 wherein each group contains exactly one channel and there are N groups.

25. The method of 22 wherein the grouping of channels comprises:
calculating at the compressor a metric of similarity between channels;
grouping similar channels together based on the similarity metric; and
sending the grouping information through the bitstream.

26. The method of claim 22 wherein at least two channels are encoded as a mono downmix with additional parametric spatial information that is applied in the first filterbank domain to reconstruct a stereo output, and wherein the parametric spatial information is used either above a defined frequency with discrete stereo information used below the defined

frequency, or parametric spatial information is used below the defined frequency with discrete stereo information used above the defined frequency.

27. A method of compressing an audio signal comprising:

receiving an initial audio signal; and

compressing the initial audio signal to substantially reduce an original dynamic range of the initial audio signal through a compression process comprising dividing the initial audio signal into a plurality of segments using a defined window shape, calculating a wideband gain in the frequency domain using a non-energy based average of frequency domain samples of the initial audio signal, and applying individual gain value to each segment of the plurality of segments to amplify segments of relatively low intensity and attenuate segments of relatively high intensity.

28. The method of claim 27 wherein the segments are overlapping and wherein a first

filterbank is used to analyze the audio signal to obtain a frequency domain representation and the defined window shape corresponds to a prototype filter for the first filterbank.

29. The method of claim 28 wherein the first filterbank is one of a quadrature modulated filter (QMF) bank or a short-time Fourier transform.

30. The method of claim 28 wherein each individual gain value is calculated using the subband samples in a subset of subbands in a respective time segment.

31. The method of claim 30 wherein the subset of subbands corresponds to the entire frequency range spanned by the first filterbank, and wherein the gain is applied in the domain of the first filterbank.

32. The method of claim 30 wherein the gain for each time segment is derived from the p -norm of the subband samples in each time segment where p is a positive real number not equal to two.

33. The method of claim 30 wherein the gain is calculated from a first subset of subbands of the first filterbank and applied to a second subset of subbands of the first filterbank, wherein the second set of subbands includes the first subset of subbands.

34. The method of claim 33 wherein the first and second subset of subbands are identical and correspond to a low frequency region of the audio signal.

5 35. The method of claim 33 wherein the first subset of subbands corresponds to a low frequency region of the audio signal and the second subset of subbands corresponds to an entire frequency range spanned by the first filterbank.

10 36. The method of claim 27 further comprising: transmitting a compressed version of the initial audio signal to a expansion component performing an expansion process comprising receiving the compressed version of audio signal; and expanding the compressed version of the audio signal to substantially restore it to an original dynamic range of the initial audio signal through process comprising dividing the initial audio signal into a plurality of time segments using the defined window shape, calculating a wideband gain in the frequency
15 domain using a non-energy based average of frequency domain samples of the initial audio signal; and applying individual gain values of the wideband gain to each time segment to amplify segments of relatively high intensity and attenuate segments of relatively low intensity.

20 37. The method of claim 36 wherein the gain calculated by the compressing step is substantially the inverse of the gain calculated by the expanding step for the same time-segment.

25 38. The method of claim 36 wherein a second filterbank is used in the expansion process to analyze the initial audio signal to obtain its frequency domain representation, and the defined window shape for segmentation is the same as the prototype filter for the filterbank, and further wherein the second filterbank is the same as the first filterbank.

30 39. The method of claim 36 wherein the received signal for the expanding step is obtained after modification of the compressed signal by an audio encoder that generates a bitstream and decoder that decodes the bitstream.

40. The method of claim 39, wherein the audio encoder and the decoder are both transform-based; and wherein the time segments of the audio signals in the compressing and

expanding steps are substantially shorter than a single window length of the transform in the audio encoder and decoder.

41. The method of claim 39 further comprising:

5 generating control information determining an activation state of the expanding step;
and

transmitting the control information in a bitstream that is transmitted from the encoder to the decoder.

10 42. The method of claim 41 wherein the audio signal in the bitstream is segmented into frames with each frame corresponding to a plurality of time-segments of the expansion process, and wherein the activation state is selected from the group consisting of: applying the expansion process on each time-segment in a frame; not applying the expansion for any time-segment in a frame; applying the expansion process on each time-segment in a frame
15 with a modified gain calculation, wherein the gain applied in each of the time-segments is the average gain of all time-segments in the frame; applying the expansion process on each time-segment in a frame with a modified gain calculation resulting in a gain value that is intermediate to not applying the expansion process at all; using a stop frame to fade out from a frame where the expansion process is applied and fade in to a frame where the expansion
20 process is not applied; using a start frame to fade out from a frame where the expansion process is not applied and fade in to a frame where the expansion process is applied; and applying the expansion process completely.

43. The method of claim 42 wherein, the control information for the expansion process is
25 determined by the compressing step based on one or more characteristics of the initial audio signal, including at least one of a content type of the audio signal and a stationary versus transient nature of the audio signal.

44. The method of claim 43 wherein the control information is determined such that
30 switching between activation states minimizes an occurrence of signal discontinuities.

45. The method of claim 44 wherein, the control information also controls the compression process and has the effect of turning off the compressing step when the expanding step will also be switched off; turning on the compressing step when the

expanding step will also be switched on; and enabling a modified gain calculation for compression when a modified gain calculation for expansion will be effected; using a stop frame when a stop frame will be used at the expander; and using a start frame when a start frame will be used at the expander.

5

46. The method of claim 39 wherein:

the compressed audio signal and the audio signal received by the expander have a number, N, channels, where N is greater than 1;

the channels are grouped into one or more discrete subsets;

10

the grouping at the compressor and the expander being the same;

the channels in each group compressed sharing the same gains at the compressor and expanded sharing the same gains at the expander.

47. The method of 46 wherein the grouping is pre-determined and known at the

15

compressor and expander.

48. The method of 47 wherein each group contains exactly one channel and there are N groups.

20

49. The method of 46 wherein the grouping of channels comprises:

calculating at the compressor a metric of similarity between channels;

grouping similar channels together based on the similarity metric; and

sending the grouping information through the bitstream.

25

50. The method of claim 49 wherein at least two channels are encoded as a mono downmix with additional parametric spatial information that is applied in the first filterbank domain to reconstruct a stereo output, and wherein the parametric spatial information is used either above a defined frequency with discrete stereo information used below the defined frequency, or parametric spatial information is used below the defined frequency with discrete stereo information used above the defined frequency.

30

51. An apparatus for compressing an audio signal comprising:

a first interface receiving an initial audio signal; and

a compressor compressing the initial audio signal to substantially reduce an original dynamic range of the initial audio signal by dividing the initial audio signal into a plurality of segments using a defined window shape, calculating a wideband gain in the frequency domain using a non-energy based average of frequency domain samples of the initial audio signal, and applying individual gain values to each segment of the plurality of segments to amplify segments of relatively low intensity and attenuate segments of relatively high intensity.

52. The apparatus of claim 51 further comprising a first filterbank analyzing the audio signal to obtain a frequency domain representation and wherein the defined window shape corresponds to a prototype filter for the first filterbank, and further wherein the first filterbank is one of a quadrature modulated filter (QMF) bank or a short-time Fourier transform.

53. The apparatus of claim 52 wherein the individual gain values are calculated using the subband samples in a subset of subbands in each respective time segment.

54. The apparatus of claim 53 wherein the subset of subbands corresponds to the entire frequency range spanned by the first filterbank, and wherein the gain is applied in the domain of the first filterbank.

55. The apparatus of claim 52 further comprising a second interface transmitting a compressed version of the initial audio signal to an expander that receives the compressed version of audio signal, and expands the compressed version of the audio signal to substantially restore it to an original dynamic range of the initial audio signal by dividing the initial audio signal into a plurality of segments using the defined window shape, calculating a wideband gain in the frequency domain using a non-energy based average of frequency domain samples of the initial audio signal; and applying a respective gain value to each segment of the plurality of segments to amplify segments of relatively high intensity and attenuate segments of relatively low intensity.

56. The apparatus of claim 55 wherein the gain calculated by the expander is substantially the inverse of the gain calculated by the compressor for the same time segment.

57. The apparatus of claim 55 further comprising a second filterbank analyzing the initial audio signal to obtain its frequency domain representation, and wherein the defined window shape for segmentation is the same as the prototype filter for the filterbank, and further wherein the second filterbank is the same as the first filterbank.

5

58. The apparatus of claim 55 further comprising encoder and decoder stages of an audio codec configured to transmit a bitstream of the compressed version of the audio signal from the compressor to the expander, and wherein the encoder and decoder are both transform-based.

10

59. The apparatus of claim 58 further comprising a control component generating control information determining an activation state of the expander and transmitting the control information in the bitstream, and wherein, the control information for the expansion process is determined by the compressor based on one or more characteristics of the initial audio signal, including at least one of a content type of the audio signal and a stationary versus transient nature of the audio signal.

15

60. The apparatus of claim 55 further comprising a parametric spatial information component applying parametric spatial information in the first filterbank domain to reconstruct a stereo output, and wherein the parametric spatial information is used either above a defined frequency with discrete stereo information used below the defined frequency, or parametric spatial information is used below the defined frequency with discrete stereo information used above the defined frequency.

20

61. An apparatus for expanding an audio signal comprising:
a first interface receiving a compressed audio signal; and
an expander expanding the compressed audio signal to substantially restore its original uncompressed dynamic range by dividing the initial audio signal into a plurality of segments using a defined window shape, calculating a wideband gain in the frequency domain using a non-energy based average of frequency domain samples of the initial audio signal, and applying individual gain values to each segment of the plurality of segments to amplify segments of relatively high intensity and attenuate segments of relatively low intensity.

25

30

62. The apparatus of claim 61 further comprising a first filterbank analyzing the audio signal to obtain a frequency domain representation and wherein the defined window shape corresponds to a prototype filter for the first filterbank, and further wherein the first filterbank is one of a quadrature modulated filter (QMF) bank or a short-time Fourier transform.

5

63. The apparatus of claim 62 wherein the wideband gain comprises an individual gain value for each time segment, and wherein each individual gain value is calculated using the subband samples in a subset of subbands in each respective time segment.

10 64. The apparatus of claim 63 wherein the subset of subbands corresponds to the entire frequency range spanned by the first filterbank, and wherein the gain is applied in the domain of the first filterbank.

15 65. The apparatus of claim 62 further comprising a second interface receiving the compressed audio signal from a compressor that receives an initial audio signal, and compresses the initial audio signal to substantially reduce the original dynamic range of the initial audio signal by dividing the initial audio signal into a plurality of segments using the defined window shape, calculating a wideband gain in the frequency domain using a non-energy based average of frequency domain samples of the initial audio signal; and applying a
20 respective gain value to each segment of the plurality of segments to amplify segments of relatively low intensity and attenuate segments of relatively high intensity.

25 66. The apparatus of claim 65 wherein the gain calculated by the compressor is substantially the inverse of the gain calculated by the expander for the same time segment.

30 67. The apparatus of claim 65 further comprising a second filterbank analyzing the initial audio signal to obtain its frequency domain representation, and wherein the defined window shape for segmentation is the same as the prototype filter for the filterbank, and further wherein the second filterbank is the same as the first filterbank.

68. The apparatus of claim 65 further comprising encoder and decoder stages of an audio codec configured to transmit a bitstream of the compressed version of the audio signal from the compressor to the expander, and wherein the encoder and decoder are both transform-based.

69. The apparatus of claim 68 further comprising a control component generating control information determining an activation state of the expander and transmitting the control information in the bitstream, and wherein, the control information for the expansion process
5 is determined by the compressor based on one or more characteristics of the initial audio signal, including at least one of a content type of the audio signal and a stationary versus transient nature of the audio signal.

70. The apparatus of claim 65 further comprising a parametric spatial information
10 component applying parametric spatial information in the first filterbank domain to reconstruct a stereo output, and wherein the parametric spatial information is used either above a defined frequency with discrete stereo information used below the defined frequency, or parametric spatial information is used below the defined frequency with discrete stereo information used above the defined frequency.

15