




(86) Date de dépôt PCT/PCT Filing Date: 2014/06/23 (87) Date publication PCT/PCT Publication Date: 2014/12/24 (45) Date de délivrance/Issue Date: 2018/05/22 (85) Entrée phase nationale/National Entry: 2015/11/25 (86) N° demande PCT/PCT Application No.: EP 2014/063173 (87) N° publication PCT/PCT Publication No.: 2014/202786 (30) Priorités/Priorities: 2013/06/21 (EP13173154.9); 2014/05/05 (EP14166998.6)	(51) Cl.Int./Int.Cl. <i>G10L 19/005</i> (2013.01), <i>G10L 19/09</i> (2013.01), <i>G10L 25/90</i> (2013.01) (72) Inventeurs/Inventors: SCHNABEL, MICHAEL, DE; MARKOVIC, GORAN, DE; SPERSCHNEIDER, RALPH, DE; LECOMTE, JEREMIE, DE; HELMRICH, CHRISTIAN, DE (73) Propriétaire/Owner: FRAUNHOFER-GESELLSCHAFT ZUR FORDERUNG DER ANGEWANDTEN FORSCHUNG E.V., DE (74) Agent: PERRY + CURRIER
---	--

(54) Titre : APPAREIL ET PROCEDE POUR GENERER UNE FORME SPECTRALE ADAPTABLE DE BRUIT DE CONFORT

(54) Title: APPARATUS AND METHOD FOR GENERATING AN ADAPTIVE SPECTRAL SHAPE OF COMFORT NOISE



(57) **Abrégé/Abstract:**
An apparatus for decoding an encoded audio signal to obtain a reconstructed audio signal is provided. The apparatus comprises a receiving interface (1110) for receiving one or more frames, a coefficient generator (1120), and a signal reconstructor (1130). The

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

coefficient generator (1120) is configured to determine, if a current frame of the one or more frames is received by the receiving interface (1110) and if the current frame being received by the receiving interface (1110) is not corrupted, one or more first audio signal coefficients, being comprised by the current frame, wherein said one or more first audio signal coefficients indicate a characteristic of the encoded audio signal, and one or more noise coefficients indicating a background noise of the encoded audio signal. Moreover, the coefficient generator (1120) is configured to generate one or more second audio signal coefficients, depending on the one or more first audio signal coefficients and depending on the one or more noise coefficients, if the current frame is not received by the receiving interface (1110) or if the current frame being received by the receiving interface (1110) is corrupted. The audio signal reconstructor (1130) is configured to reconstruct a first portion of the reconstructed audio signal depending on the one or more first audio signal coefficients, if the current frame is received by the receiving interface (1110) and if the current frame being received by the receiving interface (1110) is not corrupted. Moreover, the audio signal reconstructor (1130) is configured to reconstruct a second portion of the reconstructed audio signal depending on the one or more second audio signal coefficients, if the current frame is not received by the receiving interface (1110) or if the current frame being received by the receiving interface (1110) is corrupted.

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

(19) World Intellectual Property
Organization
International Bureau(43) International Publication Date
24 December 2014 (24.12.2014)

WIPO | PCT

(10) International Publication Number
WO 2014/202786 A1

(51) International Patent Classification:

G10L 19/005 (2013.01) *G10L 19/09* (2013.01)
G10L 25/90 (2013.01)

(21) International Application Number:

PCT/EP2014/063173

(22) International Filing Date:

23 June 2014 (23.06.2014)

(25) Filing Language:

English

(26) Publication Language:

English

(30) Priority Data:

13173154.9 21 June 2013 (21.06.2013) EP
14166998.6 5 May 2014 (05.05.2014) EP(71) Applicant: **FRAUNHOFER-GESELLSCHAFT ZUR FÖRDERUNG DER ANGEWANDTEN FORSCHUNG E.V.** [DE/DE]; Hansastraße 27c, 80686 München (DE).(72) Inventors: **SCHNABEL, Michael**; Schneidbergweg 4a, 95179 Geroldsgrün (DE). **MARKOVIC, Goran**; Aachener Straße 19, 90425 Nürnberg (DE). **SPERSCHNEIDER, Ralph**; Debert 75, 91320 Ebermannstadt (DE). **LECOMTE, Jérémie**; Turnstraße 7, 90763 Fürth (DE). **HELMRICH, Christian**; Hauptstraße 68, 91054 Erlangen (DE).(74) Agents: **ZINKLER, Franz** et al.; Hermann-Roth-Weg 1, 82049 Pullach (DE).(81) Designated States (*unless otherwise indicated, for every kind of national protection available*): AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BN, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IR, IS, JP, KE, KG, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LT, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PA, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SA, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.(84) Designated States (*unless otherwise indicated, for every kind of regional protection available*): ARIPO (BW, GH, GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, RU, TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, KM, ML, MR, NE, SN, TD, TG).

Published:

— with international search report (Art. 21(3))

(54) Title: APPARATUS AND METHOD FOR GENERATING AN ADAPTIVE SPECTRAL SHAPE OF COMFORT NOISE




FIG 11

(57) **Abstract:** An apparatus for decoding an encoded audio signal to obtain a reconstructed audio signal is provided. The apparatus comprises a receiving interface (1110) for receiving one or more frames, a coefficient generator (1120), and a signal reconstructor (1130). The coefficient generator (1120) is configured to determine, if a current frame of the one or more frames is received by the receiving interface (1110) and if the current frame being received by the receiving interface (1110) is not corrupted, one or more first audio signal coefficients, being comprised by the current frame, wherein said one or more first audio signal coefficients indicate a characteristic of the encoded audio signal, and one or more noise coefficients indicating a background noise of the encoded audio signal. Moreover, the coefficient generator (1120) is configured to generate one or more second audio signal coefficients, depending on the one or more first audio signal coefficients and depending on the one or more noise coefficients, if the current frame is not received by the receiving interface (1110) or if the current frame being received by the receiving interface (1110) is corrupted. The audio signal reconstructor (1130) is configured to reconstruct a first portion of the reconstructed audio signal depending on the one or more first audio signal coefficients, if the current frame is received by the receiving interface (1110) and if the current frame being received by the receiving interface (1110) is not corrupted. Moreover, the audio signal reconstructor (1130) is configured to reconstruct a second portion of the reconstructed audio signal depending on the one or more second audio signal coefficients, if the current frame is not received by the receiving interface (1110) or if the current frame being received by the receiving interface (1110) is corrupted.

Apparatus and Method for Generating an Adaptive Spectral Shape of Comfort NoiseDescription

5

The present invention relates to audio signal encoding, processing and decoding, and, in particular, to an apparatus and method for improved signal fade out for switched audio coding systems during error concealment.

10 In the following, the state of the art is described regarding speech and audio codecs fade out during packet loss concealment (PLC). The explanations regarding the state of the art start with the ITU-T codecs of the G-series (G.718, G.719, G.722, G.722.1, G.729, G.729.1), are followed by the 3GPP codecs (AMR, AMR-WB, AMR-WB+) and one IETF codec (OPUS), and conclude with two MPEG codecs (HE-AAC, HILN) (ITU = International
15 Telecommunication Union; 3GPP = 3rd Generation Partnership Project; AMR = Adaptive Multi-Rate; WB = Wideband; IETF = Internet Engineering Task Force). Subsequently, the state-of-the art regarding tracing the background noise level is analysed, followed by a summary which provides an overview.

20 At first, G.718 is considered. G.718 is a narrow-band and wideband speech codec, that supports DTX/CNG (DTX = Digital Theater Systems; CNG = Comfort Noise Generation). As embodiments particularly relate to low delay code, the low delay version mode will be described in more detail, here.

25 Considering ACELP (Layer 1) (ACELP = Algebraic Code Excited Linear Prediction), the ITU-T recommends for G.718 [ITU08a, section 7.11] an adaptive fade out in the linear predictive domain to control the fading speed. Generally, the concealment follows this principle:

30 According to G.718, in case of frame erasures, the concealment strategy can be summarized as a convergence of the signal energy and the spectral envelope to the estimated parameters of the background noise. The periodicity of the signal is converged to zero. The speed of the convergence is dependent on the parameters of the last correctly received frame and the number of consecutive erased frames, and is controlled
35 by an attenuation factor, α . The attenuation factor α , is further dependent on the stability, θ , of the LP filter (LP = Linear Prediction) for UNVOICED frames. In general, the convergence is slow if the last good received frame is in a stable segment and is rapid if the frame is in a transition segment.

The attenuation factor α depends on the speech signal class, which is derived by signal classification described in [ITU08a, section 6.8.1.3.1 and 7.11.1.1]. The stability factor θ is computed based on a distance measure between the adjacent ISF (Immittance Spectral
5 Frequency) filters [ITU08a, section 7.1.2.4.2].

Table 1 shows the calculation scheme of α :

last good received frame	Number of successive erased frames	α
ARTIFICIAL ONSET		0.6
ONSET, VOICED	≤ 3	1.0
	> 3	0.4
VOICED TRANSITION		0.4
UNVOICED TRANSITION		0.8
UNVOICED	$= 1$	$0.2 \cdot \theta + 0.8$
	$= 2$	0.6
	> 2	0.4

10 Table 1: Values of the attenuation factor α , the value θ is a stability factor computed from a distance measure between the adjacent LP filters. [ITU08a, section 7.1.2.4.2].

Moreover, G.718 provides a fading method in order to modify the spectral envelope. The
15 general idea is to converge the last ISF parameters towards an adaptive ISF mean vector. At first, an average ISF vector is calculated from the last 3 known ISF vectors. Then the average ISF vector is again averaged with an offline trained long term ISF vector (which is a constant vector) [ITU08a, section 7.11.1.2].

20 Moreover, G.718 provides a fading method to control the long term behavior and thus the interaction with the background noise, where the pitch excitation energy (and thus the excitation periodicity) is converging to 0, while the random excitation energy is converging to the CNG excitation energy [ITU08a, section 7.11.1.6]. The innovation gain attenuation is calculated as

25

$$g_s^{[1]} = \alpha g_s^{[0]} + (1 - \alpha) g_n \quad (1)$$

where $g_s^{[1]}$ is the innovative gain at the beginning of the next frame, $g_s^{[0]}$ is the innovative gain at the beginning of the current frame, g_n is the gain of the excitation used during the comfort noise generation and the attenuation factor α .

- 5 Similarly to the periodic excitation attenuation, the gain is attenuated linearly throughout the frame on a sample-by-sample basis starting with, $g_s^{[0]}$, and reaches $g_s^{[1]}$ at the beginning of the next frame.

Fig. 2 outlines the decoder structure of G.718. In particular, Fig. 2 illustrates a high level
10 G.718 decoder structure for PLC, featuring a high pass filter.

By the above-described approach of G.718, the innovative gain g_s converges to the gain used during comfort noise generation g_n for long bursts of packet losses. As described in [ITU08a, section 6.12.3], the comfort noise gain g_n is given as the square root of the
15 energy $\tilde{E}$. The conditions of the update of $\tilde{E}$ are not described in detail. Following the reference implementation (floating point C-code, `stat_noise_uv_mod.c`), $\tilde{E}$ is derived as follows:

```

if(unvoiced_vad == 0){
20     if( unv_cnt > 20 ){
        ftmp = lp_gainc * lp_gainc;
        lp_ener = 0.7f * lp_ener + 0.3f * ftmp;
    }
    else{
25         unv_cnt++;
    }
}
else{
    unv_cnt = 0;
30 }

```

wherein `unvoiced_vad` holds the voice activity detection, wherein `unv_cnt` holds the number of unvoiced frames in a row, wherein `lp_gainc` holds the low passed gains of the fixed codebook, and wherein `lp_ener` holds the low passed CNG energy estimate
35 $\tilde{E}$, it is initialized with 0.

Furthermore, G.718 provides a high pass filter, introduced into the signal path of the unvoiced excitation, if the signal of the last good frame was classified different from UNVOICED, see Fig. 2, also see [ITU08a, section 7.11.1.6]. This filter has a low shelf characteristic with a frequency response at DC being around 5 dB lower than at Nyquist
5 frequency.

Moreover, G.718 proposes a decoupled LTP feedback loop (LTP = Long-Term Prediction): While during normal operation the feedback loop for the adaptive codebook is updated subframe-wise ([ITU08a, section 7.1.2.1.4]) based on the full
10 excitation. During concealment this feedback loop is updated frame-wise (see [ITU08a, sections 7.11.1.4, 7.11.2.4, 7.11.1.6, 7.11.2.6; `dec_GV_exc@dec_gen_voic.c` and `syn_bfi_post@syn_bfi_pre_post.c`]) based on the voiced excitation only. With this approach, the adaptive codebook is not “polluted” with noise having its origin in by the randomly chosen innovation excitation.

15 Regarding the transform coded enhancement layers (3-5) of G.718, during concealment, the decoder behaves regarding the high layer decoding similar to the normal operation, just that the MDCT spectrum is set to zero. No special fade-out behavior is applied during concealment.

20 With respect to CNG, in G.718, the CNG synthesis is done in the following order. At first, parameters of a comfort noise frame are decoded. Then, a comfort noise frame is synthesized. Afterwards the pitch buffer is reset. Then, the synthesis for the FER (Frame Error Recovery) classification is saved. Afterwards, spectrum deemphasis is conducted.
25 Then low frequency post-filtering is conducted. Then, the CNG variables are updated.

In the case of concealment, exactly the same is performed, except the CNG parameters are not decoded from the bitstream. This means that the parameters are not updated during the frame loss, but the decoded parameters from the last good SID (Silence
30 Insertion Descriptor) frame are used.

Now, G.719 is considered. G.719, which is based on Siren 22, is a transform based full-band audio codec. The ITU-T recommends for G.719 a fade-out with frame repetition in the spectral domain [ITU08b, section 8.6]. According to G.719, a frame erasure
35 concealment mechanism is incorporated into the decoder. When a frame is correctly received, the reconstructed transform coefficients are stored in a buffer. If the decoder is informed that a frame has been lost or that a frame is corrupted, the transform coefficients reconstructed in the most recently received frame are decreasingly scaled with a factor

0.5 and then used as the reconstructed transform coefficients for the current frame. The decoder proceeds by transforming them to the time domain and performing the windowing-overlap-add operation.

5 In the following, G.722 is described. G.722 is a 50 to 7000 Hz coding system which uses subband adaptive differential pulse code modulation (SB-ADPCM) within a bitrate up to 64 kbit/s. The signal is split into a higher and a lower subband, using a QMF analysis (QMF = Quadrature Mirror Filter). The resulting two bands are ADPCM-coded (ADPCM = Adaptive Differential Pulse Code Modulation).

10

For G.722, a high-complexity algorithm for packet loss concealment is specified in Appendix III [ITU06a] and a low-complexity algorithm for packet loss concealment is specified in Appendix IV [ITU07]. G.722 - Appendix III ([ITU06a, section III.5]) proposes a gradually performed muting, starting after 20ms of frame-loss, being completed after
15 60ms of frame-loss. Moreover, G.722 - Appendix IV proposes a fade-out technique which applies "to each sample a gain factor that is computed and adapted sample by sample" [ITU07, section IV.6.1.2.7].

20

In G.722, the muting process takes place in the subband domain just before the QMF synthesis and as the last step of the PLC module. The calculation of the muting factor is performed using class information from the signal classifier which also is part of the PLC module. The distinction is made between classes TRANSIENT, UV_TRANSITION and others. Furthermore, distinction is made between single losses of 10-ms frames and other cases (multiple losses of 10-ms frames and single/multiple losses of 20-ms frames).

25

This is illustrated by Fig. 3. In particular, Fig. 3 depicts a scenario, where the fade-out factor of G.722, depends on class information and wherein 80 samples are equivalent to 10 ms.

30

According to G.722, the PLC module creates the signal for the missing frame and some additional signal (10ms) which is supposed to be cross-faded with the next good frame. The muting for this additional signal follows the same rules. In highband concealment of G.722, cross-fading does not take place.

35

In the following, G.722.1 is considered. G.722.1, which is based on Siren 7, is a transform based wide band audio codec with a super wide band extension mode, referred to as G.722.1C. G. 722.1C itself is based on Siren 14. The ITU-T recommends for G.722.1 a frame-repetition with subsequent muting [ITU05, section 4.7]. If the decoder is informed,

by means of an external signaling mechanism not defined in this recommendation, that a frame has been lost or corrupted, it repeats the previous frame's decoded MLT (Modulated Lapped Transform) coefficients. It proceeds by transforming them to the time domain, and performing the overlap and add operation with the previous and next frame's
 5 decoded information. If the previous frame was also lost or corrupted, then the decoder sets all the current frames MLT coefficients to zero.

Now, G.729 is considered. G.729 is an audio data compression algorithm for voice that compresses digital voice in packets of 10 milliseconds duration. It is officially described as
 10 Coding of speech at 8 kbit/s using code-excited linear prediction speech coding (CS-ACELP) [ITU12].

As outlined in [CPK08], G.729 recommends a fade-out in the LP domain. The PLC algorithm employed in the G.729 standard reconstructs the speech signal for the current
 15 frame based on previously-received speech information. In other words, the PLC algorithm replaces the missing excitation with an equivalent characteristic of a previously received frame, though the excitation energy gradually decays finally, the gains of the adaptive and fixed codebooks are attenuated by a constant factor.

20 The attenuated fixed-codebook gain is given by:

$$g_c^{(m)} = 0.98 \cdot g_c^{(m-1)}$$

with m is the subframe index.

25

The adaptive-codebook gain is based on an attenuated version of the previous adaptive-codebook gain:

$$g_p^{(m)} = 0.9 \cdot g_p^{(m-1)} \quad , \text{ bounded by } g_p^{(m)} < 0.9$$

30

Nam in Park et al. suggest for G.729, a signal amplitude control using prediction by means of linear regression [CPK08, PKJ+11]. It is addressed to burst packet loss and uses linear regression as a core technique. Linear regression is based on the linear model
 as

35

$$g'_i = a + bi \quad (2)$$

where g'_i is the newly predicted current amplitude, a and b are coefficients for the first order linear function, and i is the index of the frame. In order to find the optimized coefficients a^* and b^* , the summation of the squared prediction error is minimized:

$$\epsilon = \sum_{j=i-4}^{i-1} (g_j - g'_j)^2 \quad (3)$$

ϵ is the squared error, g_j is the original past j -th amplitude. To minimize this error, simply the derivative regarding a and b is set to zero. By using the optimized parameters a^* and b^* , an estimate of each g_i^* is denoted by

$$g_i^* = a^* + b^*i \quad (4)$$

Fig. 4 shows the amplitude prediction, in particular, the prediction of the amplitude g_i^* , by using linear regression.

To obtain the amplitude A'_i of the lost packet i , a ratio σ_i

$$\sigma_i = \frac{g_i^*}{g_{i-1}} \quad (5)$$

is multiplied with a scale factor S_i :

$$A'_i = S_i * \sigma_i \quad (6)$$

wherein the scale factor S_i depends on the number of consecutive concealed frames $l(i)$:

$$S_i = \begin{cases} 1.0, & \text{if } l(i) = 1, 2 \\ 0.9, & \text{if } l(i) = 3, 4 \\ 0.8, & \text{if } l(i) = 5, 6 \\ 0, & \text{otherwise} \end{cases} \quad (7)$$

In [PKJ+11], a slightly different scaling is proposed.

According to G.729, afterwards, A'_i will be smoothed to prevent discrete attenuation at frame borders. The final, smoothed amplitude $A_i(n)$ is multiplied to the excitation, obtained from the previous PLC components.

- 5 In the following, G.729.1 is considered. G.729.1 is a G.729-based embedded variable bit-rate coder: An 8-32 kbit/s scalable wideband coder bitstream inter-operable with G.729 [ITU06b].

10 According to G.729.1, as in G.718 (see above), an adaptive fade out is proposed, which depends on the stability of the signal characteristics ([ITU06b, section 7.6.1]). During concealment, the signal is usually attenuated based on an attenuation factor α which depends on the parameters of the last good received frame class and the number of consecutive erased frames. The attenuation factor α is further dependent on the stability of the LP filter for UNVOICED frames. In general, the attenuation is slow if the last good
15 received frame is in a stable segment and is rapid if the frame is in a transition segment.

Furthermore, the attenuation factor α depends on the average pitch gain per subframe $\bar{g}_p$ ([ITU06b, eq. 163, 164]):

20
$$\bar{g}_p = 0.1g_p^{(0)} + 0.2g_p^{(1)} + 0.3g_p^{(2)} + 0.4g_p^{(3)} \quad (8)$$

where $g_p^{(i)}$ is the pitch gain in subframe i .

Table 2 shows the calculation scheme of α , where

25
$$\beta = \sqrt{\bar{g}_p} \quad \text{with} \quad 0.85 \geq \beta \geq 0.98 \quad (9)$$

During the concealment process, α is used in the following concealment tools:

30

last good received frame	Number of successive erased frames	α
VOICED	1	β
	2,3	$\bar{g}_p$
	> 3	0.4
ONSET	1	0.8β
	2,3	$\bar{g}_p$
	> 3	0.4

ARTIFICIAL ONSET	1	0.6β
	2,3	$\bar{g}_p$
	> 3	0.4
VOICED TRANSITION	≤ 2	0.8
	> 2	0.2
UNVOICED TRANSITION		0.88
UNVOICED	1	0.95
	2,3	$0.6 \theta + 0.4$
	> 3	0.4

Table 2: Values of the attenuation factor α , the value θ is a stability factor computed from a distance measure between the adjacent LP filters. [ITU06b, section 7.6.1].

- 5 According to G.729.1, regarding glottal pulse resynchronization, as the last pulse of the excitation of the previous frame is used for the construction of the periodic part, its gain is approximately correct at the beginning of the concealed frame and can be set to 1. The gain is then attenuated linearly throughout the frame on a sample-by-sample basis to achieve the value of α at the end of the frame. The energy evolution of voiced segments is
- 10 extrapolated by using the pitch excitation gain values of each subframe of the last good frame. In general, if these gains are greater than 1, the signal energy is increasing, if they are lower than 1, the energy is decreasing. α is thus set to $\beta = \sqrt{\bar{g}_p}$ as described above, see [ITU06b, eq. 163, 164]. The value of β is clipped between 0.98 and 0.85 to avoid strong energy increases and decreases, see [ITU06b, section 7.6.4].

15

Regarding the construction of the random part of the excitation, according to G.729.1, at the beginning of an erased block, the innovation gain g_s is initialized by using the innovation excitation gains of each subframe of the last good frame:

20
$$g_s = 0.1g^{(0)} + 0.2g^{(1)} + 0.3g^{(2)} + 0.4g^{(3)}$$

wherein $g^{(0)}$, $g^{(1)}$, $g^{(2)}$ and $g^{(3)}$ are the fixed codebook, or innovation, gains of the four subframes of the last correctly received frame. The innovation gain attenuation is done as:

25
$$g_s^{(1)} = \alpha \cdot g_s^{(0)}$$

wherein $g_s^{(1)}$ is the innovation gain at the beginning of the next frame, $g_s^{(0)}$ is the innovation gain at the beginning of the current frame, and α is as defined in Table 2

above. Similarly to the periodic excitation attenuation, the gain is thus linearly attenuated throughout the frame on a sample by sample basis starting with $g_s^{(0)}$ and going to the value of $g_s^{(1)}$ that would be achieved at the beginning of the next frame.

- 5 According, to G.729.1, if the last good frame is UNVOICED, only the innovation excitation is used and it is further attenuated by a factor of 0.8. In this case, the past excitation buffer is updated with the innovation excitation as no periodic part of the excitation is available, see [ITU06b, section 7.6.6].
- 10 In the following, AMR is considered. 3GPP AMR [3GP12b] is a speech codec utilizing the ACELP algorithm. AMR is able to code speech with a sampling rate of 8000 samples/s and a bitrate between 4.75 and 12.2 kbit/s and supports signaling silence descriptor frames (DTX/CNG).
- 15 In AMR, during error concealment (see [3GP12a]), it is distinguished between frames which are error prone (bit errors) and frames, that are completely lost (no data at all).

For ACELP concealment, AMR introduces a state machine which estimates the quality of the channel: The larger the value of the state counter, the worse the channel quality is.

- 20 The system starts in state 0. Each time a bad frame is detected, the state counter is incremented by one and is saturated when it reaches 6. Each time a good speech frame is detected, the state counter is reset to zero, except when the state is 6, where the state counter is set to 5. The control flow of the state machine can be described by the following C code (BFI is a bad frame indicator, State is a state variable):

25

```

if(BFI != 0 ) {
    State = State + 1;
}
else if(State == 6) {
30     State = 5;
}
else {
    State = 0;
}
35 if(State > 6 ) {
    State = 6;
}

```

In addition to this state machine, in AMR, the bad frame flags from the current and the previous frames are checked (prevBFI).

5 Three different combinations are possible:

The first one of the three combinations is BFI = 0, prevBFI = 0, State = 0: No error is detected in the received or in the previous received speech frame. The received speech parameters are used in the normal way in the speech synthesis. The current frame of
10 speech parameters is saved.

The second one of the three combinations is BFI = 0, prevBFI = 1, State = 0 or 5: No error is detected in the received speech frame, but the previous received speech frame was bad. The LTP gain and fixed codebook gain are limited below the values used for the last
15 received good subframe:

$$g_p = \begin{cases} g_p, & g_p \leq g_p(-1) \\ g_p(-1), & g_p > g_p(-1) \end{cases} \quad (10)$$

where g_p = current decoded LTP gain, $g_p(-1)$ = LTP gain used for the last good subframe
20 (BFI = 0), and

$$g_c = \begin{cases} g_c, & g_c \leq g_c(-1) \\ g_c(-1), & g_c > g_c(-1) \end{cases} \quad (11)$$

where g_c = current decoded fixed codebook gain, and $g_c(-1)$ = fixed codebook gain used
25 for the last good subframe (BFI = 0).

The rest of the received speech parameters are used normally in the speech synthesis. The current frame of speech parameters is saved.

30 The third one of the three combinations is BFI = 1, prevBFI = 0 or 1, State = 1...6: An error is detected in the received speech frame and the substitution and muting procedure is started. The LTP gain and fixed codebook gain are replaced by attenuated values from the previous subframes:

$$g_p = \begin{cases} P(state) \cdot g_p(-1), & g_p(-1) \leq \text{median5}(g_p(-1), \dots, g_p(-5)) \\ P(state) \cdot \text{median5}(g_p(-1), \dots, g_p(-5)) & g_p(-1) > \text{median5}(g_p(-1), \dots, g_p(-5)) \end{cases} \quad (12)$$

where g_p indicates the current decoded LTP gain and $g_p(-1), \dots, g_p(-n)$ indicate the LTP gains used for the last n subframes and $\text{median5}()$ indicates a 5-point median operation and

$P(\text{state})$ = attenuation factor,

10 where ($P(1) = 0.98$, $P(2) = 0.98$, $P(3) = 0.8$, $P(4) = 0.3$, $P(5) = 0.2$, $P(6) = 0.2$) and state = state number, and

$$g_c = \begin{cases} C(state) \cdot g_c(-1), & g_c(-1) \leq \text{median5}(g_c(-1), \dots, g_c(-5)) \\ C(state) \cdot \text{median5}(g_c(-1), \dots, g_c(-5)) & g_c(-1) > \text{median5}(g_c(-1), \dots, g_c(-5)) \end{cases} \quad (13)$$

15

where g_c indicates the current decoded fixed codebook gain and $g_c(-1), \dots, g_c(-n)$ indicate the fixed codebook gains used for the last n subframes and $\text{median5}()$ indicates a 5-point median operation and $C(\text{state})$ = attenuation factor, where ($C(1) = 0.98$, $C(2) = 0.98$, $C(3) = 0.98$, $C(4) = 0.98$, $C(5) = 0.98$, $C(6) = 0.7$) and state = state number.

20

In AMR, the LTP-lag values (LTP = Long-Term Prediction) are replaced by the past value from the 4th subframe of the previous frame (12.2 mode) or slightly modified values based on the last correctly received value (all other modes).

25

According to AMR, the received fixed codebook innovation pulses from the erroneous frame are used in the state in which they were received when corrupted data are received. In the case when no data were received random fixed codebook indices should be employed.

30

Regarding CNG in AMR, according to [3GP12a, section 6.4], each first lost SID frame is substituted by using the SID information from earlier received valid SID frames and the procedure for valid SID frames is applied. For subsequent lost SID frames, an attenuation technique is applied to the comfort noise that will gradually decrease the output level. Therefore it is checked if the last SID update was more than 50 frames (=1 s) ago, if yes, the output will be muted (level attenuation by -6/8 dB per frame [3GP12d,

35

dtx_dec{}@sp_dec.c] which yields 37.5 dB per second). Note that the fade-out applied to CNG is performed in the LP domain.

In the following, AMR-WB is considered. Adaptive Multirate - WB [ITU03, 3GP09c] is a speech codec, ACELP, based on AMR (see section 1.8). It uses parametric bandwidth extension and also supports DTX/CNG. In the description of the standard [3GP12g] there are concealment example solutions given which are the same as for AMR [3GP12a] with minor deviations. Therefore, just the differences to AMR are described here. For the standard description, see the description above.

10

Regarding ACELP, in AMR-WB, the ACELP fade-out is performed based on the reference source code [3GP12c] by modifying the pitch gain g_p (for AMR above referred to as LTP gain) and by modifying the code gain g_c .

15 In case of lost frame, the pitch gain g_p for the first subframe is the same as in the last good frame, except that it is limited between 0.95 and 0.5. For the second, the third and the following subframes, the pitch gain g_p is decreased by a factor of 0.95 and again limited.

AMR-WB proposes that in a concealed frame, g_c is based on the last g_c :

20

$$g_{c,current} = g_{c,past} * (1.4 - g_{p,past}) \quad (14)$$

$$g_c = g_{c,current} * g_{c_{inov}} \quad (15)$$

25

$$g_{c_{inov}} = \frac{1.0}{\sqrt{\frac{ener_{inov}}{\text{subframe_size}}}} \quad (16)$$

$$ener_{inov} = \sum_{i=0}^{\text{subframe_size}-1} code[i] \quad (17)$$

30 For concealing the LTP-lags, in AMR-WB, the history of the five last good LTP-lags and LTP-gains are used for finding the best method to update, in case of a frame loss. In case the frame is received with bit errors a prediction is performed, whether the received LTP lag is usable or not [3GP12g].

Regarding CNG, in AMR-WB, if the last correctly received frame was a SID frame and a frame is classified as lost, it shall be substituted by the last valid SID frame information and the procedure for valid SID frames should be applied.

- 5 For subsequent lost SID frames, AMR-WB proposes to apply an attenuation technique to the comfort noise that will gradually decrease the output level. Therefore it is checked if the last SID update was more than 50 frames ($=1\text{ s}$) ago, if yes, the output will be muted (level attenuation by $-3/8\text{ dB}$ per frame [3GP12f, `dtx_dec{}@dtx.c`] which yields 18.75 dB per second). Note that the fade-out applied to CNG is performed in the LP
10 domain.

Now, AMR-WB+ is considered. Adaptive Multirate - WB+ [3GP09a] is a switched codec using ACELP and TCX (TCX = Transform Coded Excitation) as core codecs. It uses parametric bandwidth extension and also supports DTX/CNG.

15

In AMR-WB+, a mode extrapolation logic is applied to extrapolate the modes of the lost frames within a distorted superframe. This mode extrapolation is based on the fact that there exists redundancy in the definition of mode indicators. The decision logic (given in [3GP09a, figure 18]) proposed by AMR-WB+ is as follows:

20

- A vector mode, $(m_{-1}, m_0, m_1, m_2, m_3)$, is defined, where m_{-1} indicates the mode of the last frame of the previous superframe and m_0, m_1, m_2, m_3 indicate the modes of the frames in the current superframe (decoded from the bitstream), where $m_k = -1, 0, 1, 2$ or 3 (-1 : lost, 0 : ACELP, 1 : TCX20, 2 : TCX40, 3 : TCX80), and where the
25 number of lost frames n_{loss} may be between 0 and 4.

- If $m_{-1} = 3$ and two of the mode indicators of the frames $0 - 3$ are equal to three, all indicators will be set to three because then it is for sure that one TCX80 frame was indicated within the superframe.

30

- If only one indicator of the frames $0 - 3$ is three (and the number of lost frames n_{loss} is three), the mode will be set to $(1, 1, 1, 1)$, because then $3/4$ of the TCX80 target spectrum is lost and it is very likely that the global TCX gain is lost.

35

- If the mode is indicating $(x, 2, -1, x, x)$ or $(x, -1, 2, x, x)$, it will be extrapolated to $(x, 2, 2, x, x)$, indicating a TCX40 frame. If the mode indicates $(x, x, x, 2, -1)$ or $(x, x, -1, 2)$ it will be extrapolated to $(x, x, x, 2, 2)$, also indicating a TCX40 frame. It should be noted that $(x, [0, 1], 2, 2, [0, 1])$ are invalid configurations.

- After that, for each frame that is lost ($mode = -1$), the mode is set to ACELP ($mode = 0$) if the preceding frame was ACELP and the mode is set to TCX20 ($mode = 1$) for all other cases.

5

Regarding ACELP, according to AMR-WB+, if a lost frames mode results in $m_k = 0$ after the mode extrapolation, the same approach as in [3GP12g] is applied for this frame (see above).

- 10 In AMR-WB+, depending on the number of lost frames and the extrapolated mode, the following TCX related concealment approaches are distinguished (TCX = Transform Coded Excitation):

- 15 - If a full frame is lost, then an ACELP like concealment is applied: The last excitation is repeated and concealed ISF coefficients (slightly shifted towards their adaptive mean) are used to synthesize the time domain signal. Additionally, a fade-out factor of 0.7 per frame (20ms) [3GP09b, `dec_tcx.c`] is multiplied in the linear predictive domain, right before the LPC (Linear Predictive Coding) synthesis.

- 20 - If the last mode was TCX80 as well as the extrapolated mode of the (partially lost) superframe is TCX80 ($nloss = [1, 2]$, $mode = (3, 3, 3, 3, 3)$), concealment is performed in the FFT domain, utilizing phase and amplitude extrapolation, taking the last correctly received frame into account. The extrapolation approach of the phase information is not of any interest here (no relation to fading strategy) and therefore not described. For further details, see [3GP09a, section 6.5.1.2.4]. With respect to the amplitude modification of AMR-WB+, the approach performed for TCX concealment consists of the following steps [3GP09a, section 6.5.1.2.3]:

- 25 - The previous frame magnitude spectrum is computed:

30

$$oldA[k] = |\hat{oldX}[k]|$$

- The current frame magnitude spectrum is computed:

35

$$A[k] = |\hat{X}[k]|$$

- The gain difference of energy of non-lost spectral coefficients between the previous and the current frame is computed:

$$gain = \sqrt{\frac{\sum A[k]^2}{\sum oldA[k]^2}}$$

5

- The amplitude of the missing spectral coefficients is extrapolated using:

$$if(lost[k]) \quad A[k] = gain \cdot oldA[k]$$

- 10 - In every other case of a lost frame with $m_k = [2, 3]$, the TCX target (inverse FFT of decoded spectrum plus noise fill-in (using a noise level decoded from the bitstream)) is synthesized using all available info (including global TCX gain). No fade-out is applied in this case.

- 15 Regarding CNG in AMR-WB+, the same approach as in AMR-WB is used (see above).

In the following, OPUS is considered. OPUS [IET12] incorporates technology from two codecs: the speech-oriented SILK (known as the Skype codec) and the low-latency CELT (CELT = Constrained-Energy Lapped Transform). Opus can be adjusted seamlessly
20 between high and low bitrates, and internally, it switches between a linear prediction codec at lower bitrates (SILK) and a transform codec at higher bitrates (CELT) as well as a hybrid for a short overlap.

Regarding SILK audio data compression and decompression, in OPUS, there are several
25 parameters which are attenuated during concealment in the SILK decoder routine. The LTP gain parameter is attenuated by multiplying all LPC coefficients with either 0.99, 0.95 or 0.90 per frame, depending on the number of consecutive lost frames, where the excitation is built up using the last pitch cycle from the excitation of the previous frame. The pitch lag parameter is very slowly increased during consecutive losses. For single
30 losses it is kept constant compared to the last frame. Moreover, the excitation gain parameter is exponentially attenuated with 0.99^{lost_cnt} per frame, so that the excitation gain parameter is 0.99 for the first excitation gain parameter, so that the excitation gain parameter is 0.992 for the second excitation gain parameter, and so on. The excitation is generated using a random number generator which is generating white noise by variable
35 overflow. Furthermore, the LPC coefficients are extrapolated/averaged based on the last correctly received set of coefficients. After generating the attenuated excitation vector, the

concealed LPC coefficients are used in OPUS to synthesize the time domain output signal.

Now, in the context of OPUS, CELT is considered. CELT is a transform based codec. The concealment of CELT features a pitch based PLC approach, which is applied for up to five consecutively lost frames. Starting with frame 6, a noise like concealment approach is applied, which generating background noise, which characteristic is supposed to sound like preceding background noise.

Fig. 5 illustrates the burst loss behavior of CELT. In particular, Fig. 5 depicts a spectrogram (x-axis: time; y-axis: frequency) of a CELT concealed speech segment. The light grey box indicates the first 5 consecutively lost frames, where the pitch based PLC approach is applied. Beyond that, the noise like concealment is shown. It should be noted that the switching is performed instantly, it does not transit smoothly.

Regarding pitch based concealment, in OPUS, the pitch based concealment consists of finding the periodicity in the decoded signal by autocorrelation and repeating the windowed waveform (in the excitation domain using LPC analysis and synthesis) using the pitch offset (pitch lag). The windowed waveform is overlapped in such a way as to preserve the time-domain aliasing cancellation with the previous frame and the next frame [IET12]. Additionally a fade-out factor is derived and applied by the following code:

```
opus_val32 E1=1, E2=1;
int period;
25 if (pitch_index <= MAX_PERIOD/2) {
    period = pitch_index;
}
else {
    period = MAX_PERIOD/2;
30 }
for (i=0;i<period;i++)
{
    E1 += exc[MAX_PERIOD- period+i] * exc[MAX_PERIOD- period+i];
    E2 += exc[MAX_PERIOD-2*period+i] * exc[MAX_PERIOD-2*period+i];
35 }
if (E1 > E2) {
    E1 = E2;
}
```

```

decay = sqrt(E1/E2));
attenuation = decay;

```

5 In this code, exc contains the excitation signal up to MAX_PERIOD samples before the loss.

The excitation signal is later multiplied with attenuation, then synthesized and output via LPC synthesis.

10 The fading algorithm for the time domain approach can be summarized like this:

- Find the pitch synchronous energy of the last pitch cycle before the loss.
- Find the pitch synchronous energy of the second last pitch cycle before the loss.
- 15 - If the energy is increasing, limit it to stay constant: attenuation = 1
- If the energy is decreasing, continue with the same attenuation during concealment.

20

Regarding noise like concealment, according to OPUS, for the 6th and following consecutive lost frames a noise substitution approach in the MDCT domain is performed, in order to simulate comfort background noise.

25 Regarding tracing of the background noise level and shape, in OPUS, the background noise estimate is performed as follows: After the MDCT analysis, the square root of the MDCT band energies is calculated per frequency band, where the grouping of the MDCT bins follows the bark scale according to [IET12, Table 55]. Then the square root of the energies is transformed into the $\log_2$ domain by:

30

$$bandLogE[i] = \log_2(e) \cdot \log_e(bandE[i] - eMeans[i]) \quad \text{for } i = 0 \dots 21 \quad (18)$$

wherein e is the Euler's number, $bandE$ is the square root of the MDCT band and $eMeans$ is a vector of constants (necessary to get the result zero mean, which results in an enhanced coding gain).

35

In OPUS, the background noise is logged on the decoder side like this [IET12, amp2Log2 and log2Amp @ quant_bands.c]:

$$\begin{aligned} \text{backgroundLogE}[i] &= \min(\text{backgroundLogE}[i] + 8 \cdot 0.001, \text{bandLogE}[i]) \\ &\quad \text{for } i = 0 \dots 21 \end{aligned} \quad (19)$$

The traced minimum energy is basically determined by the square root of the energy of the band of the current frame, but the increase from one frame to the next is limited by 0.05 dB.

10

Regarding the application of the background noise level and shape, according to OPUS, if the noise like PLC is applied, *backgroundLogE* as derived in the last good frame is used and converted back to the linear domain:

$$\begin{aligned} \text{bandE}[i] &= e^{(\log_e(2) \cdot (\text{backgroundLogE}[i] + eMeans[i]))} \quad \text{for } i = 0 \dots 21 \\ &\quad (20) \end{aligned}$$

where e is the Euler's number and *eMeans* is the same vector of constants as for the "linear to log" transform.

20

The current concealment procedure is to fill the MDCT frame with white noise produced by a random number generator, and scale this white noise in a way that it matches band wise to the energy of *bandE*. Subsequently, the inverse MDCT is applied which results in a time domain signal. After the overlap add and deemphasis (like in regular decoding) it is put out.

25

In the following, MPEG-4 HE-AAC is considered (MPEG = Moving Picture Experts Group; HE-AAC = High Efficiency Advanced Audio Coding). High Efficiency Advanced Audio Coding consists of a transform based audio codec (AAC), supplemented by a parametric bandwidth extension (SBR).

30

Regarding AAC (AAC = Advanced Audio Coding), the DAB consortium specifies for AAC in DAB+, a fade-out to zero in the frequency domain [EBU10, section A1.2] (DAB = Digital Audio Broadcasting). Fade-out behavior, e.g., the attenuation ramp, might be fixed or adjustable by the user. The spectral coefficients from the last AU (AU = Access Unit) are attenuated by a factor corresponding to the fade-out characteristics and then passed to the frequency-to-time mapping. Depending on the attenuation ramp, the concealment

35

switches to muting after a number of consecutive invalid AUs, which means the complete spectrum will be set to 0.

- The DRM (DRM = Digital Rights Management) consortium specifies for AAC in DRM a
 5 fade-out in the frequency domain [EBU12, section 5.3.3]. Concealment works on the spectral data just before the final frequency to time conversion. If multiple frames are corrupted, concealment implements first a fadeout based on slightly modified spectral values from the last valid frame. Moreover, similar to DAB+, fade-out behavior, e.g., the attenuation ramp, might be fixed or adjustable by the user. The spectral coefficients from
 10 the last frame are attenuated by a factor corresponding to the fade-out characteristics and then passed to the frequency to-time mapping. Depending on the attenuation ramp, the concealment switches to muting after a number of consecutive invalid frames, which means the complete spectrum will be set to 0.
- 15 3GPP introduces for AAC in Enhanced aacPlus the fade-out in the frequency domain similar to DRM [3GP12e, section 5.1]. Concealment works on the spectral data just before the final frequency to time conversion. If multiple frames are corrupted, concealment implements first a fadeout based on slightly modified spectral values from the last good frame. A complete fading out takes 5 frames. The spectral coefficients from the last good
 20 frame are copied and attenuated by a factor of:

$$fadeOutFac = 2^{-(nFadeOutFrame/2)}$$

- with $nFadeOutFrame$ as frame counter since the last good frame. After five frames of
 25 fading out the concealment switches to muting, that means the complete spectrum will be set to 0.

- Lauber and Sperschneider introduce for AAC a frame-wise fade-out of the MDCT spectrum, based on energy extrapolation [LS01, section 4.4]. Energy shapes of a
 30 preceding spectrum might be used to extrapolate the shape of an estimated spectrum. Energy extrapolation can be performed independent of the concealment techniques as a kind of post concealment.

- Regarding AAC, the energy calculation is performed on a scale factor band basis in order
 35 to be close to the critical bands of the human auditory system. The individual energy values are decreased on a frame by frame basis in order to reduce the volume smoothly,

e.g., to fade out the signal. This becomes necessary since the probability, that the estimated values represent the current signal, decreases rapidly over time.

5 For the generation of the spectrum to be fed out they suggest frame repetition or noise substitution [LS01, sections 3.2 and 3.3].

10 Quackenbusch and Driesen suggest for AAC an exponential frame-wise fade-out to zero [QD03]. A repetition of adjacent set of time/frequency coefficients is proposed, wherein each repetition has exponentially increasing attenuation, thus fading gradually to mute in the case of extended outages.

15 Regarding SBR (SBR = Spectral Band Replication) in MPEG-4 HE-AAC, 3GPP suggests for SBR in Enhanced aacPlus to buffer the decoded envelope data and, in case of a frame loss, to reuse the buffered energies of the transmitted envelope data and to decrease them by a constant ratio of 3 dB for every concealed frame. The result is fed into the normal decoding process where the envelope adjuster uses it to calculate the gains, used for adjusting the patched highbands created by the HF generator. SBR decoding then takes place as usual. Moreover, the delta coded noise floor and sine level values are being deleted. As no difference to the previous information remains available, the
20 decoded noise floor and sine levels remain proportional to the energy of the HF generated signal [3GP12e, section 5.2].

25 The DRM consortium specified for SBR in conjunction with AAC the same technique as 3GPP [EBU12, section 5.6.3.1]. Moreover, The DAB consortium specifies for SBR in DAB+ the same technique as 3GPP [EBU10, section A2].

30 In the following, MPEG-4 CELP and MPEG-4 HVXC (HVXC = Harmonic Vector Excitation Coding) are considered. The DRM consortium specifies for SBR in conjunction with CELP and HVXC [EBU12, section 5.6.3.2] that the minimum requirement concealment for SBR for the speech codecs is to apply a predetermined set of data values, whenever a corrupted SBR frame has been detected. Those values yield a static highband spectral envelope at a low relative playback level, exhibiting a roll-off towards the higher frequencies. The objective is simply to ensure that no ill-behaved, potentially loud, audio bursts reach the listener's ears, by means of inserting "comfort noise" (as opposed to strict
35 muting). This is in fact no real fade-out but rather a jump to a certain energy level in order to insert some kind of comfort noise.

Subsequently, an alternative is mentioned [EBU12, section 5.6.3.2] which reuses the last correctly decoded data and slowly fading the levels (L) towards 0, analogously to the AAC + SBR case.

- 5 Now, MPEG-4 HILN is considered (HILN = Harmonic and Individual Lines plus Noise). Meine et al. introduce a fade-out for the parametric MPEG-4 HILN codec [ISO09] in a parametric domain [MEP01]. For continued harmonic components a good default behavior for replacing corrupted differentially encoded parameters is to keep the frequency constant, to reduce the amplitude by an attenuation factor (e.g., -6 dB), and to let the spectral envelope converge towards that of the averaged low-pass characteristic. An alternative for the spectral envelope would be to keep it unchanged. With respect to amplitudes and spectral envelopes, noise components can be treated the same way as harmonic components.
- 10
- 15 In the following, tracing of the background noise level in the prior art is considered. Rangachari and Loizou [RL06] provide a good overview of several methods and discuss some of their limitations. Methods for tracing the background noise level are, e.g., minimum tracking procedure [RL06] [Coh03] [SFB00] [Dob95], VAD based (VAD = voice activity detection); Kalman filtering [Gan05] [BJH06], subspace decompositions [BP06] [HJH08]; Soft Decision [SS98] [MPC89] [HE95], and minimum statistics.
- 20

The minimum statistics approach was chosen to be used within the scope for USAC-2, (USAC = Unified Speech and Audio Coding) and is subsequently outlined in more detail.

- 25 Noise power spectral density estimation based on optimal smoothing and minimum statistics [Mar01] introduces a noise estimator, which is capable of working independently of the signal being active speech or background noise. In contrast to other methods, the minimum statistics algorithm does not use any explicit threshold to distinguish between speech activity and speech pause and is therefore more closely related to soft-decision methods than to the traditional voice activity detection methods. Similar to soft-decision methods, it can also update the estimated noise PSD (Power Spectral Density) during speech activity.
- 30

- 35 The minimum statistics method rests on two observations namely that the speech and the noise are usually statistically independent and that the power of a noisy speech signal frequently decays to the power level of the noise. It is therefore possible to derive an accurate noise PSD (PSD = power spectral density) estimate by tracking the minimum of

the noisy signal PSD. Since the minimum is smaller than (or in other cases equal to) the average value, the minimum tracking method requires a bias compensation.

5 The bias is a function of the variance of the smoothed signal PSD and as such depends on the smoothing parameter of the PSD estimator. In contrast to earlier work on minimum tracking, which utilizes a constant smoothing parameter and a constant minimum bias correction, a time and frequency dependent PSD smoothing is used, which also requires a time and frequency dependent bias compensation.

10 Using minimum tracking provides a rough estimate of the noise power. However, there are some shortcomings. The smoothing with a fixed smoothing parameter widens the peaks of speech activity of the smoothed PSD estimate. This will lead to inaccurate noise estimates as the sliding window for the minimum search might slip into broad peaks. Thus, smoothing parameters close to one cannot be used, and, as a consequence, the
15 noise estimate will have a relatively large variance. Moreover, the noise estimate is biased toward lower values. Furthermore, in case of increasing noise power, the minimum tracking lags behind.

MMSE based noise PSD tracking with low complexity [HHJ10] introduces a background
20 noise PSD approach utilizing an MMSE search used on a DFT (Discrete Fourier Transform) spectrum. The algorithm consists of these processing steps:

- The maximum likelihood estimator is computed based on the noise PSD of the previous frame.
25
- The minimum mean square estimator is computed.
- The maximum likelihood estimator is estimated using the decision-directed approach [EM84].
30
- The inverse bias factor is computed assuming that speech and noise DFT coefficients are Gaussian distributed.
- The estimated noise power spectral density is smoothed.
35

There is also a safety-net approach applied in order to avoid a complete dead lock of the algorithm.

Tracking of non-stationary noise based on data-driven recursive noise power estimation [EH08] introduces a method for the estimation of the noise spectral variance from speech signals contaminated by highly non-stationary noise sources. This method is also using smoothing in time/frequency direction.

5

A low-complexity noise estimation algorithm based on smoothing of noise power estimation and estimation bias correction [Yu09] enhances the approach introduced in [EH08]. The main difference is, that the spectral gain function for noise power estimation is found by an iterative data-driven method.

10

Statistical methods for the enhancement of noisy speech [Mar03] combine the minimum statistics approach given in [Mar01] by soft-decision gain modification [MCA99], by an estimation of the a-priori SNR [MCA99], by an adaptive gain limiting [MC99] and by a MMSE log spectral amplitude estimator [EM85].

15

Fade out is of particular interest for a plurality of speech and audio codecs, in particular, AMR (see [3GP12b]) (including ACELP and CNG), AMR-WB (see [3GP09c]) (including ACELP and CNG), AMR-WB+ (see [3GP09a]) (including ACELP, TCX and CNG), G.718 (see [ITU08a]), G.719 (see [ITU08b]), G.722 (see [ITU07]), G.722.1 (see [ITU05]), G.729 (see [ITU12, CPK08, PKJ+11]), MPEG-4 HE-AAC / Enhanced aacPlus (see [EBU10, EBU12, 3GP12e, LS01, QD03]) (including AAC and SBR), MPEG-4 HILN (see [ISO09, MEP01]) and OPUS (see [IET12]) (including SILK and CELT).

20

Depending on the codec, fade-out is performed in different domains:

25

For codecs that utilize LPC, the fade-out is performed in the linear predictive domain (also known as the excitation domain). This holds true for codecs which are based on ACELP, e.g., AMR, AMR-WB, the ACELP core of AMR-WB+, G.718, G.729, G.729.1, the SILK core in OPUS; codecs which further process the excitation signal using a time-frequency transformation, e.g., the TCX core of AMR-WB+, the CELT core in OPUS; and for comfort noise generation (CNG) schemes, that operate in the linear predictive domain, e.g., CNG in AMR, CNG in AMR-WB, CNG in AMR-WB+.

30

For codecs that directly transform the time signal into the frequency domain, the fade-out is performed in the spectral / subband domain. This holds true for codecs which are based on MDCT or a similar transformation, such as AAC in MPEG-4 HE-AAC, G.719, G.722 (subband domain) and G.722.1.

35

For parametric codecs, fade-out is applied in the parametric domain. This holds true for MPEG-4 HILN.

Regarding fade-out speed and fade-out curve, a fade-out is commonly realized by the application of an attenuation factor, which is applied to the signal representation in the appropriate domain. The size of the attenuation factor controls the fade-out speed and the fade-out curve. In most cases the attenuation factor is applied frame wise, but also a sample wise application is utilized see, e.g., G.718 and G.722.

The attenuation factor for a certain signal segment might be provided in two manners, absolute and relative.

In the case where an attenuation factor is provided absolutely, the reference level is always the one of the last received frame. Absolute attenuation factors usually start with a value close to 1 for the signal segment immediately after the last good frame and then degrade faster or slower towards 0. The fade-out curve directly depends on these factors. This is, e.g., the case for the concealment described in Appendix IV of G.722 (see, in particular, [ITU07, figure IV.7]), where the possible fade-out curves are linear or gradually linear. Considering a gain factor $g(n)$, whereas $g(0)$ represents the gain factor of the last good frame, an absolute attenuation factor $\alpha_{abs}(n)$, the gain factor of any subsequent lost frame can be derived as

$$g(n) = \alpha_{abs}(n) \cdot g(0) \quad (21)$$

In the case where an attenuation factor is provided relatively, the reference level is the one from the previous frame. This has advantages in the case of a recursive concealment procedure, e.g., if the already attenuated signal is further processed and attenuated again.

If an attenuation factor is recursively applied, then this might be a fixed value independent of the number of consecutively lost frames, e.g., 0.5 for G.719 (see above); a fixed value relative to the number of consecutively lost frames, e.g., as proposed for G.729 in [CPK08]: 1.0 for the first two frames, 0.9 for the next two frames, 0.8 for the frames 5 and 6, and 0 for all subsequent frames (see above); or a value which is relative to the number of consecutively lost frames and which depends on signal characteristics, e.g., a faster fade-out for an instable signal and a slower fade-out for a stable signal, e.g., G.718 (see section above and [ITU08a, table 44]);

Assuming a relative fade-out factor $0 \leq \alpha_{rel}(n) \leq 1$, whereas n is the number of the lost frame ($n \geq 1$); the gain factor of any subsequent frame can be derived as

$$g(n) = \alpha_{rel}(n) \cdot g(n-1) \quad (22)$$

5

$$g(n) = \left(\prod_{m=1}^n \alpha(m) \right) \cdot g(0) \quad (23)$$

$$g(n) = \alpha_{rel}^n \cdot g(0) \quad (24)$$

10 resulting in an exponential fading.

Regarding the fade-out procedure, usually, the attenuation factor is specified, but in some application standards (DRM, DAB+) the latter is left to the manufacturer.

15 If different signal parts are faded separately, different attenuation factors might be applied, e.g., to fade tonal components with a certain speed and noise-like components with another speed (e.g., AMR, SILK).

20 Usually, a certain gain is applied to the whole frame. When the fading is performed in the spectral domain, this is the only way possible. However, if the fading is done in the time domain or the linear predictive domain, a more granular fading is possible. Such more granular fading is applied in G.718, where individual gain factors are derived for each sample by linear interpolation between the gain factor of the last frame and the gain factor of the current frame.

25

For codecs with a variable frame duration, a constant, relative attenuation factor leads to a different fade-out speed depending on the frame duration. This is, e.g., the case for AAC, where the frame duration depends on the sampling rate.

30 To adopt the applied fading curve to the temporal shape of the last received signal, the (static) fade-out factors might be further adjusted. Such further dynamic adjustment is, e.g., applied for AMR where the median of the previous five gain factors is taken into account (see [3GP12b] and section 1.8.1). Before any attenuation is performed, the current gain is set to the median, if the median is smaller than the last gain, otherwise the
35 last gain is used. Moreover, such further dynamic adjustment is, e.g., applied for G729, where the amplitude is predicted using linear regression of the previous gain factors (see

[CPK08, PKJ+11] and section 1.6). In this case, the resulting gain factor for the first concealed frames might exceed the gain factor of the last received frame.

Regarding the target level of the fade-out, with the exception of G.718 and CELT, the
5 target level is 0 for all analyzed codecs, including those codecs' comfort noise generation (CNG).

In G.718, fading of the pitch excitation (representing tonal components) and fading of the random excitation (representing noise-like components) is performed separately. While
10 the pitch gain factor is faded to zero, the innovation gain factor is faded to the CNG excitation energy.

Assuming that relative attenuation factors are given, this leads - based on formula (23) - to the following absolute attenuation factor:

15

$$g(n) = \alpha_{rel}(n) \cdot g(n-1) + (1 - \alpha_{rel}(n)) \cdot g_n \quad (25)$$

with g_n being the gain of the excitation used during the comfort noise generation. This formula corresponds to formula (23), when $g_n = 0$.

20

G.718 performs no fade-out in the case of DTX/CNG.

In CELT there is no fading towards the target level, but after 5 frames of tonal concealment (including a fade-out) the level is instantly switched to the target level at the
25 6th consecutively lost frame. The level is derived band wise using formula (19).

Regarding the target spectral shape of the fade-out, all analyzed pure transform based codecs (AAC, G.719, G.722, G.722.1) as well as SBR simply prolong the spectral shape of the last good frame during the fade-out.

30

Various speech codecs fade the spectral shape to a mean using the LPC synthesis. The mean might be static (AMR) or adaptive (AMR-WB, AMR-WB+, G.718), whereas the latter is derived from a static mean and a short term mean (derived by averaging the last n LP coefficient sets) (LP = Linear Prediction).

35

All CNG modules in the discussed codecs AMR, AMR-WB, AMR-WB+, G.718 prolong the spectral shape of the last good frame during the fade-out.

Regarding background noise level tracing, there are five different approaches known from the literature:

- 5 - Voice Activity Detector based: based on SNR/VAD, but very difficult to tune and hard to use for low SNR speech.
- Soft-decision scheme: The soft-decision approach takes the probability of speech presence into account [SS98] [MPC89] [HE95].
- 10 - Minimum statistics: The minimum of the PSD is tracked holding a certain amount of values over time in a buffer, thus enabling to find the minimal noise from the past samples [Mar01] [HHJ10] [EH08] [Yu09].
- 15 - Kalman Filtering: The algorithm uses a series of measurements observed over time, containing noise (random variations), and produces estimates of the noise PSD that tend to be more precise than those based on a single measurement alone. The Kalman filter operates recursively on streams of noisy input data to produce a statistically optimal estimate of the system state [Gan05] [BJH06].
- 20 - Subspace Decomposition: This approach tries to decompose a noise like signal into a clean speech signal and a noise part, utilizing for example the KLT (Karhunen-Loève transform, also known as principal component analysis) and/or the DFT (Discrete Time Fourier Transform). Then the eigenvectors/eigenvalues
- 25 can be traced using an arbitrary smoothing algorithm [BP06] [HJH08].

EP 2 026 330 A1 discloses a device and a method for frame lost concealment. A pitch period of a current lost frame is obtained on the basis of a pitch period of the last good frame before the current lost frame. An excitation signal of the current lost frame is recovered on the basis of the pitch period of the current lost frame and an excitation signal of the last good frame before the lost frame. Thereby, the hearing contrast of a receiver is reduced, and the quality of speech is improved. Further, in EP 2 026 330 A1, a pitch period of continual lost frames is adjusted on the basis of the change trend of the pitch period of the last good frame before the lost frame.

35

An apparatus for decoding an encoded audio signal to obtain a reconstructed audio signal is provided. The apparatus comprises a receiving interface for receiving one or more frames, a coefficient generator, and a signal reconstructor. The coefficient generator is
5 configured to determine, if a current frame of the one or more frames is received by the receiving interface and if the current frame being received by the receiving interface is not corrupted, one or more first audio signal coefficients, being comprised by the current frame, wherein said one or more first audio signal coefficients indicate a characteristic of the encoded audio signal, and one or more noise coefficients indicating a background

noise of the encoded audio signal. Moreover, the coefficient generator is configured to generate one or more second audio signal coefficients, depending on the one or more first audio signal coefficients and depending on the one or more noise coefficients, if the current frame is not received by the receiving interface or if the current frame being received by the receiving interface is corrupted. The audio signal reconstructor is configured to reconstruct a first portion of the reconstructed audio signal depending on the one or more first audio signal coefficients, if the current frame is received by the receiving interface and if the current frame being received by the receiving interface is not corrupted. Moreover, the audio signal reconstructor is configured to reconstruct a second portion of the reconstructed audio signal depending on the one or more second audio signal coefficients, if the current frame is not received by the receiving interface or if the current frame being received by the receiving interface is corrupted.

In some embodiments, the one or more first audio signal coefficients may, e.g., be one or more linear predictive filter coefficients of the encoded audio signal. In some embodiments, the one or more first audio signal coefficients may, e.g., be one or more linear predictive filter coefficients of the encoded audio signal.

According to an embodiment, the one or more noise coefficients may, e.g., be one or more linear predictive filter coefficients indicating the background noise of the encoded audio signal. In an embodiment, the one or more linear predictive filter coefficients may, e.g., represent a spectral shape of the background noise.

In an embodiment, the coefficient generator may, e.g., be configured to determine the one or more second audio signal portions such that the one or more second audio signal portions are one or more linear predictive filter coefficients of the reconstructed audio signal, or such that the one or more first audio signal coefficients are one or more immittance spectral pairs of the reconstructed audio signal.

According to an embodiment, the coefficient generator may, e.g., be configured to generate the one or more second audio signal coefficients by applying the formula:

$$f_{current}[i] = \alpha \cdot f_{last}[i] + (1 - \alpha) \cdot pt_{mean}[i]$$

wherein $f_{current}[i]$ indicates one of the one or more second audio signal coefficients, wherein $f_{last}[i]$ indicates one of the one or more first audio signal coefficients, wherein

$pt_{mean}[i]$ is one of the one or more noise coefficients, wherein α is a real number with $0 \leq \alpha \leq 1$, and wherein i is an index. In an embodiment, $0 < \alpha < 1$.

According to an embodiment, $f_{last}[i]$ indicates a linear predictive filter coefficient of the encoded audio signal, and wherein $f_{current}[i]$ indicates a linear predictive filter coefficient of the reconstructed audio signal.

In an embodiment, $pt_{mean}[i]$ may, e.g., indicate the background noise of the encoded audio signal.

10

In an embodiment, the coefficient generator may, e.g., be configured to determine, if the current frame of the one or more frames is received by the receiving interface and if the current frame being received by the receiving interface is not corrupted, the one or more noise coefficients by determining a noise spectrum of the encoded audio signal.

15

According to an embodiment, the coefficient generator may, e.g., be configured to determine LPC coefficients representing background noise by using a minimum statistics approach on the signal spectrum to determine a background noise spectrum and by calculating the LPC coefficients representing the background noise shape from the background noise spectrum.

20

Moreover, a method for decoding an encoded audio signal to obtain a reconstructed audio signal is provided. The method comprises:

25

- Receiving one or more frames.

30

- Determining, if a current frame of the one or more frames is received and if the current frame being received is not corrupted, one or more first audio signal coefficients, being comprised by the current frame, wherein said one or more first audio signal coefficients indicate a characteristic of the encoded audio signal, and one or more noise coefficients indicating a background noise of the encoded audio signal.

35

- Generating one or more second audio signal coefficients, depending on the one or more first audio signal coefficients and depending on the one or more noise coefficients, if the current frame is not received or if the current frame being received is corrupted.

- Reconstructing a first portion of the reconstructed audio signal depending on the one or more first audio signal coefficients, if the current frame is received and if the current frame being received is not corrupted. And:
- 5 - Reconstructing a second portion of the reconstructed audio signal depending on the one or more second audio signal coefficients, if the current frame is not received or if the current frame being received is corrupted.

Moreover, a computer program for implementing the above-described method when being
10 executed on a computer or signal processor is provided.

Having common means to trace and apply the spectral shape of comfort noise during fade out has several advantages. By tracing and applying the spectral shape such that it can be done similarly for both core codecs allows for a simple common approach. CELT
15 teaches only the band wise tracing of energies in the spectral domain and the band wise forming of the spectral shape in the spectral domain, which is not possible for the CELP core.

In contrast, in the prior art, the spectral shape of the comfort noise introduced during burst
20 losses is either fully static, or partly static and partly adaptive to the short term mean of the spectral shape (as realized in G.718 [ITU08a]), and will usually not match the background noise in the signal before the packet loss. This mismatch of the comfort noise characteristics might be disturbing. According to the prior art, an offline trained (static) background noise shape may be employed that may be sound pleasant for particular
25 signals, but less pleasant for others, e.g., car noise sounds totally different to office noise.

Moreover, in the prior art, an adaptation to the short term mean of the spectral shape of the previously received frames may be employed which might bring the signal characteristics closer to the signal received before, but not necessarily to the background
30 noise characteristics. In the prior art, tracing the spectral shape band wise in the spectral domain (as realized in CELT [IET12]) is not applicable for a switched codec using not only an MDCT domain based core (TCX) but also an ACELP based core. The above-mentioned embodiments are thus advantageous over the prior art.

35 Moreover, an apparatus for decoding an audio signal is provided.

The apparatus comprises a receiving interface. The receiving interface is configured to receive a plurality of frames, wherein the receiving interface is configured to receive a first

frame of the plurality of frames, said first frame comprising a first audio signal portion of the audio signal, said first audio signal portion being represented in a first domain, and wherein the receiving interface is configured to receive a second frame of the plurality of frames, said second frame comprising a second audio signal portion of the audio signal.

5

Moreover, the apparatus comprises a transform unit for transforming the second audio signal portion or a value or signal derived from the second audio signal portion from a second domain to a tracing domain to obtain a second signal portion information, wherein the second domain is different from the first domain, wherein the tracing domain is different from the second domain, and wherein the tracing domain is equal to or different from the first domain.

10

Furthermore, the apparatus comprises a noise level tracing unit, wherein the noise level tracing unit is configured to receive a first signal portion information being represented in the tracing domain, wherein the first signal portion information depends on the first audio signal portion. The noise level tracing unit is configured to receive the second signal portion being represented in the tracing domain, and wherein the noise level tracing unit is configured to determine noise level information depending on the first signal portion information being represented in the tracing domain and depending on the second signal portion information being represented in the tracing domain.

15

20

Moreover, the apparatus comprises a reconstruction unit for reconstructing a third audio signal portion of the audio signal depending on the noise level information, if a third frame of the plurality of frames is not received by the receiving interface but is corrupted.

25

An audio signal may, for example, be a speech signal, or a music signal, or signal that comprises speech and music, etc.

The statement that the first signal portion information depends on the first audio signal portion means that the first signal portion information either is the first audio signal portion, or that the first signal portion information has been obtained/generated depending on the first audio signal portion or in some other way depends on the first audio signal portion. For example, the first audio signal portion may have been transformed from one domain to another domain to obtain the first signal portion information.

30

35

Likewise, a statement that the second signal portion information depends on a second audio signal portion means that the second signal portion information either is the second audio signal portion, or that the second signal portion information has been

obtained/generated depending on the second audio signal portion or in some other way depends on the second audio signal portion. For example, the second audio signal portion may have been transformed from one domain to another domain to obtain second signal portion information.

5

In an embodiment, the first audio signal portion may, e.g., be represented in a time domain as the first domain. Moreover, transform unit may, e.g., be configured to transform the second audio signal portion or the value derived from the second audio signal portion from an excitation domain being the second domain to the time domain being the tracing domain. Furthermore, the noise level tracing unit may, e.g., be configured to receive the first signal portion information being represented in the time domain as the tracing domain. Moreover, the noise level tracing unit may, e.g., be configured to receive the second signal portion being represented in the time domain as the tracing domain.

15 According to an embodiment, the first audio signal portion may, e.g., be represented in an excitation domain as the first domain. Moreover, the transform unit may, e.g., be configured to transform the second audio signal portion or the value derived from the second audio signal portion from a time domain being the second domain to the excitation domain being the tracing domain. Furthermore, the noise level tracing unit may, e.g., be configured to receive the first signal portion information being represented in the excitation domain as the tracing domain. Moreover, the noise level tracing unit may, e.g., be configured to receive the second signal portion being represented in the excitation domain as the tracing domain.

25 In an embodiment, the first audio signal portion may, e.g., be represented in an excitation domain as the first domain, wherein the noise level tracing unit may, e.g., be configured to receive the first signal portion information, wherein said first signal portion information is represented in the FFT domain, being the tracing domain, and wherein said first signal portion information depends on said first audio signal portion being represented in the excitation domain, wherein the transform unit may, e.g., be configured to transform the second audio signal portion or the value derived from the second audio signal portion from a time domain being the second domain to an FFT domain being the tracing domain, and wherein the noise level tracing unit may, e.g., be configured to receive the second audio signal portion being represented in the FFT domain.

35

In an embodiment, the apparatus may, e.g., further comprise a first aggregation unit for determining a first aggregated value depending on the first audio signal portion. Moreover, the apparatus may, e.g., further comprise a second aggregation unit for determining,

depending on the second audio signal portion, a second aggregated value as the value derived from the second audio signal portion. Furthermore, the noise level tracing unit may, e.g., be configured to receive the first aggregated value as the first signal portion information being represented in the tracing domain, wherein the noise level tracing unit may, e.g., be configured to receive the second aggregated value as the second signal portion information being represented in the tracing domain, and wherein the noise level tracing unit may, e.g., be configured to determine noise level information depending on the first aggregated value being represented in the tracing domain and depending on the second aggregated value being represented in the tracing domain.

10

According to an embodiment, the first aggregation unit may, e.g., be configured to determine the first aggregated value such that the first aggregated value indicates a root mean square of the first audio signal portion or of a signal derived from the first audio signal portion. Moreover, the second aggregation unit may, e.g., be configured to determine the second aggregated value such that the second aggregated value indicates a root mean square of the second audio signal portion or of a signal derived from the second audio signal portion.

15

In an embodiment, the transform unit may, e.g., be configured to transform the value derived from the second audio signal portion from the second domain to the tracing domain by applying a gain value on the value derived from the second audio signal portion.

20

According to embodiments, the gain value may, e.g., indicate a gain introduced by Linear predictive coding synthesis, or the gain value may, e.g., indicate a gain introduced by Linear predictive coding synthesis and deemphasis.

25

In an embodiment, the noise level tracing unit may, e.g., be configured to determine noise level information by applying a minimum statistics approach.

30

According to an embodiment, the noise level tracing unit may, e.g., be configured to determine a comfort noise level as the noise level information. The reconstruction unit may, e.g., be configured to reconstruct the third audio signal portion depending on the noise level information, if said third frame of the plurality of frames is not received by the receiving interface or if said third frame is received by the receiving interface but is corrupted.

35

In an embodiment, the noise level tracing unit may, e.g., be configured to determine a comfort noise level as the noise level information derived from a noise level spectrum, wherein said noise level spectrum is obtained by applying the minimum statistics approach. The reconstruction unit may, e.g., be configured to reconstruct the third audio
5 signal portion depending on a plurality of Linear Predictive coefficients, if said third frame of the plurality of frames is not received by the receiving interface or if said third frame is received by the receiving interface but is corrupted.

According to another embodiment, the noise level tracing unit may, e.g., be configured to
10 determine a plurality of Linear Predictive coefficients indicating a comfort noise level as the noise level information, and the reconstruction unit may, e.g., be configured to reconstruct the third audio signal portion depending on the plurality of Linear Predictive coefficients.

15 In an embodiment, the noise level tracing unit is configured to determine a plurality of FFT coefficients indicating a comfort noise level as the noise level information, and the first reconstruction unit is configured to reconstruct the third audio signal portion depending on a comfort noise level derived from said FFT coefficients, if said third frame of the plurality of frames is not received by the receiving interface or if said third frame is received by the
20 receiving interface but is corrupted.

In an embodiment, the reconstruction unit may, e.g., be configured to reconstruct the third audio signal portion depending on the noise level information and depending on the first audio signal portion, if said third frame of the plurality of frames is not received by the
25 receiving interface or if said third frame is received by the receiving interface but is corrupted.

According to an embodiment, the reconstruction unit may, e.g., be configured to reconstruct the third audio signal portion by attenuating or amplifying a signal derived from
30 the first or the second audio signal portion.

In an embodiment, the apparatus may, e.g., further comprise a long-term prediction unit comprising a delay buffer. Moreover, the long-term prediction unit may, e.g., be configured to generate a processed signal depending on the first or the second audio signal portion,
35 depending on a delay buffer input being stored in the delay buffer and depending on a long-term prediction gain. Furthermore, the long-term prediction unit may, e.g., be configured to fade the long-term prediction gain towards zero , if said third frame of the

plurality of frames is not received by the receiving interface or if said third frame is received by the receiving interface but is corrupted.

According to an embodiment, the long-term prediction unit may, e.g., be configured to
5 fade the long-term prediction gain towards zero, wherein a speed with which the long-term prediction gain is faded to zero depends on a fade-out factor.

In an embodiment, the long-term prediction unit may, e.g., be configured to update the delay buffer input by storing the generated processed signal in the delay buffer, if said
10 third frame of the plurality of frames is not received by the receiving interface or if said third frame is received by the receiving interface but is corrupted.

According to an embodiment, the transform unit may, e.g., be a first transform unit, and the reconstruction unit is a first reconstruction unit. The apparatus further comprises a
15 second transform unit and a second reconstruction unit. The second transform unit may, e.g., be configured to transform the noise level information from the tracing domain to the second domain, if a fourth frame of the plurality of frames is not received by the receiving interface or if said fourth frame is received by the receiving interface but is corrupted. Moreover, the second reconstruction unit may, e.g., be configured to reconstruct a fourth
20 audio signal portion of the audio signal depending on the noise level information being represented in the second domain if said fourth frame of the plurality of frames is not received by the receiving interface or if said fourth frame is received by the receiving interface but is corrupted.

25 In an embodiment, the second reconstruction unit may, e.g., be configured to reconstruct the fourth audio signal portion depending on the noise level information and depending on the second audio signal portion.

According to an embodiment, the second reconstruction unit may, e.g., be configured to
30 reconstruct the fourth audio signal portion by attenuating or amplifying a signal derived from the first or the second audio signal portion.

Moreover, a method for decoding an audio signal is provided.

35 The method comprises:

- Receiving a first frame of a plurality of frames, said first frame comprising a first audio signal portion of the audio signal, said first audio signal portion being represented in a first domain.
- 5 - Receiving a second frame of the plurality of frames, said second frame comprising a second audio signal portion of the audio signal.
- Transforming the second audio signal portion or a value or signal derived from the second audio signal portion from a second domain to a tracing domain to obtain a second signal portion information, wherein the second domain is different from the first domain, wherein the tracing domain is different from the second domain, and
10 wherein the tracing domain is equal to or different from the first domain.
- Determining noise level information depending on first signal portion information, being represented in the tracing domain, and depending on the second signal portion information being represented in the tracing domain, wherein the first signal portion information depends on the first audio signal portion. And:
15
- Reconstructing a third audio signal portion of the audio signal depending on the noise level information being represented in the tracing domain, if a third frame of the plurality of frames is not received or if said third frame is received but is corrupted.
20

Furthermore, a computer program for implementing the above-described method when
25 being executed on a computer or signal processor is provided.

Some of embodiments of the present invention provide a time varying smoothing parameter such that the tracking capabilities of the smoothed periodogram and its variance are better balanced, to develop an algorithm for bias compensation, and to
30 speed up the noise tracking in general.

Embodiments of the present invention are based on the finding that with regard to the fade-out, the following parameters are of interest: The fade-out domain; the fade-out speed, or, more general, fade-out curve; the target level of the fade-out; the target
35 spectral shape of the fade-out; and/or the background noise level tracing. In this context, embodiments are based on the finding that the prior art has significant drawbacks.

An apparatus and method for improved signal fade out for switched audio coding systems during error concealment is provided.

Moreover, a computer program for implementing the above-described method when being
5 executed on a computer or signal processor is provided.

Embodiments realize a fade-out to comfort noise level. According to embodiments, a common comfort noise level tracing in the excitation domain is realized. The comfort noise level being targeted during burst packet loss will be the same, regardless of the core
10 coder (ACELP/TCX) in use, and it will always be up to date. There is no prior art known, where a common noise level tracing is necessary. Embodiments provide the fading of a switched codec to a comfort noise like signal during burst packet losses.

Moreover, embodiments realize that the overall complexity will be lower compared to
15 having two independent noise level tracing modules, since functions (PROM) and memory can be shared.

In embodiments, the level derivation in the excitation domain (compared to the level derivation in the time domain) provides more minima during active speech, since part of
20 the speech information is covered by the LP coefficients.

In the case of ACELP, according to embodiments, the level derivation takes place in the excitation domain. In the case of TCX, in embodiments, the level is derived in the time domain, and the gain of the LPC synthesis and de-emphasis is applied as a correction
25 factor in order to model the energy level in the excitation domain. Tracing the level in the excitation domain, e.g., before the FDNS, would theoretically also be possible, but the level compensation between the TCX excitation domain and the ACELP excitation domain is deemed to be rather complex.

30 No prior art incorporates such a common background level tracing in different domains. The prior art techniques do not have such a common comfort noise level tracing, e.g., in the excitation domain, in a switched codec system. Thus, embodiments are advantageous over the prior art, as for the prior art techniques, the comfort noise level that is targeted during burst packet losses may be different, depending on the preceding coding mode
35 (ACELP/TCX), where the level was traced; as in the prior art, tracing which is separate for each coding mode will cause unnecessary overhead and additional computational complexity; and as in the prior art, no up-to-date comfort noise level might be available in either core due to recent switching to this core.

According to some embodiments, level tracing is conducted in the excitation domain, but TCX fade-out is conducted in the time domain. By fading in the time domain, failures of the TDAC are avoided, which would cause aliasing. This becomes of particular interest
5 when tonal signal components are concealed. Moreover, level conversion between the ACELP excitation domain and the MDCT spectral domain is avoided and thus, e.g., computation resources are saved. Because of switching between the excitation domain and the time domain, a level adjustment is required between the excitation domain and the time domain. This is resolved by the derivation of the gain that would be introduced by
10 the LPC synthesis and the preemphasis and to use this gain as a correction factor to convert the level between the two domains.

In contrast, prior art techniques do not conduct level tracing in the excitation domain and TCX Fade-Out in the Time Domain. Regarding state of the art transform based codecs,
15 the attenuation factor is applied either in the excitation domain (for time-domain/ACELP like concealment approaches, see [3GP09a]) or in the frequency domain (for frequency domain approaches like frame repetition or noise substitution, see [LS01]). A drawback of the approach of the prior art to apply the attenuation factor in the frequency domain is that aliasing will be caused in the overlap-add region in the time domain. This will be the case
20 for adjacent frames to which different attenuation factors are applied, because the fading procedure causes the TDAC (time domain alias cancellation) to fail. This is particularly relevant when tonal signal components are concealed. The above-mentioned embodiments are thus advantageous over the prior art.

25 Embodiments compensate the influence of the high pass filter on the LPC synthesis gain. According to embodiments, to compensate for the unwanted gain change of the LPC analysis and emphasis caused by the high pass filtered unvoiced excitation, a correction factor is derived. This correction factor takes this unwanted gain change into account and modifies the target comfort noise level in the excitation domain such that the correct target
30 level is reached in the time domain.

In contrast, the prior art, for example, G.718 [ITU08a], introduces a high pass filter into the signal path of the unvoiced excitation, as depicted in Fig. 2, if the signal of the last good frame was not classified as UNVOICED. By this, the prior art techniques cause unwanted
35 side effects, since the gain of the subsequent LPC synthesis depends on the signal characteristics, which are altered by this high pass filter. Since the background level is traced and applied in the excitation domain, the algorithm relies on the LPC synthesis gain, which in return again depends on the characteristics of the excitation signal. In other

words: The modification of the signal characteristics of the excitation due to the high pass filtering, as conducted by the prior art, might lead to a modified (usually reduced) gain of the LPC synthesis. This leads to a wrong output level even though the excitation level is correct.

5

Embodiments overcome these disadvantages of the prior art.

In particular, embodiments realize an adaptive spectral shape of comfort noise. In contrast to G.718, by tracing the spectral shape of the background noise, and by applying (fading to) this shape during burst packet losses, the noise characteristic of preceding background noise will be matched, leading to a pleasant noise characteristic of the comfort noise. This avoids obtrusive mismatches of the spectral shape that may be introduced by using a spectral envelope which was derived by offline training and/or the spectral shape of the last received frames.

15

Moreover, an apparatus for decoding an audio signal is provided. The apparatus comprises a receiving interface , wherein the receiving interface is configured to receive a first frame comprising a first audio signal portion of the audio signal, and wherein the receiving interface is configured to receive a second frame comprising a second audio signal portion of the audio signal.

20

Moreover, the apparatus comprises a noise level tracing unit, wherein the noise level tracing unit is configured to determine noise level information depending on at least one of the first audio signal portion and the second audio signal portion (this means: depending on the first audio signal portion and/or the second audio signal portion), wherein the noise level information is represented in a tracing domain.

25

Furthermore, the apparatus comprises a first reconstruction unit for reconstructing, in a first reconstruction domain, a third audio signal portion of the audio signal depending on the noise level information, if a third frame of the plurality of frames is not received by the receiving interface or if said third frame is received by the receiving interface but is corrupted, wherein the first reconstruction domain is different from or equal to the tracing domain.

30

Moreover, the apparatus comprises a transform unit for transforming the noise level information from the tracing domain to a second reconstruction domain, if a fourth frame of the plurality of frames is not received by the receiving interface or if said fourth frame is received by the receiving interface but is corrupted, wherein the second reconstruction

35

domain is different from the tracing domain, and wherein the second reconstruction domain is different from the first reconstruction domain, and

Furthermore, the apparatus comprises a second reconstruction unit for reconstructing, in
 5 the second reconstruction domain, a fourth audio signal portion of the audio signal depending on the noise level information being represented in the second reconstruction domain, if said fourth frame of the plurality of frames is not received by the receiving interface or if said fourth frame is received by the receiving interface but is corrupted.

10 According to some embodiments, the tracing domain may, e.g., be wherein the tracing domain is a time domain, a spectral domain, an FFT domain, an MDCT domain, or an excitation domain. The first reconstruction domain may, e.g., be the time domain, the spectral domain, the FFT domain, the MDCT domain, or the excitation domain. The
 15 second reconstruction domain may, e.g., be the time domain, the spectral domain, the FFT domain, the MDCT domain, or the excitation domain.

In an embodiment, the tracing domain may, e.g., be the FFT domain, the first reconstruction domain may, e.g., be the time domain, and the second reconstruction domain may, e.g., be the excitation domain.

20

In another embodiment, the tracing domain may, e.g., be the time domain, the first reconstruction domain may, e.g., be the time domain, and the second reconstruction domain may, e.g., be the excitation domain.

25 According to an embodiment, said first audio signal portion may, e.g., be represented in a first input domain, and said second audio signal portion may, e.g., be represented in a second input domain. The transform unit may, e.g., be a second transform unit. The apparatus may, e.g., further comprise a first transform unit for transforming the second audio signal portion or a value or signal derived from the second audio signal portion from
 30 the second input domain to the tracing domain to obtain a second signal portion information. The noise level tracing unit may, e.g., be configured to receive a first signal portion information being represented in the tracing domain, wherein the first signal portion information depends on the first audio signal portion, wherein the noise level tracing unit is configured to receive the second signal portion being represented in the
 35 tracing domain, and wherein the noise level tracing unit is configured to determine the noise level information depending on the first signal portion information being represented in the tracing domain and depending on the second signal portion information being represented in the tracing domain.

According to an embodiment, the first input domain may, e.g., be the excitation domain, and the second input domain may, e.g., be the MDCT domain.

- 5 In another embodiment, the first input domain may, e.g., be the MDCT domain, and wherein the second input domain may, e.g., be the MDCT domain.

According to an embodiment, the first reconstruction unit may, e.g., be configured to reconstruct the third audio signal portion by conducting a first fading to a noise like
10 spectrum. The second reconstruction unit may, e.g., be configured to reconstruct the fourth audio signal portion by conducting a second fading to a noise like spectrum and/or a second fading of an LTP gain. Moreover, the first reconstruction unit and the second reconstruction unit may, e.g., be configured to conduct the first fading and the second fading to a noise like spectrum and/or a second fading of an LTP gain with the same
15 fading speed.

In an embodiment, the apparatus may, e.g., further comprise a first aggregation unit for determining a first aggregated value depending on the first audio signal portion. Moreover, the apparatus further may, e.g., comprise a second aggregation unit for determining,
20 depending on the second audio signal portion, a second aggregated value as the value derived from the second audio signal portion. The noise level tracing unit may, e.g., be configured to receive the first aggregated value as the first signal portion information being represented in the tracing domain, wherein the noise level tracing unit may, e.g., be configured to receive the second aggregated value as the second signal portion
25 information being represented in the tracing domain, and wherein the noise level tracing unit is configured to determine the noise level information depending on the first aggregated value being represented in the tracing domain and depending on the second aggregated value being represented in the tracing domain.

30 According to an embodiment, the first aggregation unit may, e.g., be configured to determine the first aggregated value such that the first aggregated value indicates a root mean square of the first audio signal portion or of a signal derived from the first audio signal portion. The second aggregation unit is configured to determine the second aggregated value such that the second aggregated value indicates a root mean square of
35 the second audio signal portion or of a signal derived from the second audio signal portion.

In an embodiment, the first transform unit may, e.g., be configured to transform the value derived from the second audio signal portion from the second input domain to the tracing domain by applying a gain value on the value derived from the second audio signal portion.

5

According to an embodiment, the gain value may, e.g., indicate a gain introduced by Linear predictive coding synthesis, or wherein the gain value indicates a gain introduced by Linear predictive coding synthesis and deemphasis.

- 10 In an embodiment, the noise level tracing unit may, e.g., be configured to determine the noise level information by applying a minimum statistics approach.

- 15 According to an embodiment, the noise level tracing unit may, e.g., be configured to determine a comfort noise level as the noise level information. The reconstruction unit may, e.g., be configured to reconstruct the third audio signal portion depending on the noise level information, if said third frame of the plurality of frames is not received by the receiving interface or if said third frame is received by the receiving interface but is corrupted.

- 20 In an embodiment, the noise level tracing unit may, e.g., be configured to determine a comfort noise level as the noise level information derived from a noise level spectrum, wherein said noise level spectrum is obtained by applying the minimum statistics approach. The reconstruction unit may, e.g., be configured to reconstruct the third audio signal portion depending on a plurality of Linear Predictive coefficients, if said third frame
25 of the plurality of frames is not received by the receiving interface or if said third frame is received by the receiving interface but is corrupted.

- 30 According to an embodiment, the first reconstruction unit may, e.g., be configured to reconstruct the third audio signal portion depending on the noise level information and depending on the first audio signal portion, if said third frame of the plurality of frames is not received by the receiving interface or if said third frame is received by the receiving interface but is corrupted.

- 35 In an embodiment, the first reconstruction unit may, e.g., be configured to reconstruct the third audio signal portion by attenuating or amplifying the first audio signal portion.

According to an embodiment, the second reconstruction unit may, e.g., be configured to reconstruct the fourth audio signal portion depending on the noise level information and depending on the second audio signal portion.

- 5 In an embodiment, the second reconstruction unit may, e.g., be configured to reconstruct the fourth audio signal portion by attenuating or amplifying the second audio signal portion.

10 According to an embodiment, the apparatus may, e.g., further comprise a long-term prediction unit comprising a delay buffer, wherein the long-term prediction unit may, e.g., be configured to generate a processed signal depending on the first or the second audio signal portion, depending on a delay buffer input being stored in the delay buffer and depending on a long-term prediction gain, and wherein the long-term prediction unit is
15 configured to fade the long-term prediction gain towards zero, if said third frame of the plurality of frames is not received by the receiving interface or if said third frame is received by the receiving interface but is corrupted.

In an embodiment, the long-term prediction unit may, e.g., be configured to fade the long-term prediction gain towards zero, wherein a speed with which the long-term prediction
20 gain is faded to zero depends on a fade-out factor.

In an embodiment, the long-term prediction unit may, e.g., be configured to update the delay buffer input by storing the generated processed signal in the delay buffer, if said third frame of the plurality of frames is not received by the receiving interface or if said
25 third frame is received by the receiving interface but is corrupted.

Moreover, a method for decoding an audio signal is provided. The method comprises:

- Receiving a first frame comprising a first audio signal portion of the audio signal,
30 and receiving a second frame comprising a second audio signal portion of the audio signal.
- Determining noise level information depending on at least one of the first audio signal portion and the second audio signal portion, wherein the noise level information is represented in a tracing domain.
35
- Reconstructing, in a first reconstruction domain, a third audio signal portion of the audio signal depending on the noise level information, if a third frame of the

plurality of frames is not received or if said third frame is received but is corrupted, wherein the first reconstruction domain is different from or equal to the tracing domain.

- 5 - Transforming the noise level information from the tracing domain to a second reconstruction domain, if a fourth frame of the plurality of frames is not received or if said fourth frame is received but is corrupted, wherein the second reconstruction domain is different from the tracing domain, and wherein the second reconstruction domain is different from the first reconstruction domain. And:
- 10 - Reconstructing, in the second reconstruction domain, a fourth audio signal portion of the audio signal depending on the noise level information being represented in the second reconstruction domain, if said fourth frame of the plurality of frames is not received or if said fourth frame is received but is corrupted.

15

Moreover, a computer program for implementing the above-described method when being executed on a computer or signal processor is provided.

Moreover, an apparatus for decoding an encoded audio signal to obtain a reconstructed audio signal is provided. The apparatus comprises a receiving interface for receiving one or more frames comprising information on a plurality of audio signal samples of an audio signal spectrum of the encoded audio signal, and a processor for generating the reconstructed audio signal. The processor is configured to generate the reconstructed audio signal by fading a modified spectrum to a target spectrum, if a current frame is not received by the receiving interface or if the current frame is received by the receiving interface but is corrupted, wherein the modified spectrum comprises a plurality of modified signal samples, wherein, for each of the modified signal samples of the modified spectrum, an absolute value of said modified signal sample is equal to an absolute value of one of the audio signal samples of the audio signal spectrum. Moreover, the processor is configured to not fade the modified spectrum to the target spectrum, if the current frame of the one or more frames is received by the receiving interface and if the current frame being received by the receiving interface is not corrupted.

According to an embodiment, the target spectrum may, e.g., be a noise like spectrum.

35

In an embodiment, the noise like spectrum may, e.g., represent white noise.

According to an embodiment, the noise like spectrum may, e.g., be shaped.

In an embodiment, the shape of the noise like spectrum may, e.g., depend on an audio signal spectrum of a previously received signal.

- 5 According to an embodiment, the noise like spectrum may, e.g., be shaped depending on the shape of the audio signal spectrum.

In an embodiment, the processor may, e.g., employ a tilt factor to shape the noise like spectrum.

10

According to an embodiment, the processor may, e.g., employ the formula

$$\text{shaped_noise}[i] = \text{noise} * \text{power}(\text{tilt_factor}, i/N)$$

- 15 wherein N indicates the number of samples, wherein i is an index, wherein $0 \leq i < N$, with $\text{tilt_factor} > 0$, and wherein power is a power function.

$\text{power}(x, y)$ indicates x^y

- 20 $\text{power}(\text{tilt_factor}, i/N)$ indicates $\text{tilt_factor}^{\frac{i}{N}}$

If the tilt_factor is smaller 1 this means attenuation with increasing i . If the tilt_factor is larger 1 means amplification with increasing i .

- 25 According to another embodiment, the processor may, e.g., employ the formula

$$\text{shaped_noise}[i] = \text{noise} * (1 + i / (N-1) * (\text{tilt_factor} - 1))$$

- 30 wherein N indicates the number of samples, wherein i is an index, wherein $0 \leq i < N$, with $\text{tilt_factor} > 0$.

If the tilt_factor is smaller 1 this means attenuation with increasing i . If the tilt_factor is larger 1 means amplification with increasing i .

- 35 According to an embodiment, the processor may, e.g., be configured to generate the modified spectrum, by changing a sign of one or more of the audio signal samples of the

audio signal spectrum, if the current frame is not received by the receiving interface or if the current frame being received by the receiving interface is corrupted.

In an embodiment, each of the audio signal samples of the audio signal spectrum may,
5 e.g., be represented by a real number but not by an imaginary number.

According to an embodiment, the audio signal samples of the audio signal spectrum may, e.g., be represented in a Modified Discrete Cosine Transform domain.

10 In another embodiment, the audio signal samples of the audio signal spectrum may, e.g., be represented in a Modified Discrete Sine Transform domain.

According to an embodiment, the processor may, e.g., be configured to generate the modified spectrum by employing a random sign function which randomly or pseudo-
15 randomly outputs either a first or a second value.

In an embodiment, the processor may, e.g., be configured to fade the modified spectrum to the target spectrum by subsequently decreasing an attenuation factor.

20 According to an embodiment, the processor may, e.g., be configured to fade the modified spectrum to the target spectrum by subsequently increasing an attenuation factor.

In an embodiment, if the current frame is not received by the receiving interface or if the current frame being received by the receiving interface is corrupted, the processor may,
25 e.g., be configured to generate the reconstructed audio signal by employing the formula:

$$x[i] = (1 - \text{cum_damping}) * \text{noise}[i] + \text{cum_damping} * \text{random_sign}() * x_old[i]$$

30 wherein i is an index, wherein $x[i]$ indicates a sample of the reconstructed audio signal, wherein cum_damping is an attenuation factor, wherein $x_old[i]$ indicates one of the audio signal samples of the audio signal spectrum of the encoded audio signal, wherein $\text{random_sign}()$ returns 1 or -1, and wherein noise is a random vector indicating the target spectrum.

35

In an embodiment, said random vector noise may, e.g., be scaled such that its quadratic mean is similar to the quadratic mean of the spectrum of the encoded audio signal being comprised by one of the frames being last received by the receiving interface.

According to a general embodiment, the processor may, e.g., be configured to generate the reconstructed audio signal, by employing a random vector which is scaled such that its quadratic mean is similar to the quadratic mean of the spectrum of the encoded audio signal being comprised by one of the frames being last received by the receiving interface.

Moreover, a method for decoding an encoded audio signal to obtain a reconstructed audio signal is provided. The method comprises:

- 10 - Receiving one or more frames comprising information on a plurality of audio signal samples of an audio signal spectrum of the encoded audio signal. And:
- Generating the reconstructed audio signal.

15 Generating the reconstructed audio signal is conducted by fading a modified spectrum to a target spectrum, if a current frame is not received or if the current frame is received but is corrupted, wherein the modified spectrum comprises a plurality of modified signal samples, wherein, for each of the modified signal samples of the modified spectrum, an absolute value of said modified signal sample is equal to an absolute value of one of the
20 audio signal samples of the audio signal spectrum. The modified spectrum is not faded to a white noise spectrum, if the current frame of the one or more frames is received and if the current frame being received is not corrupted.

Moreover, a computer program for implementing the above-described method when being
25 executed on a computer or signal processor is provided.

Embodiments realize a fade MDCT spectrum to white noise prior to FDNS Application (FDNS = Frequency Domain Noise Substitution).

30 According to the prior art, in ACELP based codecs, the innovative codebook is replaced with a random vector (e.g., with noise). In embodiments, the ACELP approach, which consists of replacing the innovative codebook with a random vector (e.g., with noise) is adopted to the TCX decoder structure. Here, the equivalent of the innovative codebook is the MDCT spectrum usually received within the bitstream and fed into the FDNS.

35

The classical MDCT concealment approach would be to simply repeat this spectrum as is or to apply a certain randomization process, which basically prolongs the spectral shape of the last received frame [LS01]. This has the drawback that the short-term spectral

shape is prolonged, leading frequently to a repetitive, metallic sound which is not background noise like, and thus cannot be used as comfort noise.

Using the proposed method the short term spectral shaping is performed by the FDNS
5 and the TCX LTP, the spectral shaping on the long run is performed by the FDNS only. The shaping by the FDNS is faded from the short-term spectral shape to the traced long-term spectral shape of the background noise, and the TCX LTP is faded to zero.

10 Fading the FDNS coefficients to traced background noise coefficients leads to having a smooth transition between the last good spectral envelope and the spectral background envelope which should be targeted in the long run, in order to achieve a pleasant background noise in case of long burst frame losses.

15 In contrast, according to the state of the art, for transform based codecs, noise like concealment is conducted by frame repetition or noise substitution in the frequency domain [LS01]. In the prior art, the noise substitution is usually performed by sign scrambling of the spectral bins. If in the prior art TCX (frequency domain) sign scrambling is used during concealment, the last received MDCT coefficients are re-used and each
20 sign is randomized before the spectrum is inversely transformed to the time domain. The drawback of this procedure of the prior art is, that for consecutively lost frames the same spectrum is used again and again, just with different sign randomizations and global attenuation. When looking to the spectral envelope over time on a coarse time grid, it can be seen that the envelope is approximately constant during consecutive frame loss, because the band energies are kept constant relatively to each other within a frame and
25 are just globally attenuated. In the used coding system, according to the prior art, the spectral values are processed using FDNS, in order to restore the original spectrum. This means, that if one wants to fade the MDCT spectrum to a certain spectral envelope (using FDNS coefficients, e.g., describing the current background noise), the result is not just dependent on the FDNS coefficients, but also dependent on the previously decoded
30 spectrum which was sign scrambled. The above-mentioned embodiments overcome these disadvantages of the prior art.

Embodiments are based on the finding that it is necessary to fade the spectrum used for the sign scrambling to white noise, before feeding it into the FDNS processing. Otherwise
35 the outputted spectrum will never match the targeted envelope used for FDNS processing.

In embodiments, the same fading speed is used for LTP gain fading as for the white noise fading.

Moreover, an apparatus for decoding an encoded audio signal to obtain a reconstructed audio signal is provided. The apparatus comprises a receiving interface for receiving a plurality of frames, a delay buffer for storing audio signal samples of the decoded audio signal, a sample selector for selecting a plurality of selected audio signal samples from the audio signal samples being stored in the delay buffer, and a sample processor for processing the selected audio signal samples to obtain reconstructed audio signal samples of the reconstructed audio signal. The sample selector is configured to select, if a current frame is received by the receiving interface and if the current frame being received by the receiving interface is not corrupted, the plurality of selected audio signal samples from the audio signal samples being stored in the delay buffer depending on a pitch lag information being comprised by the current frame. Moreover, the sample selector is configured to select, if the current frame is not received by the receiving interface or if the current frame being received by the receiving interface is corrupted, the plurality of selected audio signal samples from the audio signal samples being stored in the delay buffer depending on a pitch lag information being comprised by another frame being received previously by the receiving interface.

According to an embodiment, the sample processor may, e.g., be configured to obtain the reconstructed audio signal samples, if the current frame is received by the receiving interface and if the current frame being received by the receiving interface is not corrupted, by rescaling the selected audio signal samples depending on the gain information being comprised by the current frame. Moreover, the sample selector may, e.g., be configured to obtain the reconstructed audio signal samples, if the current frame is not received by the receiving interface or if the current frame being received by the receiving interface is corrupted, by rescaling the selected audio signal samples depending on the gain information being comprised by said another frame being received previously by the receiving interface.

In an embodiment, the sample processor may, e.g., be configured to obtain the reconstructed audio signal samples, if the current frame is received by the receiving interface and if the current frame being received by the receiving interface is not corrupted, by multiplying the selected audio signal samples and a value depending on the gain information being comprised by the current frame. Moreover, the sample selector is configured to obtain the reconstructed audio signal samples, if the current frame is not received by the receiving interface or if the current frame being received by the receiving interface is corrupted, by multiplying the selected audio signal samples and a value

depending on the gain information being comprised by said another frame being received previously by the receiving interface.

According to an embodiment, the sample processor may, e.g., be configured to store the
5 reconstructed audio signal samples into the delay buffer.

In an embodiment, the sample processor may, e.g., be configured to store the reconstructed audio signal samples into the delay buffer before a further frame is received by the receiving interface.

10

According to an embodiment, the sample processor may, e.g., be configured to store the reconstructed audio signal samples into the delay buffer after a further frame is received by the receiving interface.

15 In an embodiment, the sample processor may, e.g., be configured to rescale the selected audio signal samples depending on the gain information to obtain rescaled audio signal samples and by combining the rescaled audio signal samples with input audio signal samples to obtain the processed audio signal samples.

20 According to an embodiment, the sample processor may, e.g., be configured to store the processed audio signal samples, indicating the combination of the rescaled audio signal samples and the input audio signal samples, into the delay buffer, and to not store the rescaled audio signal samples into the delay buffer, if the current frame is received by the receiving interface and if the current frame being received by the receiving interface is not
25 corrupted. Moreover, the sample processor is configured to store the rescaled audio signal samples into the delay buffer and to not store the processed audio signal samples into the delay buffer, if the current frame is not received by the receiving interface or if the current frame being received by the receiving interface is corrupted.

30 According to another embodiment, the sample processor may, e.g., be configured to store the processed audio signal samples into the delay buffer, if the current frame is not received by the receiving interface or if the current frame being received by the receiving interface is corrupted.

35 In an embodiment, the sample selector may, e.g., be configured to obtain the reconstructed audio signal samples by rescaling the selected audio signal samples depending on a modified gain, wherein the modified gain is defined according to the formula:

gain = gain_past * damping;

wherein gain is the modified gain, wherein the sample selector may, e.g., be configured
 5 to set gain_past to gain after gain and has been calculated, and wherein damping
 is a real value.

According to an embodiment, the sample selector may, e.g., be configured to calculate
 the modified gain.

10

In an embodiment, damping may, e.g., be defined according to: $0 \leq \text{damping} \leq 1$.

According to an embodiment, the modified gain gain may, e.g., be set to zero, if at least
 a predefined number of frames have not been received by the receiving interface since a
 15 frame last has been received by the receiving interface.

Moreover, a method for decoding an encoded audio signal to obtain a reconstructed audio
 signal is provided. The method comprises:

- 20 - Receiving a plurality of frames.
- Storing audio signal samples of the decoded audio signal.
- Selecting a plurality of selected audio signal samples from the audio signal
 25 samples being stored in the delay buffer. And:
- Processing the selected audio signal samples to obtain reconstructed audio signal
 samples of the reconstructed audio signal.

30 If a current frame is received and if the current frame being received is not corrupted, the
 step of selecting the plurality of selected audio signal samples from the audio signal
 samples being stored in the delay buffer is conducted depending on a pitch lag
 information being comprised by the current frame. Moreover, if the current frame is not
 received or if the current frame being received is corrupted, the step of selecting the
 35 plurality of selected audio signal samples from the audio signal samples being stored in
 the delay buffer is conducted depending on a pitch lag information being comprised by
 another frame being received previously by the receiving interface.

Moreover, a computer program for implementing the above-described method when being executed on a computer or signal processor is provided.

Embodiments employ TCX LTP (TxC LTP = Transform Coded Excitation Long-Term Prediction). During normal operation, the TCX LTP memory is updated with the synthesized signal, containing noise and reconstructed tonal components.

Instead of disabling the TCX LTP during concealment, its normal operation may be continued during concealment with the parameters received in the last good frame. This preserves the spectral shape of the signal, particularly those tonal components which are modelled by the LTP filter.

Moreover, embodiments decouple the TCX LTP feedback loop. A simple continuation of the normal TCX LTP operation introduces additional noise, since with each update step further randomly generated noise from the LTP excitation is introduced. The tonal components are hence getting distorted more and more over time by the added noise.

To overcome this, only the updated TCX LTP buffer may be fed back (without adding noise), in order to not pollute the tonal information with undesired random noise.

Furthermore, according to embodiments, the TCX LTP gain is faded to zero.

These embodiments are based on the finding that continuing the TCX LTP helps to preserve the signal characteristics on the short term, but has drawbacks on the long term: The signal played out during concealment will include the voicing/tonal information which was present preceding to the loss. Especially for clean speech or speech over background noise, it is extremely unlikely that a tone or harmonic will decay very slowly over a very long time. By continuing the TCX LTP operation during concealment, particularly if the LTP memory update is decoupled (just tonal components are fed back and not the sign scrambled part), the voicing/tonal information will stay present in the concealed signal for the whole loss, being attenuated just by the overall fade-out to the comfort noise level. Moreover, it is impossible to reach the comfort noise envelope during burst packet losses, if the TCX LTP is applied during the burst loss without being attenuated over time, because the signal will then always incorporate the voicing information of the LTP.

Therefore, the TCX LTP gain is faded towards zero, such that tonal components represented by the LTP will be faded to zero, at the same time the signal is faded to the

background signal level and shape, and such that the fade-out reaches the desired spectral background envelope (comfort noise) without incorporating undesired tonal components.

- 5 In embodiments, the same fading speed is used for LTP gain fading as for the white noise fading.

In contrast, in the prior art, there is no transform codec known that uses LTP during concealment. For the MPEG-4 LTP [ISO09] no concealment approaches exist in the prior art. Another MDCT based codec of the prior art which makes use of an LTP is CELT, but this codec uses an ACELP-like concealment for the first five frames, and for all subsequent frames background noise is generated, which does not make use of the LTP. A drawback of the prior art of not using the TCX LTP is, that all tonal components being modelled with the LTP disappear abruptly. Moreover, in ACELP based codecs of the prior art, the LTP operation is prolonged during concealment, and the gain of the adaptive codebook is faded towards zero. With regard to the feedback loop operation, the prior art employs two approaches, either the whole excitation, e.g., the sum of the innovative and the adaptive excitation, is fed back (AMR-WB); or only the updated adaptive excitation, e.g., the tonal signal parts, is fed back (G.718). The above-mentioned embodiments overcome the disadvantages of the prior art.

In the following, embodiments of the present invention are described in more detail with reference to the figures, in which:

- 25 Fig. 1a illustrates an apparatus for decoding an audio signal according to an embodiment,
- Fig. 1b illustrates an apparatus for decoding an audio signal according to another embodiment,
- 30 Fig. 1c illustrates an apparatus for decoding an audio signal according to another embodiment, wherein the apparatus further comprises a first and a second aggregation unit,
- 35 Fig. 1d illustrates an apparatus for decoding an audio signal according to a further embodiment, wherein the apparatus moreover comprises a long-term prediction unit comprising a delay buffer,

- Fig. 2 illustrates the decoder structure of G.718,
- Fig. 3 depicts a scenario, where the fade-out factor of G.722 depends on class information,
- 5 Fig. 4 shows an approach for amplitude prediction using linear regression,
- Fig. 5 illustrates the burst loss behavior of Constrained-Energy Lapped Transform (CELT),
- 10 Fig. 6 shows a background noise level tracing according to an embodiment in the decoder during an error-free operation mode,
- Fig. 7 illustrates gain derivation of LPC synthesis and deemphasis according to an embodiment,
- 15 Fig. 8 depicts comfort noise level application during packet loss according to an embodiment,
- 20 Fig. 9 illustrates advanced high pass gain compensation during ACELP concealment according to an embodiment,
- Fig. 10 depicts the decoupling of the LTP feedback loop during concealment according to an embodiment,
- 25 Fig. 11 illustrates an apparatus for decoding an encoded audio signal to obtain a reconstructed audio signal according to an embodiment,
- Fig. 12 shows an apparatus for decoding an encoded audio signal to obtain a reconstructed audio signal according to another embodiment, and
- 30 Fig. 13 illustrates an apparatus for decoding an encoded audio signal to obtain a reconstructed audio signal a further embodiment, and
- 35 Fig. 14 illustrates an apparatus for decoding an encoded audio signal to obtain a reconstructed audio signal another embodiment.

Fig. 1a illustrates an apparatus for decoding an audio signal according to an embodiment.

The apparatus comprises a receiving interface 110. The receiving interface is configured to receive a plurality of frames, wherein the receiving interface 110 is configured to receive a first frame of the plurality of frames, said first frame comprising a first audio signal portion of the audio signal, said first audio signal portion being represented in a first domain. Moreover, the receiving interface 110 is configured to receive a second frame of the plurality of frames, said second frame comprising a second audio signal portion of the audio signal.

10

Moreover, the apparatus comprises a transform unit 120 for transforming the second audio signal portion or a value or signal derived from the second audio signal portion from a second domain to a tracing domain to obtain a second signal portion information, wherein the second domain is different from the first domain, wherein the tracing domain is different from the second domain, and wherein the tracing domain is equal to or different from the first domain.

15

Furthermore, the apparatus comprises a noise level tracing unit 130, wherein the noise level tracing unit is configured to receive a first signal portion information being represented in the tracing domain, wherein the first signal portion information depends on the first audio signal portion, wherein the noise level tracing unit is configured to receive the second signal portion being represented in the tracing domain, and wherein the noise level tracing unit is configured to determine noise level information depending on the first signal portion information being represented in the tracing domain and depending on the second signal portion information being represented in the tracing domain.

20

25

Moreover, the apparatus comprises a reconstruction unit for reconstructing a third audio signal portion of the audio signal depending on the noise level information, if a third frame of the plurality of frames is not received by the receiving interface but is corrupted.

30

Regarding the first and/or the second audio signal portion, for example, the first and/or the second audio signal portion may, e.g., be fed into one or more processing units (not shown) for generating one or more loudspeaker signals for one or more loudspeakers, so that the received sound information comprised by the first and/or the second audio signal portion can be replayed.

35

Moreover, however, the first and second audio signal portion are also used for concealment, e.g., in case subsequent frames do not arrive at the receiver or in case that subsequent frames are erroneous.

5 Inter alia, the present invention is based on the finding that noise level tracing should be conducted in a common domain, herein referred to as "tracing domain". The tracing domain, may, e.g., be an excitation domain, for example, the domain in which the signal is represented by LPCs (LPC = Linear Predictive Coefficient) or by ISPs (ISP = Immittance Spectral Pair) as described in AMR-WB and AMR-WB+ (see [3GP12a], [3GP12b],
10 [3GP09a], [3GP09b], [3GP09c]). Tracing the noise level in a single domain has inter alia the advantage that aliasing effects are avoided when the signal switches between a first representation in a first domain and a second representation in a second domain (for example, when the signal representation switches from ACELP to TCX or vice versa).

15 Regarding the transform unit 120, what is transformed is either the second audio signal portion itself, or a signal derived from the second audio signal portion (e.g., the second audio signal portion has been processed to obtain the derived signal), or a value derived from the second audio signal portion (e.g., the second audio signal portion has been processed to obtain the derived value).

20

Regarding the first audio signal portion, in some embodiments, the first audio signal portion may be processed and/or transformed to the tracing domain.

25 In other embodiments, however, the first audio signal portion may be already represented in the tracing domain.

In some embodiments, the first signal portion information is identical to the first audio signal portion. In other embodiments, the first signal portion information is, e.g., an aggregated value depending on the first audio signal portion.

30

Now, at first, fade-out to a comfort noise level is considered in more detail.

The fade-out approach described may, e.g., be implemented in a low-delay version of xHE-AAC [NMR+12] (xHE-AAC = Extended High Efficiency AAC), which is able to switch
35 seamlessly between ACELP (speech) and MDCT (music / noise) coding on a per-frame basis.

Regarding common level tracing in a tracing domain, for example, an excitation domain, as to apply a smooth fade-out to an appropriate comfort noise level during packet loss, such comfort noise level needs to be identified during the normal decoding process. It may, e.g., be assumed, that a noise level similar to the background noise is most comfortable. Thus, the background noise level may be derived and constantly updated during normal decoding.

The present invention is based on the finding that when having a switched core codec (e.g., ACELP and TCX), considering a common background noise level independent from the chosen core coder is particularly suitable.

Fig. 6 depicts a background noise level tracing according to a preferred embodiment in the decoder during the error-free operation mode, e.g., during normal decoding.

The tracing itself may, e.g., be performed using the minimum statistics approach (see [Mar01]).

This traced background noise level may, e.g., be considered as the noise level information mentioned above.

For example, the minimum statistics noise estimation presented in the document: "Rainer Martin, *Noise power spectral density estimation based on optimal smoothing and minimum statistics*, IEEE Transactions on Speech and Audio Processing 9 (2001), no. 5, 504 –512" [Mar01] may be employed for background noise level tracing.

Correspondingly, in some embodiments, the noise level tracing unit 130 is configured to determine noise level information by applying a minimum statistics approach, e.g., by employing the minimum statistics noise estimation of [Mar01].

Subsequently, some considerations and details of this tracing approach are described.

Regarding level tracing, the background is supposed to be noise-like. Hence it is preferable to perform the level tracing in the excitation domain to avoid tracing foreground tonal components which are taken out by the LPC. For example, ACELP noise filling may also employ the background noise level in the excitation domain. With tracing in the excitation domain, only one single tracing of the background noise level can serve two purposes, which saves computational complexity. In a preferred embodiment, the tracing is performed in the ACELP excitation domain.

Fig. 7 illustrates gain derivation of LPC synthesis and deemphasis according to an embodiment.

- 5 Regarding level derivation, the level derivation may, for example, be conducted either in time domain or in excitation domain, or in any other suitable domain. If the domains for the level derivation and the level tracing differ, a gain compensation may, e.g., be needed.

10 In the preferred embodiment, the level derivation for ACELP is performed in the excitation domain. Hence, no gain compensation is required.

For TCX, a gain compensation may, e.g., be needed to adjust the derived level to the ACELP excitation domain.

- 15 In the preferred embodiment, the level derivation for TCX takes place in the time domain. A manageable gain compensation was found for this approach: The gain introduced by LPC synthesis and deemphasis is derived as shown in Fig. 7 and the derived level is divided by this gain.

- 20 Alternatively, the level derivation for TCX could be performed in the TCX excitation domain. However, the gain compensation between the TCX excitation domain and the ACELP excitation domain was deemed too complicated.

25 Thus, returning to Fig. 1a, in some embodiments, the first audio signal portion is represented in a time domain as the first domain. The transform unit 120 is configured to transform the second audio signal portion or the value derived from the second audio signal portion from an excitation domain being the second domain to the time domain being the tracing domain. In such embodiments, the noise level tracing unit 130 is configured to receive the first signal portion information being represented in the time domain as the tracing domain. Moreover, the noise level tracing unit 130 is configured to receive the second signal portion being represented in the time domain as the tracing domain.

35 In other embodiments, the first audio signal portion is represented in an excitation domain as the first domain. The transform unit 120 is configured to transform the second audio signal portion or the value derived from the second audio signal portion from a time domain being the second domain to the excitation domain being the tracing domain. In such embodiments, the noise level tracing unit 130 is configured to receive the first signal

portion information being represented in the excitation domain as the tracing domain. Moreover, the noise level tracing unit 130 is configured to receive the second signal portion being represented in the excitation domain as the tracing domain.

- 5 In an embodiment, the first audio signal portion may, e.g., be represented in an excitation domain as the first domain, wherein the noise level tracing unit 130 may, e.g., be configured to receive the first signal portion information, wherein said first signal portion information is represented in the FFT domain, being the tracing domain, and wherein said first signal portion information depends on said first audio signal portion being represented
- 10 in the excitation domain, wherein the transform unit 120 may, e.g., be configured to transform the second audio signal portion or the value derived from the second audio signal portion from a time domain being the second domain to an FFT domain being the tracing domain, and wherein the noise level tracing unit 130 may, e.g., be configured to receive the second audio signal portion being represented in the FFT domain.

- 15 Fig. 1b illustrates an apparatus according to another embodiment. In Fig. 1b, the transform unit 120 of Fig. 1a is a first transform unit 120, and the reconstruction unit 140 of Fig. 1a is a first reconstruction unit 140. The apparatus further comprises a second transform unit 121 and a second reconstruction unit 141.

- 20 The second transform unit 121 is configured to transform the noise level information from the tracing domain to the second domain, if a fourth frame of the plurality of frames is not received by the receiving interface or if said fourth frame is received by the receiving interface but is corrupted.

- 25 Moreover, the second reconstruction unit 141 is configured to reconstruct a fourth audio signal portion of the audio signal depending on the noise level information being represented in the second domain if said fourth frame of the plurality of frames is not received by the receiving interface or if said fourth frame is received by the receiving
- 30 interface but is corrupted.

- Fig. 1c illustrates an apparatus for decoding an audio signal according to another embodiment. The apparatus further comprises a first aggregation unit 150 for determining a first aggregated value depending on the first audio signal portion. Moreover, the
- 35 apparatus of Fig. 1c further comprises a second aggregation unit 160 for determining a second aggregated value as the value derived from the second audio signal portion depending on the second audio signal portion. In the embodiment of Fig. 1c, the noise level tracing unit 130 is configured to receive first aggregated value as the first signal

portion information being represented in the tracing domain, wherein the noise level tracing unit 130 is configured to receive the second aggregated value as the second signal portion information being represented in the tracing domain. The noise level tracing unit 130 is configured to determine noise level information depending on the first aggregated value being represented in the tracing domain and depending on the second aggregated value being represented in the tracing domain.

In an embodiment, the first aggregation unit 150 is configured to determine the first aggregated value such that the first aggregated value indicates a root mean square of the first audio signal portion or of a signal derived from the first audio signal portion. Moreover, the second aggregation unit 160 is configured to determine the second aggregated value such that the second aggregated value indicates a root mean square of the second audio signal portion or of a signal derived from the second audio signal portion.

Fig. 6 illustrates an apparatus for decoding an audio signal according to a further embodiment.

In Fig. 6, background level tracing unit 630 implements a noise level tracing unit 130 according to Fig. 1a.

Moreover, in Fig. 6, RMS unit 650 (RMS = root mean square) is a first aggregation unit and RMS unit 660 is a second aggregation unit.

According to some embodiments, the (first) transform unit 120 of Fig. 1a, Fig. 1b and Fig. 1c is configured to transform the value derived from the second audio signal portion from the second domain to the tracing domain by applying a gain value (x) on the value derived from the second audio signal portion, e.g., by dividing the value derived from the second audio signal portion by a gain value (x). In other embodiments, a gain value may, e.g., be multiplied.

In some embodiments, the gain value (x) may, e.g., indicate a gain introduced by Linear predictive coding synthesis, or the gain value (x) may, e.g., indicate a gain introduced by Linear predictive coding synthesis and deemphasis.

In Fig. 6, unit 622 provides the value (x) which indicates the gain introduced by Linear predictive coding synthesis and deemphasis. Unit 622 then divides the value, provided by the second aggregation unit 660, which is a value derived from the second audio signal

portion, by the provided gain value (x) (e.g., either by dividing by x , or by multiplying the value $1/x$). Thus, unit 620 of Fig. 6 which comprises units 621 and 622 implements the first transform unit of Fig. 1a, Fig. 1b or Fig. 1c.

- 5 The apparatus of Fig. 6 receives a first frame with a first audio signal portion being a voiced excitation and/or an unvoiced excitation and being represented in the tracing domain, in Fig. 6 an (ACELP) LPC domain. The first audio signal portion is fed into an LPC Synthesis and De-Emphasis unit 671 for processing to obtain a time-domain first audio signal portion output. Moreover, the first audio signal portion is fed into RMS module
10 650 to obtain a first value indicating a root mean square of the first audio signal portion. This first value (first RMS value) is represented in the tracing domain. The first RMS value, being represented in the tracing domain, is then fed into the noise level tracing unit 630.
- 15 Moreover, the apparatus of Fig. 6 receives a second frame with a second audio signal portion comprising an MDCT spectrum and being represented in an MDCT domain. Noise filling is conducted by a noise filling module 681, frequency-domain noise shaping is conducted by a frequency-domain noise shaping module 682, transformation to the time domain is conducted by an iMDCT/OLA module 683 (OLA = overlap-add) and long-term
20 prediction is conducted by a long-term prediction unit 684. The long-term prediction unit may, e.g., comprise a delay buffer (not shown in Fig. 6).

The signal derived from the second audio signal portion is then fed into RMS module 660 to obtain a second value indicating a root mean square of that signal derived from the
25 second audio signal portion is obtained. This second value (second RMS value) is still represented in the time domain. Unit 620 then transforms the second RMS value from the time domain to the tracing domain, here, the (ACELP) LPC domain. The second RMS value, being represented in the tracing domain, is then fed into the noise level tracing unit 630.

30

In embodiments, level tracing is conducted in the excitation domain, but TCX fade-out is conducted in the time domain.

Whereas during normal decoding the background noise level is traced, it may, e.g., be
35 used during packet loss as an indicator of an appropriate comfort noise level, to which the last received signal is smoothly faded level-wise.

Deriving the level for tracing and applying the level fade-out are in general independent from each other and could be performed in different domains. In the preferred embodiment, the level application is performed in the same domains as the level derivation, leading to the same benefits that for ACELP, no gain compensation is needed, and that for TCX, the inverse gain compensation as for the level derivation (see Fig. 6) is
5 needed and hence the same gain derivation can be used, as illustrated by Fig. 7.

In the following, compensation of an influence of the high pass filter on the LPC synthesis gain according to embodiments is described.
10 Fig. 8 outlines this approach. In particular, Fig. 8 illustrates comfort noise level application during packet loss.

In Fig. 8, high pass gain filter unit 643, multiplication unit 644, fading unit 645, high pass
15 filter unit 646, fading unit 647 and combination unit 648 together form a first reconstruction unit.

Moreover, in Fig. 8, background level provision unit 631 provides the noise level information. For example, background level provision unit 631 may be equally
20 implemented as background level tracing unit 630 of Fig. 6.

Furthermore, in Fig. 8, LPC Synthesis & De-Emphasis Gain Unit 649 and multiplication unit 641 together for a second transform unit 640.

25 Moreover, in Fig. 8, fading unit 642 represents a second reconstruction unit.

In the embodiment of Fig. 8, voiced and unvoiced excitation are faded separately: The voiced excitation is faded to zero, but the unvoiced excitation is faded towards the comfort noise level. Fig. 8 furthermore depicts a high pass filter, which is introduced into the signal
30 chain of the unvoiced excitation to suppress low frequency components for all cases except when the signal was classified as unvoiced.

As to model the influence of the high pass filter, the level after LPC synthesis and de-emphasis is computed once with and once without the high pass filter. Subsequently the
35 ratio of those two levels is derived and used to alter the applied background level.

This is illustrated by Fig. 9. In particular, Fig. 9 depicts advanced high pass gain compensation during ACELP concealment according to an embodiment.

Instead of the current excitation signal just a simple impulse is used as input for this computation. This allows for a reduced complexity, since the impulse response decays quickly and so the RMS derivation can be performed on a shorter time frame. In practice,
5 just one subframe is used instead of the whole frame.

According to an embodiment, the noise level tracing unit 130 is configured to determine a comfort noise level as the noise level information. The reconstruction unit 140 is configured to reconstruct the third audio signal portion depending on the noise level
10 information, if said third frame of the plurality of frames is not received by the receiving interface 110 or if said third frame is received by the receiving interface 110 but is corrupted.

According to an embodiment, the noise level tracing unit 130 is configured to determine a
15 comfort noise level as the noise level information. The reconstruction unit 140 is configured to reconstruct the third audio signal portion depending on the noise level information, if said third frame of the plurality of frames is not received by the receiving interface 110 or if said third frame is received by the receiving interface 110 but is corrupted.

20 In an embodiment, the noise level tracing unit 130 is configured to determine a comfort noise level as the noise level information derived from a noise level spectrum, wherein said noise level spectrum is obtained by applying the minimum statistics approach. The reconstruction unit 140 is configured to reconstruct the third audio signal portion
25 depending on a plurality of Linear Predictive coefficients, if said third frame of the plurality of frames is not received by the receiving interface 110 or if said third frame is received by the receiving interface 110 but is corrupted.

In an embodiment, the (first and/or second) reconstruction unit 140, 141 may, e.g., be
30 configured to reconstruct the third audio signal portion depending on the noise level information and depending on the first audio signal portion, if said third (fourth) frame of the plurality of frames is not received by the receiving interface 110 or if said third (fourth) frame is received by the receiving interface 110 but is corrupted.

35 According to an embodiment, the (first and/or second) reconstruction unit 140, 141 may, e.g., be configured to reconstruct the third (or fourth) audio signal portion by attenuating or amplifying the first audio signal portion.

Fig. 14 illustrates an apparatus for decoding an audio signal. The apparatus comprises a receiving interface 110, wherein the receiving interface 110 is configured to receive a first frame comprising a first audio signal portion of the audio signal, and wherein the receiving interface 110 is configured to receive a second frame comprising a second audio signal portion of the audio signal.

Moreover, the apparatus comprises a noise level tracing unit 130, wherein the noise level tracing unit 130 is configured to determine noise level information depending on at least one of the first audio signal portion and the second audio signal portion (this means: depending on the first audio signal portion and/or the second audio signal portion), wherein the noise level information is represented in a tracing domain.

Furthermore, the apparatus comprises a first reconstruction unit 140 for reconstructing, in a first reconstruction domain, a third audio signal portion of the audio signal depending on the noise level information, if a third frame of the plurality of frames is not received by the receiving interface 110 or if said third frame is received by the receiving interface 110 but is corrupted, wherein the first reconstruction domain is different from or equal to the tracing domain.

Moreover, the apparatus comprises a transform unit 121 for transforming the noise level information from the tracing domain to a second reconstruction domain, if a fourth frame of the plurality of frames is not received by the receiving interface 110 or if said fourth frame is received by the receiving interface 110 but is corrupted, wherein the second reconstruction domain is different from the tracing domain, and wherein the second reconstruction domain is different from the first reconstruction domain, and

Furthermore, the apparatus comprises a second reconstruction unit 141 for reconstructing, in the second reconstruction domain, a fourth audio signal portion of the audio signal depending on the noise level information being represented in the second reconstruction domain, if said fourth frame of the plurality of frames is not received by the receiving interface 110 or if said fourth frame is received by the receiving interface 110 but is corrupted.

According to some embodiments, the tracing domain may, e.g., be wherein the tracing domain is a time domain, a spectral domain, an FFT domain, an MDCT domain, or an excitation domain. The first reconstruction domain may, e.g., be the time domain, the spectral domain, the FFT domain, the MDCT domain, or the excitation domain. The

second reconstruction domain may ,e.g., be the time domain, the spectral domain, the FFT domain, the MDCT domain, or the excitation domain.

5 In an embodiment, the tracing domain may, e.g., be the FFT domain, the first reconstruction domain may, e.g., be the time domain, and the second reconstruction domain may, e.g., be the excitation domain.

10 In another embodiment, the tracing domain may, e.g., be the time domain, the first reconstruction domain may, e.g., be the time domain, and the second reconstruction domain may, e.g., be the excitation domain.

15 According to an embodiment, said first audio signal portion may, e.g., be represented in a first input domain, and said second audio signal portion may, e.g., be represented in a second input domain. The transform unit may, e.g., be a second transform unit. The apparatus may, e.g., further comprise a first transform unit for transforming the second audio signal portion or a value or signal derived from the second audio signal portion from the second input domain to the tracing domain to obtain a second signal portion information. The noise level tracing unit may, e.g., be configured to receive a first signal portion information being represented in the tracing domain, wherein the first signal portion information depends on the first audio signal portion, wherein the noise level tracing unit is configured to receive the second signal portion being represented in the tracing domain, and wherein the noise level tracing unit is configured to determine the noise level information depending on the first signal portion information being represented in the tracing domain and depending on the second signal portion information being represented in the tracing domain.

According to an embodiment, the first input domain may, e.g., be the excitation domain, and the second input domain may, e.g., be the MDCT domain.

30 In another embodiment, the first input domain may, e.g., be the MDCT domain, and wherein the second input domain may, e.g., be the MDCT domain.

35 If, for example, a signal is represented in a time domain, it may, e.g., be represented by time domain samples of the signal. Or, for example, if a signal is represented in a spectral domain, it may, e.g., be represented by spectral samples of a spectrum of the signal.

In an embodiment, the tracing domain may, e.g., be the FFT domain, the first reconstruction domain may, e.g., be the time domain, and the second reconstruction domain may, e.g., be the excitation domain.

- 5 In another embodiment, the tracing domain may, e.g., be the time domain, the first reconstruction domain may, e.g., be the time domain, and the second reconstruction domain may, e.g., be the excitation domain.

10 In some embodiments, the units illustrated in Fig. 14, may, for example, be configured as described for Fig. 1a, 1b, 1c and 1d.

Regarding particular embodiments, in, for example, a low rate mode, an apparatus according to an embodiment may, for example, receive ACELP frames as an input, which are represented in an excitation domain, and which are then transformed to a time domain via LPC synthesis. Moreover, in the low rate mode, the apparatus according to an embodiment may, for example, receive TCX frames as an input, which are represented in an MDCT domain, and which are then transformed to a time domain via an inverse MDCT.

20 Tracing is then conducted in an FFT-Domain, wherein the FFT signal is derived from the time domain signal by conducting an FFT (Fast Fourier Transform). Tracing may, for example, be conducted by conducting a minimum statistics approach, separate for all spectral lines to obtain a comfort noise spectrum.

25 Concealment is then conducted by conducting level derivation based on the comfort noise spectrum. Level derivation is conducted based on the comfort noise spectrum. Level conversion into the time domain is conducted for FD TCX PLC. A fading in the time domain is conducted. A level derivation into the excitation domain is conducted for ACELP PLC and for TD TCX PLC (ACELP like). A fading in the excitation domain is then conducted.

The following list summarizes this:

low rate:

- 35 • input:
- acelp (excitation domain -> time domain, via lpc synthesis)

- o tcx (mdct domain -> time domain, via inverse MDCT)
- tracing:
 - o fft-domain, derived from time domain via FFT
 - o minimum statistics, separate for all spectral lines
-> comfort noise spectrum
- concealment:
 - o level derivation based on the comfort noise spectrum
 - o level conversion into time domain for
 - FD TCX PLC
-> fading in the time domain
 - o level conversion into excitation domain for
 - ACELP PLC
 - TD TCX PLC (ACELP like)
-> fading in the excitation domain

In, for example, a high rate mode, may, for example, receive TCX frames as an input, which are represented in the MDCT domain, and which are then transformed to the time domain via an inverse MDCT.

Tracing may then be conducted in the time domain. Tracing may, for example, be conducted by conducting a minimum statistics approach based on the energy level to obtain a comfort noise level.

For concealment, for FD TCX PLC, the level may be used as is and only a fading in the time domain may be conducted. For TD TCX PLC (ACELP like), level conversion into the excitation domain and fading in the excitation domain is conducted.

The following list summarizes this:

high rate:

- input:
 - o tcx (mdct domain -> time domain, via inverse MDCT)
- tracing:

- o time-domain
 - o minimum statistics on the energy level -> comfort noise level
- concealment:
 - 5 o level usage "as is"
 - FD TCX PLC
 - > fading in the time domain
 - o level conversion into excitation domain for
 - 10 ▪ TD TCX PLC (ACELP like)
 - > fading in the excitation domain

The FFT domain and the MDCT domain are both spectral domains, whereas the excitation domain is some kind of time domain.

- 15 According to an embodiment, the first reconstruction unit 140 may, e.g., be configured to reconstruct the third audio signal portion by conducting a first fading to a noise like spectrum. The second reconstruction unit 141 may, e.g., be configured to reconstruct the fourth audio signal portion by conducting a second fading to a noise like spectrum and/or a second fading of an LTP gain. Moreover, the first reconstruction unit 140 and the
- 20 second reconstruction unit 141 may, e.g., be configured to conduct the first fading and the second fading to a noise like spectrum and/or a second fading of an LTP gain with the same fading speed.

Now adaptive spectral shaping of comfort noise is considered.

- 25 To achieve adaptive shaping to comfort noise during burst packet loss, as a first step, finding appropriate LPC coefficients which represent the background noise may be conducted. These LPC coefficients may be derived during active speech using a minimum statistics approach for finding the background noise spectrum and then calculating LPC
- 30 coefficients from it by using an arbitrary algorithm for LPC derivation known from the literature. Some embodiments, for example, may directly convert the background noise spectrum into a representation which can be used directly for FDNS in the MDCT domain.

- The fading to comfort noise can be done in the ISF domain (also applicable in LSF
- 35 domain; LSF Line spectral frequency):

$$f_{current}[i] = \alpha \cdot f_{last}[i] + (1 - \alpha) \cdot pt_{mean}[i] \quad i = 0..16 \quad (26)$$

by setting pt_{mean} to appropriate LP coefficients describing the comfort noise.

- 5 Regarding the above-described adaptive spectral shaping of the comfort noise, a more general embodiment is illustrated by Fig. 11.

Fig. 11 illustrates an apparatus for decoding an encoded audio signal to obtain a reconstructed audio signal according to an embodiment.

10

The apparatus comprises a receiving interface 1110 for receiving one or more frames, a coefficient generator 1120, and a signal reconstructor 1130.

- 15 The coefficient generator 1120 is configured to determine, if a current frame of the one or more frames is received by the receiving interface 1110 and if the current frame being received by the receiving interface 1110 is not corrupted/erroneous, one or more first audio signal coefficients, being comprised by the current frame, wherein said one or more first audio signal coefficients indicate a characteristic of the encoded audio signal, and one or more noise coefficients indicating a background noise of the encoded audio signal.
- 20 Moreover, the coefficient generator 1120 is configured to generate one or more second audio signal coefficients, depending on the one or more first audio signal coefficients and depending on the one or more noise coefficients, if the current frame is not received by the receiving interface 1110 or if the current frame being received by the receiving interface 1110 is corrupted/erroneous.

25

- The audio signal reconstructor 1130 is configured to reconstruct a first portion of the reconstructed audio signal depending on the one or more first audio signal coefficients, if the current frame is received by the receiving interface 1110 and if the current frame being received by the receiving interface 1110 is not corrupted. Moreover, the audio signal
- 30 reconstructor 1130 is configured to reconstruct a second portion of the reconstructed audio signal depending on the one or more second audio signal coefficients, if the current frame is not received by the receiving interface 1110 or if the current frame being received by the receiving interface 1110 is corrupted.

- 35 Determining a background noise is well known in the art (see, for example, [Mar01]: Rainer Martin, *Noise power spectral density estimation based on optimal smoothing and*

minimum statistics, IEEE Transactions on Speech and Audio Processing 9 (2001), no. 5, 504 –512), and in an embodiment, the apparatus proceeds accordingly.

In some embodiments, the one or more first audio signal coefficients may, e.g., be one or more linear predictive filter coefficients of the encoded audio signal. In some
5 embodiments, the one or more first audio signal coefficients may, e.g., be one or more linear predictive filter coefficients of the encoded audio signal.

It is well known in the art how to reconstruct an audio signal, e.g., a speech signal, from
10 linear predictive filter coefficients or from immittance spectral pairs (see, for example, [3GP09c]: *Speech codec speech processing functions; adaptive multi-rate - wideband (AMRWB) speech codec; transcoding functions*, 3GPP TS 26.190, 3rd Generation Partnership Project, 2009), and in an embodiment, the signal reconstructor proceeds accordingly.

15 According to an embodiment, the one or more noise coefficients may, e.g., be one or more linear predictive filter coefficients indicating the background noise of the encoded audio signal. In an embodiment, the one or more linear predictive filter coefficients may, e.g., represent a spectral shape of the background noise.

20 In an embodiment, the coefficient generator 1120 may, e.g., be configured to determine the one or more second audio signal portions such that the one or more second audio signal portions are one or more linear predictive filter coefficients of the reconstructed audio signal, or such that the one or more first audio signal coefficients are one or more
25 immittance spectral pairs of the reconstructed audio signal.

According to an embodiment, the coefficient generator 1120 may, e.g., be configured to generate the one or more second audio signal coefficients by applying the formula:

$$30 \quad f_{current}[i] = \alpha \cdot f_{last}[i] + (1 - \alpha) \cdot pt_{mean}[i]$$

wherein $f_{current}[i]$ indicates one of the one or more second audio signal coefficients, wherein $f_{last}[i]$ indicates one of the one or more first audio signal coefficients, wherein $pt_{mean}[i]$ is one of the one or more noise coefficients, wherein α is a real number with $0 \leq \alpha$
35 ≤ 1 , and wherein i is an index.

According to an embodiment, $f_{last}[i]$ indicates a linear predictive filter coefficient of the encoded audio signal, and wherein $f_{current}[i]$ indicates a linear predictive filter coefficient of the reconstructed audio signal.

- 5 In an embodiment, $pt_{mean}[i]$ may, e.g., be a linear predictive filter coefficient indicating the background noise of the encoded audio signal.

According to an embodiment, the coefficient generator 1120 may, e.g., be configured to generate at least 10 second audio signal coefficients as the one or more second audio
10 signal coefficients.

In an embodiment, the coefficient generator 1120 may, e.g., be configured to determine, if the current frame of the one or more frames is received by the receiving interface 1110 and if the current frame being received by the receiving interface 1110 is not corrupted,
15 the one or more noise coefficients by determining a noise spectrum of the encoded audio signal.

In the following, fading the MDCT Spectrum to White Noise prior to FDNS Application is considered.
20

Instead of randomly modifying the sign of an MDCT bin (sign scrambling), the complete spectrum is filled with white noise, being shaped using the FDNS. To avoid an instant change in the spectrum characteristics, a cross-fade between sign scrambling and noise filling is applied. The cross fade can be realized as follows:

25

```

for(i=0; i<L_frame; i++) {
    if (old_x[i] != 0) {
        x[i] = (1 - cum_damping)*noise[i] + cum_damping *
random_sign() * x_old[i];
30    }
}

```

where:

- 35 $cum_damping$ is the (absolute) attenuation factor - it decreases from frame to frame, starting from 1 and decreasing towards 0

x_old is the spectrum of the last received frame

random_sign returns 1 or -1

noise contains a random vector (white noise) which is scaled such that its quadratic
5 mean (RMS) is similar to the last good spectrum.

The term `random_sign()*old_x[i]` characterizes the sign-scrambling process to randomize the phases and such avoid harmonic repetitions.

10 Subsequently, another normalization of the energy level might be performed after the cross-fade to make sure that the summation energy does not deviate due to the correlation of the two vectors.

According to embodiments, the first reconstruction unit 140 may, e.g., be configured to
15 reconstruct the third audio signal portion depending on the noise level information and depending on the first audio signal portion. In a particular embodiment, the first reconstruction unit 140 may, e.g., be configured to reconstruct the third audio signal portion by attenuating or amplifying the first audio signal portion.

20 In some embodiments, the second reconstruction unit 141 may, e.g., be configured to reconstruct the fourth audio signal portion depending on the noise level information and depending on the second audio signal portion. In a particular embodiment, the second reconstruction unit 141 may, e.g., be configured to reconstruct the fourth audio signal portion by attenuating or amplifying the second audio signal portion.

25 Regarding the above-described fading of the MDCT Spectrum to white noise prior to the FDNS application, a more general embodiment is illustrated by Fig. 12.

Fig. 12 illustrates an apparatus for decoding an encoded audio signal to obtain a
30 reconstructed audio signal according to an embodiment.

The apparatus comprises a receiving interface 1210 for receiving one or more frames comprising information on a plurality of audio signal samples of an audio signal spectrum of the encoded audio signal, and a processor 1220 for generating the reconstructed audio
35 signal.

The processor 1220 is configured to generate the reconstructed audio signal by fading a modified spectrum to a target spectrum, if a current frame is not received by the receiving

interface 1210 or if the current frame is received by the receiving interface 1210 but is corrupted, wherein the modified spectrum comprises a plurality of modified signal samples, wherein, for each of the modified signal samples of the modified spectrum, an absolute value of said modified signal sample is equal to an absolute value of one of the audio signal samples of the audio signal spectrum.

Moreover, the processor 1220 is configured to not fade the modified spectrum to the target spectrum, if the current frame of the one or more frames is received by the receiving interface 1210 and if the current frame being received by the receiving interface 1210 is not corrupted.

According to an embodiment, the target spectrum is a noise like spectrum.

In an embodiment, the noise like spectrum represents white noise.

According to an embodiment, the noise like spectrum is shaped.

In an embodiment, the shape of the noise like spectrum depends on an audio signal spectrum of a previously received signal.

According to an embodiment, the noise like spectrum is shaped depending on the shape of the audio signal spectrum.

In an embodiment, the processor 1220 employs a tilt factor to shape the noise like spectrum.

According to an embodiment, the processor 1220 employs the formula

$$\text{shaped_noise}[i] = \text{noise} * \text{power}(\text{tilt_factor}, i/N)$$

wherein N indicates the number of samples,

wherein i is an index,

wherein $0 \leq i < N$, with $\text{tilt_factor} > 0$,

wherein power is a power function.

If the `tilt_factor` is smaller 1 this means attenuation with increasing `i`. If the `tilt_factor` is larger 1 means amplification with increasing `i`.

According to another embodiment, the processor 1220 may employ the formula

5

$$\text{shaped_noise}[i] = \text{noise} * (1 + i / (N-1) * (\text{tilt_factor}-1))$$

wherein `N` indicates the number of samples,

10

wherein `i` is an index, wherein $0 \leq i < N$,

with `tilt_factor` > 0.

15

According to an embodiment, the processor 1220 is configured to generate the modified spectrum, by changing a sign of one or more of the audio signal samples of the audio signal spectrum, if the current frame is not received by the receiving interface 1210 or if the current frame being received by the receiving interface 1210 is corrupted.

20

In an embodiment, each of the audio signal samples of the audio signal spectrum is represented by a real number but not by an imaginary number.

According to an embodiment, the audio signal samples of the audio signal spectrum are represented in a Modified Discrete Cosine Transform domain.

25

In another embodiment, the audio signal samples of the audio signal spectrum are represented in a Modified Discrete Sine Transform domain.

30

According to an embodiment, the processor 1220 is configured to generate the modified spectrum by employing a random sign function which randomly or pseudo-randomly outputs either a first or a second value.

In an embodiment, the processor 1220 is configured to fade the modified spectrum to the target spectrum by subsequently decreasing an attenuation factor.

35

According to an embodiment, the processor 1220 is configured to fade the modified spectrum to the target spectrum by subsequently increasing an attenuation factor.

In an embodiment, if the current frame is not received by the receiving interface 1210 or if the current frame being received by the receiving interface 1210 is corrupted, the processor 1220 is configured to generate the reconstructed audio signal by employing the formula:

5

$$x[i] = (1 - \text{cum_damping}) * \text{noise}[i] + \text{cum_damping} * \text{random_sign}() * x_old[i]$$

wherein i is an index, wherein $x[i]$ indicates a sample of the reconstructed audio signal, wherein cum_damping is an attenuation factor, wherein $x_old[i]$ indicates one of the audio signal samples of the audio signal spectrum of the encoded audio signal, wherein $\text{random_sign}()$ returns 1 or -1, and wherein noise is a random vector indicating the target spectrum.

15 Some embodiments continue a TCX LTP operation. In those embodiments, the TCX LTP operation is continued during concealment with the LTP parameters (LTP lag and LTP gain) derived from the last good frame.

The LTP operations can be summarized as:

20

- Feed the LTP delay buffer based on the previously derived output.
- Based on the LTP lag: choose the appropriate signal portion out of the LTP delay buffer that is used as LTP contribution to shape the current signal.
- 25 - Rescale this LTP contribution using the LTP gain.
- Add this rescaled LTP contribution to the LTP input signal to generate the LTP output signal.

30

Different approaches could be considered with respect to the time, when the LTP delay buffer update is performed:

As the first LTP operation in frame n using the output from the last frame $n-1$. This updates the LTP delay buffer in frame n to be used during the LTP processing in frame n .

35

As the last LTP operation in frame n using the output from the current frame n . This updates the LTP delay buffer in frame n to be used during the LTP processing in frame $n + 1$.

- 5 In the following, decoupling of the TCX LTP feedback loop is considered.

Decoupling the TCX LTP feedback loop avoids the introduction of additional noise (resulting from the noise substitution applied to the LPT input signal) during each feedback loop of the LTP decoder when being in concealment mode.

10

Fig. 10 illustrates this decoupling. In particular, Fig. 10 depicts the decoupling of the LTP feedback loop during concealment ($bfi=1$).

- 15 Fig. 10 illustrates a delay buffer 1020, a sample selector 1030, and a sample processor 1040 (the sample processor 1040 is indicated by the dashed line).

Towards the time, when the LTP delay buffer 1020 update is performed, some embodiments proceed as follows:

- 20 - For the normal operation: To update the LTP delay buffer 1020 as the first LTP operation might be preferred, since the summed output signal is usually stored persistently. With this approach, a dedicated buffer can be omitted.
- 25 - For the decoupled operation: To update the LTP delay buffer 1020 as the last LTP operation might be preferred, since the LTP contribution to the signal is usually just stored temporarily. With this approach, the transitorily LTP contribution signal is preserved. Implementation-wise this LTP contribution buffer could just be made persistent.

- 30 Assuming that the latter approach is used in any case (normal operation and concealment), embodiments, may, e.g., implement the following:

- 35 - During normal operation: The time domain signal output of the LTP decoder after its addition to the LTP input signal is used to feed the LTP delay buffer.
- During concealment: The time domain signal output of the LTP decoder prior to its addition to the LTP input signal is used to feed the LTP delay buffer.

Some embodiments fade the TCX LTP gain towards zero. In such embodiment, the TCX LTP gain may, e.g., be faded towards zero with a certain, signal adaptive fade-out factor. This may, e.g., be done iteratively, for example, according to the following pseudo-code:

```

5  gain = gain_past * damping;
    [...]
    gain_past = gain;

```

where:

10

gain is the TCX LTP decoder gain applied in the current frame;

gain_past is the TCX LTP decoder gain applied in the previous frame;

15

damping is the (relative) fade-out factor.

Fig. 1d illustrates an apparatus according to a further embodiment, wherein the apparatus further comprises a long-term prediction unit 170 comprising a delay buffer 180. The long-term prediction unit 170 is configured to generate a processed signal depending on the second audio signal portion, depending on a delay buffer input being stored in the delay buffer 180 and depending on a long-term prediction gain. Moreover, the long-term prediction unit is configured to fade the long-term prediction gain towards zero, if said third frame of the plurality of frames is not received by the receiving interface 110 or if said third frame is received by the receiving interface 110 but is corrupted.

25

In other embodiments (not shown), the long-term prediction unit may, e.g., be configured to generate a processed signal depending on the first audio signal portion, depending on a delay buffer input being stored in the delay buffer and depending on a long-term prediction gain.

30

In Fig. 1d, the first reconstruction unit 140 may, e.g., generate the third audio signal portion furthermore depending on the processed signal.

In an embodiment, the long-term prediction unit 170 may, e.g., be configured to fade the long-term prediction gain towards zero, wherein a speed with which the long-term prediction gain is faded to zero depends on a fade-out factor.

Alternatively or additionally, the long-term prediction unit 170 may, e.g., be configured to update the delay buffer 180 input by storing the generated processed signal in the delay buffer 180 if said third frame of the plurality of frames is not received by the receiving interface 110 or if said third frame is received by the receiving interface 110 but is
5 corrupted.

Regarding the above-described usage of TCX LTP, a more general embodiment is illustrated by Fig. 13.

10 Fig. 13 illustrates an apparatus for decoding an encoded audio signal to obtain a reconstructed audio signal.

The apparatus comprises a receiving interface 1310 for receiving a plurality of frames, a delay buffer 1320 for storing audio signal samples of the decoded audio signal, a sample
15 selector 1330 for selecting a plurality of selected audio signal samples from the audio signal samples being stored in the delay buffer 1320, and a sample processor 1340 for processing the selected audio signal samples to obtain reconstructed audio signal samples of the reconstructed audio signal.

20 The sample selector 1330 is configured to select, if a current frame is received by the receiving interface 1310 and if the current frame being received by the receiving interface 1310 is not corrupted, the plurality of selected audio signal samples from the audio signal samples being stored in the delay buffer 1320 depending on a pitch lag information being comprised by the current frame. Moreover, the sample selector 1330 is configured to
25 select, if the current frame is not received by the receiving interface 1310 or if the current frame being received by the receiving interface 1310 is corrupted, the plurality of selected audio signal samples from the audio signal samples being stored in the delay buffer 1320 depending on a pitch lag information being comprised by another frame being received previously by the receiving interface 1310.

30 According to an embodiment, the sample processor 1340 may, e.g., be configured to obtain the reconstructed audio signal samples, if the current frame is received by the receiving interface 1310 and if the current frame being received by the receiving interface 1310 is not corrupted, by rescaling the selected audio signal samples depending on the
35 gain information being comprised by the current frame. Moreover, the sample selector 1330 may, e.g., be configured to obtain the reconstructed audio signal samples, if the current frame is not received by the receiving interface 1310 or if the current frame being received by the receiving interface 1310 is corrupted, by rescaling the selected audio

signal samples depending on the gain information being comprised by said another frame being received previously by the receiving interface 1310.

5 In an embodiment, the sample processor 1340 may, e.g., be configured to obtain the reconstructed audio signal samples, if the current frame is received by the receiving interface 1310 and if the current frame being received by the receiving interface 1310 is not corrupted, by multiplying the selected audio signal samples and a value depending on the gain information being comprised by the current frame. Moreover, the sample selector 1330 is configured to obtain the reconstructed audio signal samples, if the current frame is
10 not received by the receiving interface 1310 or if the current frame being received by the receiving interface 1310 is corrupted, by multiplying the selected audio signal samples and a value depending on the gain information being comprised by said another frame being received previously by the receiving interface 1310.

15 According to an embodiment, the sample processor 1340 may, e.g., be configured to store the reconstructed audio signal samples into the delay buffer 1320.

In an embodiment, the sample processor 1340 may, e.g., be configured to store the reconstructed audio signal samples into the delay buffer 1320 before a further frame is
20 received by the receiving interface 1310.

According to an embodiment, the sample processor 1340 may, e.g., be configured to store the reconstructed audio signal samples into the delay buffer 1320 after a further frame is received by the receiving interface 1310.

25 In an embodiment, the sample processor 1340 may, e.g., be configured to rescale the selected audio signal samples depending on the gain information to obtain rescaled audio signal samples and by combining the rescaled audio signal samples with input audio signal samples to obtain the processed audio signal samples.

30 According to an embodiment, the sample processor 1340 may, e.g., be configured to store the processed audio signal samples, indicating the combination of the rescaled audio signal samples and the input audio signal samples, into the delay buffer 1320, and to not store the rescaled audio signal samples into the delay buffer 1320, if the current
35 frame is received by the receiving interface 1310 and if the current frame being received by the receiving interface 1310 is not corrupted. Moreover, the sample processor 1340 is configured to store the rescaled audio signal samples into the delay buffer 1320 and to not store the processed audio signal samples into the delay buffer 1320, if the current frame is

not received by the receiving interface 1310 or if the current frame being received by the receiving interface 1310 is corrupted.

5 According to another embodiment, the sample processor 1340 may, e.g., be configured to store the processed audio signal samples into the delay buffer 1320, if the current frame is not received by the receiving interface 1310 or if the current frame being received by the receiving interface 1310 is corrupted.

10 In an embodiment, the sample selector 1330 may, e.g., be configured to obtain the reconstructed audio signal samples by rescaling the selected audio signal samples depending on a modified gain, wherein the modified gain is defined according to the formula:

15
$$\text{gain} = \text{gain_past} * \text{damping};$$

wherein gain is the modified gain, wherein the sample selector 1330 may, e.g., be configured to set gain_past to gain after gain and has been calculated, and wherein damping is a real number.

20 According to an embodiment, the sample selector 1330 may, e.g., be configured to calculate the modified gain.

In an embodiment, damping may, e.g., be defined according to: $0 < \text{damping} < 1$.

25 According to an embodiment, the modified gain gain may, e.g., be set to zero, if at least a predefined number of frames have not been received by the receiving interface 1310 since a frame last has been received by the receiving interface 1310.

30 In the following, the fade-out speed is considered. There are several concealment modules which apply a certain kind of fade-out. While the speed of this fade-out might be differently chosen across those modules, it is beneficial to use the same fade-out speed for all concealment modules for one core (ACELP or TCX). For example:

35 For ACELP, the same fade out speed should be used, in particular, for the adaptive codebook (by altering the gain), and/or for the innovative codebook signal (by altering the gain).

Also, for TCX, the same fade out speed should be used, in particular, for time domain signal, and/or for the LTP gain (fade to zero), and/or for the LPC weighting (fade to one), and/or for the LP coefficients (fade to background spectral shape), and/or for the cross-fade to white noise.

5

It might further be preferable to also use the same fade-out speed for ACELP and TCX, but due to the different nature of the cores it might also be chosen to use different fade-out speeds.

- 10 This fade-out speed might be static, but is preferably adaptive to the signal characteristics. For example, the fade-out speed may, e.g., depend on the LPC stability factor (TCX) and/or on a classification, and/or on a number of consecutively lost frames.

- 15 The fade-out speed may, e.g., be determined depending on the attenuation factor, which might be given absolutely or relatively, and which might also change over time during a certain fade-out.

In embodiments, the same fading speed is used for LTP gain fading as for the white noise fading.

20

An apparatus, method and computer program for generating a comfort noise signal as described above have been provided.

- 25 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.

- 30 The inventive decomposed 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.

- 35 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 disk, a DVD, a CD, a ROM, a PROM, an 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.

5 Some embodiments according to the invention comprise a non-transitory 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.

10 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.

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

20 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.

A further embodiment of the inventive methods is, therefore, a data carrier (or a digital storage medium, or a computer-readable medium) comprising, recorded thereon, the computer program for performing one of the methods described herein.

25 A further embodiment of the inventive 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.

30 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.

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

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,
5 the methods are preferably performed by any hardware apparatus.

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,
10 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 herein.

References

- 5 [3GP09a] 3GPP; Technical Specification Group Services and System Aspects, *Extended adaptive multi-rate - wideband (AMR-WB+) codec*, 3GPP TS 26.290, 3rd Generation Partnership Project, 2009.
- [3GP09b] *Extended adaptive multi-rate - wideband (AMR-WB+) codec; floating-point ANSI-C code*, 3GPP TS 26.304, 3rd Generation Partnership Project, 2009.
- 10 [3GP09c] *Speech codec speech processing functions; adaptive multi-rate - wideband (AMRWB) speech codec; transcoding functions*, 3GPP TS 26.190, 3rd Generation Partnership Project, 2009.
- 15 [3GP12a] *Adaptive multi-rate (AMR) speech codec; error concealment of lost frames (release 11)*, 3GPP TS 26.091, 3rd Generation Partnership Project, Sep 2012.
- 20 [3GP12b] *Adaptive multi-rate (AMR) speech codec; transcoding functions (release 11)*, 3GPP TS 26.090, 3rd Generation Partnership Project, Sep 2012.
[3GP12c] , *ANSI-C code for the adaptive multi-rate - wideband (AMR-WB) speech codec*, 3GPP TS 26.173, 3rd Generation Partnership Project, Sep 2012.
- 25 [3GP12d] *ANSI-C code for the floating-point adaptive multi-rate (AMR) speech codec (release11)*, 3GPP TS 26.104, 3rd Generation Partnership Project, Sep 2012.
- 30 [3GP12e] *General audio codec audio processing functions; Enhanced aacPlus general audio codec; additional decoder tools (release 11)*, 3GPP TS 26.402, 3rd Generation Partnership Project, Sep 2012.
- 35 [3GP12f] *Speech codec speech processing functions; adaptive multi-rate - wideband (amr-wb) speech codec; ansi-c code*, 3GPP TS 26.204, 3rd Generation Partnership Project, 2012.
- [3GP12g] *Speech codec speech processing functions; adaptive multi-rate - wideband (AMR-WB) speech codec; error concealment of erroneous or lost frames*, 3GPP TS 26.191, 3rd Generation Partnership Project, Sep 2012.

- [BJH06] I. Batina, J. Jensen, and R. Heusdens, *Noise power spectrum estimation for speech enhancement using an autoregressive model for speech power spectrum dynamics*, in Proc. IEEE Int. Conf. Acoust., Speech, Signal Process. 3 (2006), 1064–1067.
- [BP06] A. Borowicz and A. Petrovsky, *Minima controlled noise estimation for klt-based speech enhancement*, CD-ROM, 2006, Italy, Florence.
- [Coh03] I. Cohen, *Noise spectrum estimation in adverse environments: Improved minima controlled recursive averaging*, IEEE Trans. Speech Audio Process. 11 (2003), no. 5, 466–475.
- [CPK08] Choong Sang Cho, Nam In Park, and Hong Kook Kim, *A packet loss concealment algorithm robust to burst packet loss for celp-type speech coders*, Tech. report, Korea Electronics Technology Institute, Gwang Institute of Science and Technology, 2008, The 23rd International Technical Conference on Circuits/Systems, Computers and Communications (ITC-CSCC 2008).
- [Dob95] G. Doblinger, *Computationally efficient speech enhancement by spectral minima tracking in subbands*, in Proc. Eurospeech (1995), 1513–1516.
- [EBU10] EBU/ETSI JTC Broadcast, *Digital audio broadcasting (DAB); transport of advanced audio coding (AAC) audio*, ETSI TS 102 563, European Broadcasting Union, May 2010.
- [EBU12] *Digital radio mondiale (DRM); system specification*, ETSI ES 201 980, ETSI, Jun 2012.
- [EH08] Jan S. Erkelens and Richards Heusdens, *Tracking of Nonstationary Noise Based on Data-Driven Recursive Noise Power Estimation*, Audio, Speech, and Language Processing, IEEE Transactions on 16 (2008), no. 6, 1112 – 1123.
- [EM84] Y. Ephraim and D. Malah, *Speech enhancement using a minimum mean-square error short-time spectral amplitude estimator*, IEEE Trans. Acoustics, Speech and Signal Processing 32 (1984), no. 6, 1109–1121.

- [EM85] *Speech enhancement using a minimum mean-square error log-spectral amplitude estimator*, IEEE Trans. Acoustics, Speech and Signal Processing 33 (1985), 443–445.
- 5 [Gan05] S. Gannot, *Speech enhancement: Application of the kalman filter in the estimate-maximize (em framework)*, Springer, 2005.
- 10 [HE95] H. G. Hirsch and C. Ehrlicher, *Noise estimation techniques for robust speech recognition*, Proc. IEEE Int. Conf. Acoustics, Speech, Signal Processing, no. pp. 153-156, IEEE, 1995.
- 15 [HHJ10] Richard C. Hendriks, Richard Heusdens, and Jesper Jensen, *MMSE based noise PSD tracking with low complexity*, Acoustics Speech and Signal Processing (ICASSP), 2010 IEEE International Conference on, Mar 2010, pp. 4266 –4269.
- 20 [HJH08] Richard C. Hendriks, Jesper Jensen, and Richard Heusdens, *Noise tracking using dft domain subspace decompositions*, IEEE Trans. Audio, Speech, Lang. Process. 16 (2008), no. 3, 541–553.
- [IET12] IETF, *Definition of the Opus Audio Codec*, Tech. Report RFC 6716, Internet Engineering Task Force, Sep 2012.
- 25 [ISO09] ISO/IEC JTC1/SC29/WG11, *Information technology – coding of audio-visual objects – part 3: Audio*, ISO/IEC IS 14496-3, International Organization for Standardization, 2009.
- 30 [ITU03] ITU-T, *Wideband coding of speech at around 16 kbit/s using adaptive multi-rate wideband (amr-wb)*, Recommendation ITU-T G.722.2, Telecommunication Standardization Sector of ITU, Jul 2003.
- [ITU05] *Low-complexity coding at 24 and 32 kbit/s for hands-free operation in systems with low frame loss*, Recommendation ITU-T G.722.1, Telecommunication Standardization Sector of ITU, May 2005.
- 35 [ITU06a] *G.722 Appendix III: A high-complexity algorithm for packet loss concealment for G.722*, ITU-T Recommendation, ITU-T, Nov 2006.

- 5 [ITU06b] *G.729.1: G.729-based embedded variable bit-rate coder: An 8-32 kbit/s scalable wideband coder bitstream interoperable with g.729*, Recommendation ITU-T G.729.1, Telecommunication Standardization Sector of ITU, May 2006.
- [ITU07] *G.722 Appendix IV: A low-complexity algorithm for packet loss concealment with G.722*, ITU-T Recommendation, ITU-T, Aug 2007.
- 10 [ITU08a] *G.718: Frame error robust narrow-band and wideband embedded variable bit-rate coding of speech and audio from 8-32 kbit/s*, Recommendation ITU-T G.718, Telecommunication Standardization Sector of ITU, Jun 2008.
- 15 [ITU08b] *G.719: Low-complexity, full-band audio coding for high-quality, conversational applications*, Recommendation ITU-T G.719, Telecommunication Standardization Sector of ITU, Jun 2008.
- [ITU12] *G.729: Coding of speech at 8 kbit/s using conjugate-structure algebraic-code-excited linear prediction (cs-acelp)*, Recommendation ITU-T G.729, Telecommunication Standardization Sector of ITU, June 2012.
- 20 [LS01] Pierre Lauber and Ralph Sperschneider, *Error concealment for compressed digital audio*, Audio Engineering Society Convention 111, no. 5460, Sep 2001.
- 25 [Mar01] Rainer Martin, *Noise power spectral density estimation based on optimal smoothing and minimum statistics*, IEEE Transactions on Speech and Audio Processing 9 (2001), no. 5, 504–512.
- 30 [Mar03] *Statistical methods for the enhancement of noisy speech*, International Workshop on Acoustic Echo and Noise Control (IWAENC2003), Technical University of Braunschweig, Sep 2003.
- 35 [MC99] R. Martin and R. Cox, *New speech enhancement techniques for low bit rate speech coding*, in Proc. IEEE Workshop on Speech Coding (1999), 165–167.

- [MCA99] D. Malah, R. V. Cox, and A. J. Accardi, *Tracking speech-presence uncertainty to improve speech enhancement in nonstationary noise environments*, Proc. IEEE Int. Conf. on Acoustics Speech and Signal Processing (1999), 789–792.
- 5 [MEP01] Nikolaus Meine, Bernd Edler, and Heiko Purnhagen, *Error protection and concealment for HILN MPEG-4 parametric audio coding*, Audio Engineering Society Convention 110, no. 5300, May 2001.
- 10 [MPC89] Y. Mahieux, J.-P. Petit, and A. Charbonnier, *Transform coding of audio signals using correlation between successive transform blocks*, Acoustics, Speech, and Signal Processing, 1989. ICASSP-89., 1989 International Conference on, 1989, pp. 2021–2024 vol.3.
- 15 [NMR+12] Max Neuendorf, Markus Multrus, Nikolaus Rettelbach, Guillaume Fuchs, Julien Robilliard, Jérémie Lecomte, Stephan Wilde, Stefan Bayer, Sascha Disch, Christian Helmrich, Roch Lefebvre, Philippe Gournay, Bruno Bessette, Jimmy Lapierre, Kristopfer Kjörling, Heiko Purnhagen, Lars Villemoes, Werner Oomen, Erik Schuijers, Kei Kikuri, Toru Chinen, Takeshi Norimatsu, Chong Kok Seng, Eunmi Oh, Miyoung Kim, Schuyler Quackenbush, and Berndhard Grill, *MPEG Unified Speech and Audio Coding - The ISO / MPEG Standard for High-Efficiency Audio Coding of all Content Types*, Convention Paper 8654, AES, April 2012, Presented at the 132nd Convention Budapest, Hungary.
- 20 [PKJ+11] Nam In Park, Hong Kook Kim, Min A Jung, Seong Ro Lee, and Seung Ho Choi, *Burst packet loss concealment using multiple codebooks and comfort noise for celp-type speech coders in wireless sensor networks*, Sensors 11 (2011), 5323–5336.
- 25 [QD03] Schuyler Quackenbush and Peter F. Driessen, *Error mitigation in MPEG-4 audio packet communication systems*, Audio Engineering Society Convention 115, no. 5981, Oct 2003.
- 30 [RL06] S. Rangachari and P. C. Loizou, *A noise-estimation algorithm for highly non-stationary environments*, Speech Commun. 48 (2006), 220–231.

- [SFB00] V. Stahl, A. Fischer, and R. Bippus, *Quantile based noise estimation for spectral subtraction and wiener filtering*, in Proc. IEEE Int. Conf. Acoust., Speech and Signal Process. (2000), 1875–1878.
- 5 [SS98] J. Sohn and W. Sung, *A voice activity detector employing soft decision based noise spectrum adaptation*, Proc. IEEE Int. Conf. Acoustics, Speech, Signal Processing, no. pp. 365-368, IEEE, 1998.
- 10 [Yu09] Rongshan Yu, *A low-complexity noise estimation algorithm based on smoothing of noise power estimation and estimation bias correction*, Acoustics, Speech and Signal Processing, 2009. ICASSP 2009. IEEE International Conference on, Apr 2009, pp. 4421–4424.

Claims

1. An apparatus for decoding an encoded audio signal to obtain a reconstructed audio signal, wherein the apparatus comprises:
 - 5 a receiving interface for receiving one or more frames,
 - a coefficient generator, and
 - 10 a signal reconstructor,

wherein the coefficient generator is configured to determine, if a current frame of the one or more frames is received by the receiving interface and if the current frame being received by the receiving interface is not corrupted, one or more first audio signal coefficients, being comprised by the current frame, wherein said one

15 or more first audio signal coefficients indicate a characteristic of the encoded audio signal, and one or more noise coefficients indicating a spectral shape of a background noise of the encoded audio signal,

20 wherein the coefficient generator is configured to generate one or more second audio signal coefficients, depending on the one or more first audio signal coefficients and depending on the one or more noise coefficients, if the current frame is not received by the receiving interface or if the current frame being received by the receiving interface is corrupted,

25 wherein the audio signal reconstructor is configured to reconstruct a first portion of the reconstructed audio signal depending on the one or more first audio signal coefficients, if the current frame is received by the receiving interface and if the current frame being received by the receiving interface is not corrupted, and

30 wherein the audio signal reconstructor is configured to reconstruct a second portion of the reconstructed audio signal depending on the one or more second audio signal coefficients, if the current frame is not received by the receiving interface or if the current frame being received by the receiving interface is

35 corrupted.

2. An apparatus according to claim 1, wherein the one or more first audio signal coefficients are one or more linear predictive filter coefficients of the encoded audio signal.
- 5 3. An apparatus according to claim 2, wherein the one or more linear predictive filter coefficients are represented by one or more immittance spectral pairs or by one or more line spectral pairs, or by one or more immittance spectral frequencies, or by one or more line spectral frequencies of the encoded audio signal.
- 10 4. An apparatus according to any one of claims 1 to 3, wherein the one or more noise coefficients are one or more linear predictive filter coefficients indicating the background noise of the encoded audio signal.
- 15 5. An apparatus according to claim 4, wherein the one or more linear predictive filter coefficients represent a spectral shape of the background noise.
6. An apparatus according to any one of claims 1 to 5, wherein the coefficient generator is configured to determine the one or more second audio signal coefficients such that the one or more second audio signal coefficients are one or more linear predictive filter coefficients of the reconstructed audio signal.
- 20 7. An apparatus according to claim 1,

25 wherein the coefficient generator is configured to generate the one or more second audio signal coefficients by applying the formula:

$$f_{current}[i] = \alpha \cdot f_{last}[i] + (1 - \alpha) \cdot pt_{mean}[i]$$

30 wherein $f_{current}[i]$ indicates one of the one or more second audio signal coefficients,

wherein $f_{last}[i]$ indicates one of the one or more first audio signal coefficients,

wherein $pt_{mean}[i]$ is one of the one or more noise coefficients,

35 wherein α is a real number with $0 \leq \alpha \leq 1$, and

wherein i is an index.

8. An apparatus according to claim 7,
 wherein $f_{last}[i]$ indicates a linear predictive filter coefficient of the encoded audio
 5 signal, and
 wherein $f_{current}[i]$ indicates a linear predictive filter coefficient of the reconstructed
 audio signal.
- 10 9. An apparatus according to claim 8, wherein $pt_{mean}[i]$ indicates the background
 noise of the encoded audio signal.
10. An apparatus according to any one of claims 1 to 9, wherein the coefficient
 15 generator is configured to determine, if the current frame of the one or more
 frames is received by the receiving interface and if the current frame being
 received by the receiving interface is not corrupted, the one or more noise
 coefficients by determining a noise spectrum of the encoded audio signal.
11. An apparatus according to any one of claims 1 to 10, wherein the coefficient
 20 generator is configured to determine LPC coefficients representing background
 noise by using a minimum statistics approach on a signal spectrum to determine a
 background noise spectrum and by calculating the LPC coefficients representing a
 background noise shape from the background noise spectrum.
- 25 12. A method for decoding an encoded audio signal to obtain a reconstructed audio
 signal, wherein the method comprises:
 receiving one or more frames,
 30 determining, if a current frame of the one or more frames is received and if the
 current frame being received is not corrupted, one or more first audio signal
 coefficients, being comprised by the current frame, wherein said one or more first
 audio signal coefficients indicate a characteristic of the encoded audio signal, and
 one or more noise coefficients indicating a spectral shape of a background noise of
 35 the encoded audio signal,
 generating one or more second audio signal coefficients, depending on the one or
 more first audio signal coefficients and depending on the one or more noise

coefficients, if the current frame is not received or if the current frame being received is corrupted,

5 reconstructing a first portion of the reconstructed audio signal depending on the one or more first audio signal coefficients, if the current frame is received and if the current frame being received is not corrupted, and

10 reconstructing a second portion of the reconstructed audio signal depending on the one or more second audio signal coefficients, if the current frame is not received or if the current frame being received is corrupted.

13. A program product comprising a computer readable memory storing computer executable instructions thereon that when executed by a computer perform the method steps of claim 12.




FIG 1A




FIG 1B

3/16




FIG 1C




FIG 1D





FIG 2 (PRIOR ART)



G.722APPV_F07

----- TRANSIENT
-.-.-.- VUV_TRANSITION
_____ other classes

FIG 3 (PRIOR ART)




FIG 4 (PRIOR ART)




FIG 5 (PRIOR ART)




FIG 6





FIG 7



12/16




FIG 9




FIG 10




FIG 11




FIG 12




FIG 13




FIG 14

