

Apparatus and Method for Processing an Audio Signal Using a Harmonic Post-Filter

5

Specification

The present invention is related to audio processing and, particularly, to audio processing using a harmonic post filter.

10

Transform-based audio codecs generally introduce inter-harmonic noise when processing harmonic audio signals, particularly at low bitrates.

15

This effect is further worsen when the transform-based audio codec operates at low delay, due to the worse frequency resolution and/or selectivity introduced by a shorter transform size and/or a worse window frequency response.

20

This inter-harmonic noise is generally perceived as a very annoying artifact, significantly reducing the performance of the transform-based audio codec when subjectively evaluated on highly tonal audio material.

25

Several solutions exist to improve the subjective quality of transform-based audio codecs on harmonics audio signals. All of them are based on prediction-based techniques, either in the transform-domain or in the time-domain.

30

Examples of transform-domain approaches are:

- [1] H. Fuchs, "Improving MPEG Audio Coding by Backward Adaptive Linear Stereo Prediction", 99th AES Convention, New York 1995, Preprint 4086.
- [2] L. Yin, M. Suonio, M. Väänänen, "A New Backward Predictor for MPEG Audio Coding", 103rd AES Convention, New York 1997, Preprint 4521
- [3] Juha Ojanperä, Mauri Väänänen, Lin Yin, "Long Term Predictor for Transform Domain Perceptual Audio Coding", 107th AES Convention, New York 1999, Preprint 5036.

35

Examples of time-domain approaches are:

- [4] Philip J. Wilson, Harprit Chhatwal, "Adaptive transform coder having long term predictor", U.S. Patent 5,012,517, April 30, 1991.
- [5] Jeongook Song, Chang-Heon Lee, Hyen-O Oh, Hong-Goo Kang, "Harmonic Enhancement in Low Bitrate Audio Coding Using and Efficient Long-Term Predictor", EURASIP Journal on Advances in Signal Processing 2010.
- [6] Juin-Hwey Chen, "Pitch-based pre-filtering and post-filtering for compression of audio signals", U.S. Patent 8,738,385, May 27, 2014.

It is an object of the present invention to provide an improved concept for processing an audio signal.

This object is achieved by an apparatus for processing an audio signal of claim 1, a method for processing an audio signal of claim 12, a system of claim 13, a method for operating a system of claim 17 or a computer program of claim 18.

The present invention is based on the finding that the subjective quality of an audio signal can be substantially improved by using a harmonic post-filter having a transfer function comprising a numerator and a denominator. The numerator of the transfer function comprises a gain value indicated by a transmitted gain information and the denominator comprises an integer part of a pitch lag indicated by a pitch lag information and a multi-tap filter depending on a fractional part of the pitch lag.

Hence, it is possible to remove inter-harmonic noise introduced by a typical domain-changing audio decoder as an artifact. This harmonic post-filter is particularly useful in that it relies on transmitted information, i.e., the pitch gain and the pitch lag which are available anyway in a decoder, since this information is received from a corresponding encoder via a decoder input signal. Furthermore, the post-filtering is of specific accuracy due to the fact that not only the integer part of the pitch lag is accounted for, but, in addition, the fractional part of the pitch lag is accounted for. The fractional part of the pitch lag can be particularly introduced into the post-filter via a multi-tap filter which has filter coefficients actually depending on the fractional part of the pitch lag. This filter can be implemented as an FIR filter or can also be implemented as any other filter such as an IIR filter or a different filter implementation. Any domain change such as a time to frequency change or an LPC to time change or a time to LPC change or a frequency to time change can be advantageously improved by the post-filter concept of the invention. Preferably, however, the domain change is a frequency to time domain change.

Hence, embodiments of the present invention reduces inter-harmonic noise introduced by a transform audio codec based on a long-term predictor working in the time domain. Contrary to [04] - [6], where both pre-filter before the transform coding and a post-filter after the transform decoding are used, the present invention preferably applies a post-filter only.

Furthermore, it has been noticed that the pre-filter employed in [04] - [6] has the tendency to introduce instabilities in the input signal given to the transform encoder. These instabilities are due to changes in gain and/or pitch lag from frame to frame. The transform coder has difficulties in encoding such instabilities, particularly at low bitrates, and one will sometimes introduce even more noise in the decoded signal compared to a situation without any pre- or post-filter.

Preferably, the present invention does not employ any pre-filter at all and, therefore, completely avoids the problems involved with a pre-filter.

Furthermore, the present invention relies on a post-filter that is applied on the decoded signal after transform coding. This post-filter is based on a long-term prediction filter accounting for the integer part and the fractional part of the pitch lag that reduces the inter-harmonic noise introduced by the transform audio codec.

For better robustness, the post-filter parameters pitch lag and pitch gain are estimated at the encoder-side and transmitted in the bitstream. However, in other implementations, the pitch lag and pitch gain can also be estimated on the decoder-side based on the decoded audio signal obtained by an audio decoder comprising a frequency-time converter for converting a frequency-representation of the audio signal into a time-domain representation of the audio signal.

In a preferred embodiment, the numerator additionally comprises a multi-tap filter for a zero fractional part of the pitch lag in order to compensate for a spectral tilt introduced by the multi-tap filter in the denominator, which depends on the fractional part of the pitch lag.

Preferably, the post-filter is configured to suppress an amount of energy between harmonics in a frame, wherein the amount of energy suppressed is smaller than 20% of a total energy of the time-domain representation in the frame.

In a further embodiment, the denominator comprises a product between the multi-tap filter and the gain value.

- 5 In a further embodiment, the filter numerator further comprises a product of a first scalar value and a second scalar value, wherein the denominator only comprises the second scalar value rather than the first scalar value. These scalar values are set to predetermined values and have values greater than 0 and lower than 1; and, additionally, the second scalar value is lower than the first scalar value. Hence, it is possible in a very efficient way to set the energy removal characteristics which are typically unwanted and to additionally set the filter strength, i.e., how strong the filter attenuates inter-harmonic artifacts in a transform-domain decoder output signal.
- 10

- The apparatus further comprises, in a preferred embodiment, a filter controller for setting at least the second scalar value depending on a bitrate so that a higher value is set for a lower bitrate and vice versa.
- 15

- Furthermore, the filter controller is configured for selecting, depending on the fractional part of the pitch lag, the corresponding multi-tap filter in a signal-dependent way in order to set the harmonic post-filter signal-adaptively, i.e., dependent on the actually provided fractional part value of the pitch lag.
- 20

- Subsequently, preferred embodiments of the present invention are discussed in the context of the accompanying drawings, in which:
- 25

- Fig. 1 illustrates an embodiment of an inventive apparatus for processing an audio signal;
- Fig 2 illustrates a preferred implementation of the harmonic post-filter represented as transfer functions in the z domain;
- 30
- Fig. 3 illustrates a further preferred embodiment for the harmonic post-filter represented by a transfer function in the z domain;

- Fig. 4 illustrates a preferred implementation of an encoder for generating an encoded signal to be decoded by a transform-domain audio decoder illustrated in Fig. 1;
- 5 Fig. 5 illustrates a preferred implementation of the multi-tap filter as an FIR filter controlled by a filter controller;
- Fig. 6 illustrates a cooperation between the filter controller and a memory having pre-stored tap weights depending on the fractional part;
- 10 Fig. 7a illustrates a frequency response of a filter having a zero α value.
- Fig. 7b illustrates a frequency response of a preferred harmonic post-filter having an α value equal to 1;
- 15 Fig. 7c illustrates a frequency response of a preferred harmonic post-filter having an α value of 0.8;
- Fig. 8a illustrates a preferred embodiment of a harmonic post-filter having a β value equal to 0.4; and
- 20 Fig. 8b illustrates a frequency response of a harmonic post-filter having a β value of 0.2.
- 25 Fig. 1 illustrates an apparatus for processing an audio signal having associated therewith a pitch lag information and a gain information. This gain information can be transmitted to a decoder 100 via a decoder input 102 receiving an encoded signal or, alternatively, this information can be calculated in the decoder itself, when this information is not available. However, for a more robust operation, it is preferred to calculate the pitch lag information and the pitch gain information on the encoder-side.
- 30

The decoder 100 comprises e.g. a frequency-time converter for converting a frequency-time representation of the audio signal into a time-domain representation of the audio signal. Thus, the decoder is not a pure time-domain speech codec, but comprises a pure transform domain decoder or a mixed transform domain decoder or any other coder

35

operating in a domain different from a time domain. Furthermore, it is preferred that the second domain is the time domain.

The apparatus furthermore comprises a harmonic post-filter 104 for filtering the time-domain representation of the audio signal, and this harmonic post-filter is based on a transfer function comprising a numerator and a denominator. Particularly, the numerator comprises a gain value indicated by the gain information and the denominator comprises an integer part of a pitch lag indicated by the pitch lag information and, importantly, further comprises a multi-tap filter depending on a fractional part of the pitch lag.

A preferred implementation of this harmonic post filter with a transfer function $H(z)$ is illustrated in Fig. 2. This filter receives the decoder output signal 106 and subjects this decoded output signal to a post-filtering operation to obtain a post-filtered output signal 108. This post-filtered output signal can be output as the processed signal or can be further processed by any procedure for removing any discontinuities introduced by the post-filtering operation which, of course, is signal-dependent, i.e., can vary from frame to frame. This discontinuity removal operation can be any of the well-known discontinuity removal operation such as cross-fading, which means that an earlier frame is faded out and, at the same time, a new frame is faded in and, preferably, the fading characteristic is so that the fading factors add up to one throughout the cross-fading operation. However, other discontinuity removal such as low-pass filtering or LPC filtering can be applied as well.

The apparatus for processing an audio signal illustrated in Fig. 1 furthermore comprises a multi-tap filter information storage 112 and a filter controller 114. Particularly, the filter controller 114 receives side information 116 from the decoder 100, and this side information can, for example, be the pitch gain information g and the pitch lag information, i.e., information on the integer part T_{int} of the pitch lag and the fractional part T_{fr} of the pitch lag. This information is useful for setting the harmonic post-filter from frame to frame and, additionally, for selecting a multi-tap filter information $B(z, T_{\text{fr}})$. Furthermore, additional information such as the bitrate applied by the decoder or the sampling rate underlying the decoded signal can also be used by the filter control 114 in order to particularly set the scalar values α , β for a certain encoder and/or decoder setting with respect to bitrate and sampling rate.

Fig. 2 illustrates a pole/zero representation of a filter transfer function $H(z)$ in the z domain as known in the art. Naturally, there are numerous other representations of the harmonic post-filter, which are all filter representations, which can be converted to the kind of pole/zero representation in the z domain. Hence, the present invention is applicable for each filter which is describable in any way by such a transfer function as illustrated in the specification.

Fig. 3 illustrates a preferred embodiment of the harmonic post-filter again described as a transfer function in the pole/zero notation in the z domain.

The filter can be described as follows:

$$H(z) = \frac{1 - \alpha\beta g B(z, 0)}{1 - \beta g B(z, T_{fr}) z^{-T_{int}}}$$

with g the decoded gain, T_{int} and T_{fr} the integer and fractional part of the decoded pitch lag, α and β two scalars that weight the gain, and $B(z, T_{fr})$ a low-pass FIR filter whose coefficients depends on the fractional part of the decoded pitch lag.

Note that $B(z, 0)$ in the numerator of $H(z)$ is used to compensate for the tilt introduced by $B(z, T_{fr})$.

β is used to control the strength of the post-filter. A β equals to 1 produces full effects, suppressing the maximum possible amount of energy between the harmonics. A β equals to 0 disables the post-filter. Generally, a quite low value is used to not suppress too much energy between the harmonics. The value can also depend on the bitrate with a higher value at a lower bitrate, e.g. 0.4 at low bitrate and 0.2 at a high bitrate.

α is used to add a slight tilt to the frequency response of $H(z)$, in order to compensate for the slight loss in energy in the low frequencies. The value of α is generally chosen close to 1, e.g. 0.8.

An example of $B(z, T_{fr})$ is given in Fig. 6. The order and the coefficients of $B(z, T_{fr})$ can also depend on the bitrate and the output sampling rate. A different frequency response can be designed and tuned for each combination of bitrate and output sampling rate.

Particularly, it has been found out that even values for α between 0.6 and lower than 1.0 are useful and that, additionally, values for β between 0.1 and 0.5 have been proved to be useful as well.

5 Furthermore, the multi-tap filter can have a variable number of taps. It has been found that for certain implementations, four taps are sufficient, where one tap is z^{+1} . However, smaller filters with only two taps or even larger filters with more than four taps are useful for certain implementations.

10 Fig. 6 illustrates a preferred implementation of filters $B(z)$ for different fractional values of the pitch lag and, particularly, for a pitch lag resolution of $\frac{1}{4}$. For this implementation, four different filter descriptions for the multi-tap filter in the denominator of the transfer function of the harmonic post-filter are illustrated. However, it has been found that the filter coefficients do not necessarily have to indicate exactly the illustrated values in Fig. 6, but
15 certain variations of ± 0.05 can be useful in other implementations as well.

Particularly, as illustrated in Fig. 1, the tap weights illustrated in Fig. 6 are stored within the memory 112 for the multi-tap filter information. The filter controller 114 receives the fractional part T_{fr} from line 116 of Fig. 1 and, in response to this value, addresses the
20 memory 112 in order to retrieve, via a retrieval line 200 the specific filter information for the specific fractional part of the pitch lag. This information is then forwarded via an output line 202 to the harmonic post-filter 104 so that the harmonic post-filter is correctly set. A certain implementation of the multi-tap FIR filter is illustrated in Fig. 5. The weight indication w_1 to w_4 corresponds to the notation in Fig. 6 and the filter controller 114
25 applies, in response to the actual fractional part of the pitch lag the corresponding weights for a certain audio frame. The other portions such as delay portions 501, 502, 503 and the combiner 505 can be implemented as illustrated. In this context, it is emphasized that the delay value 501 is, in the z notation a negative delay value, since it has been found out that an FIR filter representation having a negative delay value in addition to a positive
30 delay value such as 503 and 504 is particularly useful.

Subsequently, a preferred encoder implementation having certain functional blocks and operating without any pre-filter is illustrated in Fig. 4. The filter portion illustrated in Fig. 4 comprises a pitch estimator 402, a pitch refiner 404, a fractional part estimator 406, a
35 transient detector 408, a gain estimator 410 and a gain quantizer 412. The information provided by the gain quantizer 412, the fractional part estimator 406, the pitch refiner 404

and the decision bit generated by the transient detector 408 are input into an encoded signal former 414. The encoded signal former provides an encoded signal 102, which is then input into the decoder 100 illustrated in Fig. 1. The encoded signal 102 will comprise additional signal information not illustrated in Fig. 4.

5

Subsequently, the functionality of the pitch estimator 402 is described.

One pitch lag (integer part + fractional part) per frame is estimated (frame size e.g. 20ms). This is done in 3 steps to reduce complexity and improves estimation accuracy.

10

A pitch analysis algorithm that produces a smooth pitch evolution contour is used (e.g. Open-loop pitch analysis described in Rec. ITU-T G.718, sec. 6.6). This analysis is generally done on a subframe basis (subframe size e.g. 10ms), and produces one pitch lag estimate per subframe. Note that these pitch lag estimates do not have any fractional part and are generally estimated on a downsampled signal (sampling rate e.g. 6400Hz). The signal used can be any audio signal, e.g. a LPC weighted audio signal as described in Rec. ITU-T G.718, sec. 6.5.

15

The pitch refiner operates as follows:

20

The final integer part of the pitch lag is estimated on an audio signal $x[n]$ running at the core encoder sampling rate, which is generally higher than the sampling rate of the downsampled signal used in a. (e.g. 12.8kHz, 16kHz, 32kHz...). The signal $x[n]$ can be any audio signal e.g. an LPC weighted audio signal.

25

The integer part of the pitch lag is then the lag d_m that maximizes the autocorrelation function

$$C(d) = \sum_{n=0}^N x[n]x[n-d]$$

with d around a pitch lag T estimated in step 1.a.

30

$$T - \delta_1 \leq d \leq T + \delta_2$$

The fractional part estimator 406 operates as follows:

35

The fractional part is found by interpolating the autocorrelation function $C(d)$ computed in step 2.b. and selecting the fractional pitch lag which maximizes the interpolated

autocorrelation function. The interpolation can be performed using a low-pass FIR filter as described in e.g. Rec. ITU-T G.718, sec. 6.6.7.

The transient detector 408 illustrated in Fig. 4 is configured for generating a decision bit.

5

If the input audio signal does not contain any harmonic content, then no parameters are encoded in the bitstream. Only 1 bit is sent such that the decoder knows whether he has to decode the post-filter parameters or not. The decision is made based on several parameters:

10

a. Normalized correlation at the integer pitch lag estimated in step 1.b.

$$\text{norm. corr.} = \frac{\sum_{n=0}^N x[n]x[n-d_m]}{\sqrt{\sum_{n=0}^N x[n]x[n]} \sqrt{\sum_{n=0}^N x[n-d_m]x[n-d_m]}}$$

15 The normalized correlation is 1 if the input signal is perfectly predictable by the integer pitch lag, and 0 if it is not predictable at all. A high value (close to 1) would then indicate a harmonic signal. For a more robust decision, the normalized correlation of the past frame can also be used in the decision., e.g.:

20 If (norm.corr(curr.)*norm.corr(prev.)) > 0.25, then the current frame contains some harmonic content (bit=1)

b. Features computed by a transient detector (e.g. Temporal flatness measure, Maximal energy change), to avoid activating the post-filter on a signal containing a transient. e.g.

25 If (tempFlatness>3.5 or maxEnergychange>3.5) then set bit=0 and do not send any parameters

Furthermore, the gain estimator 410 calculates a gain to be input into the gain quantizer 412

30

The gain is generally estimated on the input audio signal at the core encoder sampling rate, but it can also be any audio signal like the LPC weighted audio signal. This signal is noted $y[n]$ and can be the same or different than $x[n]$.

35 The prediction $y_P[n]$ of $y[n]$ is first found by filtering $y[n]$ with the following filter

$$P(z) = B(z, T_{fr})z^{-T_{int}}$$

with T_{int} the integer part of the pitch lag (estimated in 1.b.) and $B(z, T_{fr})$ a low-pass FIR filter whose coefficients depend on the fractional part of the pitch lag T_{fr} (estimated
5 in 1.c.).

One example of $B(z)$ when the pitch lag resolution is $\frac{1}{4}$:

$T_{fr} = \frac{0}{4}$	$B(z) = 0.0000z^{-2} + 0.2325z^{-1} + 0.5349z^0 + 0.2325z^1$
$T_{fr} = \frac{1}{4}$	$B(z) = 0.0152z^{-2} + 0.3400z^{-1} + 0.5094z^0 + 0.1353z^1$
$T_{fr} = \frac{2}{4}$	$B(z) = 0.0609z^{-2} + 0.4391z^{-1} + 0.4391z^0 + 0.0609z^1$
$T_{fr} = \frac{3}{4}$	$B(z) = 0.1353z^{-2} + 0.5094z^{-1} + 0.3400z^0 + 0.0152z^1$

The gain g is then computed as follows:

$$g = \frac{\sum_{n=0}^{N-1} y[n]y_p[n]}{\sum_{n=0}^{N-1} y_p[n]y_p[n]}$$

and limited between 0 and 1.

Finally, the gain is quantized e.g. on 2 bits, using e.g. uniform quantization.

If the gain is quantized to 0, then no parameters are encoded in the bitstream, only the
20 one decision bit (bit=0).

As outlined before, the post-filter is applied on the output audio signal after the transform decoder. It processes the signal on the frame-by-frame basis, with the same frame size as used in the encoder-side such as 20ms. As illustrated, it is based on a long-term prediction
25 filter $H(z)$ whose parameters are determined from the parameters estimated at the encoder-side and decoded from the bitstream. This information comprises the decision bit, the pitch lag and the gain. If the decision bit is 0, then the pitch lag and the gain are not decoded and are assumed to be 0 not written at all into the bitstream.

30 As discussed, if the filter parameters are different from one frame to the next frame, a discontinuity can be introduced at the border between the two frames. To avoid discontinuity, a discontinuity remover is applied such as a cross-fader or any other implementation for that purpose.

Furthermore, several different ways to set the harmonic post-filter are illustrated in Fig. 7a to 8b. The plots illustrate the frequency domain transfer function. The horizontal axis is related to the normalized frequency 1 and the vertical axis is the magnitude of the filter response in dB. It is emphasized that in all illustrations but Fig. 7b, the filter introduces an amplification for low frequencies, i.e., a certain positive dB magnitude value.

Particularly, Fig. 7a illustrates a transfer function, implementing the filter in Fig. 3, with the certain parameter values as indicated above. Furthermore, the α value, i.e., the first scalar value is set to 0. Fig. 7b illustrates a similar situation, but now with an α value equal to 1.

The other parameters are identical to Fig. 7a.

Fig. 7c illustrates a further implementation where α is equal to 0.8 which has a slight tilt and a boosting of the lower frequencies. Again, Fig. 7 has the same other parameters as indicated in Fig. 7a. It becomes clear that α equal to 1 removes the tilt and all harmonic frequencies have a gain of 1. The drawback of this setting is a loss of energy at the frequencies between the harmonics. Therefore, a value of α equal to 0.8 as in Fig. 7c is preferred. This value adds a slight tilt compared to the α equal to 1 situation of Fig. 7b. In order to compensate the loss of energies at the frequencies between the harmonics, this slight tilt is preferably used.

Furthermore, Fig. 8a and 8b illustrate filter settings for a value of α equal to 0.8 and different β -values, i.e., a β -value of 0.4 in Fig. 8a and a β -value of 0.2 in Fig. 8b. It becomes clear that a β -value of 0.4 has a stronger post-filtering effect compared to a β -value of 0.2 and, therefore, a β -value of 0.4 is used at lower bitrates in order to remove inter-harmonic noise introduced by such a low bitrate.

On the other hand, β equal 0.2 has a less strong effect for suppressing energy between the harmonics and, therefore, this β -value is preferred for high bitrates due to the fact that at such higher bitrates, not so much inter-harmonic noise exists.

Although some aspects have been described in the context of an apparatus, it is clear that these aspects also represent a description of the corresponding method, where a block or device corresponds to a method step or a feature of a method step. Analogously, aspects described in the context of a method step also represent a description of a corresponding block or item or feature of a corresponding apparatus. Some or all of the method steps may be executed by (or using) a hardware apparatus, like for example, a microprocessor,

a programmable computer or an electronic circuit. In some embodiments, some one or more of the most important method steps may be executed by such an apparatus.

5 The inventive transmitted or encoded signal can be stored on a digital storage medium or can be transmitted on a transmission medium such as a wireless transmission medium or a wired transmission medium such as the Internet.

10 Depending on certain implementation requirements, embodiments of the invention can be implemented in hardware or in software. The implementation can be performed using a digital storage medium, for example a floppy disc, a DVD, a Blu-Ray, a CD, a ROM, a PROM, and EPROM, an EEPROM or a FLASH memory, having electronically readable control signals stored thereon, which cooperate (or are capable of cooperating) with a programmable computer system such that the respective method is performed. Therefore, the digital storage medium may be computer readable.

15 Some embodiments according to the invention comprise a data carrier having electronically readable control signals, which are capable of cooperating with a programmable computer system, such that one of the methods described herein is performed.

20 Generally, embodiments of the present invention can be implemented as a computer program product with a program code, the program code being operative for performing one of the methods when the computer program product runs on a computer. The program code may, for example, be stored on a machine readable carrier.

25 Other embodiments comprise the computer program for performing one of the methods described herein, stored on a machine readable carrier.

30 In other words, an embodiment of the inventive method is, therefore, a computer program having a program code for performing one of the methods described herein, when the computer program runs on a computer.

35 A further embodiment of the inventive method is, therefore, a data carrier (or a non-transitory storage medium such as a digital storage medium, or a computer-readable medium) comprising, recorded thereon, the computer program for performing one of the

methods described herein. The data carrier, the digital storage medium or the recorded medium are typically tangible and/or non-transitory.

5 A further embodiment of the invention method is, therefore, a data stream or a sequence of signals representing the computer program for performing one of the methods described herein. The data stream or the sequence of signals may, for example, be configured to be transferred via a data communication connection, for example, via the internet.

10 A further embodiment comprises a processing means, for example, a computer or a programmable logic device, configured to, or adapted to, perform one of the methods described herein.

15 A further embodiment comprises a computer having installed thereon the computer program for performing one of the methods described herein.

20 A further embodiment according to the invention comprises an apparatus or a system configured to transfer (for example, electronically or optically) a computer program for performing one of the methods described herein to a receiver. The receiver may, for example, be a computer, a mobile device, a memory device or the like. The apparatus or system may, for example, comprise a file server for transferring the computer program to the receiver .

25 In some embodiments, a programmable logic device (for example, a field programmable gate array) may be used to perform some or all of the functionalities of the methods described herein. In some embodiments, a field programmable gate array may cooperate with a microprocessor in order to perform one of the methods described herein. Generally, the methods are preferably performed by any hardware apparatus.

30 The above described embodiments are merely illustrative for the principles of the present invention. It is understood that modifications and variations of the arrangements and the details described herein will be apparent to others skilled in the art. It is the intent, therefore, to be limited only by the scope of the impending patent claims and not by the specific details presented by way of description and explanation of the embodiments
35 herein.

Claims

1. Apparatus for processing an audio signal having associated therewith a pitch lag information and a gain information, comprising:

5

a domain converter (100) for converting a first domain representation of the audio signal into a second domain representation of the audio signal; and

10

a harmonic post-filter (104) for filtering the second domain representation of the audio signal, wherein the post-filter is based on a transfer function comprising a numerator and a denominator, wherein the numerator comprises a gain value indicated by the gain information, and wherein the denominator comprises an integer part of a pitch lag indicated by the pitch lag information and a multi-tap filter depending on a fractional part of the pitch lag.

15

2. Apparatus of claim 1, wherein the transfer function of the post-filter comprises, in the numerator, a further multi-tap FIR filter for a zero fractional part of the pitch lag.

20

3. Apparatus of claim 1 or 2, wherein the denominator comprises a product between the multi-tap filter and the gain value.

4. Apparatus of one of the preceding claims, wherein the numerator furthermore comprises a product of a first scalar value and a second scalar value, wherein the denominator comprises the second scalar value and not the first scalar value, wherein the first and second scalar values are predetermined and have values greater than 0 and lower than, and wherein the second scalar value is lower than the first scalar value.

25

5. Apparatus of claim 4, further comprising:

30

a filter controller (114) configured for setting the second scalar value depending on a bitrate, by which the frequency-time converter (100) is operated, wherein the second scalar value is set to a first value, when the bitrate has a first value, wherein the second scalar value is set to a second value, when the bitrate has a second value, wherein the second value of the bitrate is lower than the first value

35

of the bitrate, and wherein the second value of the second scalar value is greater than the first value of the second scalar value.

6. The apparatus in accordance with claim 4 or 5, wherein the first scalar value is set
5 between 0.6 and 1.0 and wherein the second scalar value is set between 0.1 and 0.5.

7. Apparatus of one of the preceding claims,
10 wherein the post-filter has the transfer function $H(z)$ in a pole-zero representation based on the following equation:

$$H(z) = \frac{1 - \alpha\beta g B(z, 0)}{1 - \beta g B(z, T_{fr}) z^{-T_{int}}}$$

15 wherein α is a first scalar value, wherein β is a second scalar value, wherein $B(z, 0)$ is a multi-tap filter for a zero fractional part pitch lag, wherein $B(z, T_{fr})$ is a multi-tap filter depending on the fractional part of the pitch lag, wherein T_{int} is the integer part of the pitch lag, wherein T_{fr} is the fractional part of the pitch lag, and wherein g is the gain value indicated by the gain information z is a variable in a z -plane.

20 8. Apparatus of one of the preceding claims, wherein the multi-tap filter is a finite impulse response (FIR) filter and has at least three taps.

9. Apparatus of one of the preceding claims,
25 wherein the multi-tap filter in the denominator comprises four taps, wherein, for a zero fractional part, the first tap is between 0.0 and 0.1, the second tap is between 0.2 and 0.3, the third tap is between 0.5 and 0.6, and the fourth tap is between 0.2 and 0.3,

30 wherein the multi-tap filter comprises, for a first fractional part, four filter taps, wherein the first tap is between 0.0 and 0.1, the second tap is between 0.3 and 0.4, the third tap is between 0.45 and 0.55, and the fourth tap is between 0.1 and 0.2,

35

wherein the multi-tap filter comprises, for a second fractional part, four filter taps, wherein the first tap is between 0.0 and 0.1, the second tap is between 0.35 and 0.45, the third tap is between 0.35 and 0.45, and the fourth tap is between 0.0 and 0.1,

5

wherein the multi-tap filter comprises, for a third fractional part, four filter taps, wherein the first tap is between 0.1 and 0.2, the second tap is between 0.45 and 0.55, the third tap is between 0.3 and 0.4, and the fourth tap is between 0.0 and 0.1,

10

wherein the third fractional part is greater than the second fractional part, and wherein the second fractional part is greater than the first fractional part.

10. Apparatus of one of the preceding claims,

15

wherein the post-filter is configured to have a negative spectral tilt for compensating a loss in energy by the harmonic post-filter, or

wherein the post-filter is configured to suppress an amount of energy between harmonics in a frame, wherein the amount of energy suppressed is smaller than 20% of a total energy of the time-domain representation in the frame.

20

11. Apparatus of one of preceding claims,

25

wherein the domain converter is a frequency-time converter, wherein the first domain is a frequency domain and the second domain is a time domain, or

wherein the domain converter is an LPC residual-time converter, wherein the first domain is an LPC residual domain and the second domain is a time domain."

30

12. Method of processing an audio signal having associated therewith a pitch lag information and a gain information, comprising:

converting (100) a frequency representation of the audio signal into a time-domain representation of the audio signal; and

35

filtering the time-domain representation of the audio signal by a harmonic post-filter (104), wherein the post-filter is based on a transfer function comprising a

numerator and a denominator, wherein the numerator comprises a gain value indicated by the gain information, and wherein the denominator comprises an integer part of a pitch lag indicated by the pitch lag information and a multi-tap filter depending on a fractional part of the pitch lag.

5

13. System for processing an audio signal comprising an encoder for encoding an audio signal and a decoder comprising a processor, the processor comprising:

10

a domain converter (100) for converting a frequency representation of the audio signal into a time-domain representation of the audio signal; and

a harmonic post-filter (104) for filtering the time-domain representation of the audio signal,

15

wherein the post-filter is based on a transfer function comprising a numerator and a denominator, wherein the numerator comprises a gain value indicated by a gain information, and wherein the denominator comprises an integer part of a pitch lag indicated by a pitch lag information and a multi-tap filter depending on a fractional part of the pitch lag.

20

14. System of claim 13, wherein the encoder comprises a pitch lag calculator (402, 404, 406) for calculating an integer part and a fractional part of the pitch lag and a gain calculator (410, 412) for calculating the gain value, and an encoded signal former (414) for generating an encoded signal (102) comprising the pitch lag information and the gain information.

25

15. Method of processing an audio signal comprising a method of encoding an audio signal and a method of decoding comprising:

30

converting (100) a frequency representation of the audio signal into a time-domain representation of the audio signal; and

35

filtering the time-domain representation of the audio signal using a harmonic post-filter (104), wherein the post-filter is based on a transfer function comprising a numerator and a denominator, wherein the numerator comprises a gain value indicated by a gain information, and wherein the denominator comprises an integer

part of a pitch lag indicated by a pitch lag information and a multi-tap filter depending on a fractional part of the pitch lag.

16. Computer program for performing a method of claim 12 or claim 15, when the
5 computer program is running on a computer or a processor.