

Downscaled Decoding

Description

5 The present application is concerned with a downscaled decoding concept.

The MPEG-4 Enhanced Low Delay AAC (AAC-ELD) usually operates at sample rates up to 48 kHz, which results in an algorithmic delay of 15ms. For some applications, e.g. lip-sync transmission of audio, an even lower delay is desirable. AAC-ELD already provides
10 such an option by operating at higher sample rates, e.g. 96 kHz, and therefore provides operation modes with even lower delay, e.g. 7.5 ms. However, this operation mode comes along with an unnecessary high complexity due to the high sample rate.

The solution to this problem is to apply a downscaled version of the filter bank and
15 therefore, to render the audio signal at a lower sample rate, e.g. 48kHz instead of 96 kHz. The downscaling operation is already part of AAC-ELD as it is inherited from the MPEG-4 AAC-LD codec, which serves as a basis for AAC-ELD.

The question which remains, however, is how to find the downscaled version of a specific
20 filter bank. That is, the only uncertainty is the way the window coefficients are derived whilst enabling clear conformance testing of the downscaled operation modes of the AAC-ELD decoder.

In the following the principles of the down-scaled operation mode of the AAC-(E)LD
25 codecs are described.

The downscaled operation mode or AAC-LD is described for AAC-LD in ISO/IEC 14496-3:2009 in section 4.6.17.2.7 "Adaptation to systems using lower sampling rates" as follows:

30

*"In certain applications it may be necessary to integrate the low delay decoder into an audio system running at lower sampling rates (e.g. 16 kHz) while the nominal sampling rate of the bitstream payload is much higher (e.g. 48 kHz, corresponding to an algorithmic codec delay of approx. 20 ms). In such cases, it is favorable to decode the output of the
35 low delay codec directly at the target sampling rate rather than using an additional sampling rate conversion operation after decoding.*

This can be approximated by appropriate downscaling of both, the frame size and the sampling rate, by some integer factor (e.g. 2, 3), resulting in the same time/frequency resolution of the codec. For example, the codec output can be generated at 16 kHz sampling rate instead of the nominal 48 kHz by retaining only the lowest third (i.e. $480/3 = 160$) of the spectral coefficients prior to the synthesis filterbank and reducing the inverse transform size to one third (i.e. window size $960/3 = 320$).

As a consequence, decoding for lower sampling rates reduces both memory and computational requirements, but may not produce exactly the same output as a full-bandwidth decoding, followed by band limiting and sample rate conversion.

Please note that decoding at a lower sampling rate, as described above, does not affect the interpretation of levels, which refers to the nominal sampling rate of the AAC low delay bitstream payload.”

Please note that AAC-LD works with a standard MDCT framework and two window shapes, i.e. sine-window and low-overlap-window. Both windows are fully described by formulas and therefore, window coefficients for any transformation lengths can be determined.

Compared to AAC-LD, the AAC-ELD codec shows two major differences:

- The Low Delay MDCT window (LD-MDCT)
- The possibility of utilizing the Low Delay SBR tool

The IMDCT algorithm using the low delay MDCT window is described in 4.6.20.2 in [1], which is very similar to the standard IMDCT version using e.g. the sine window. The coefficients of the low delay MDCT windows (480 and 512 samples frame size) are given in Table 4.A.15 and 4.A.16 in [1]. Please note that the coefficients cannot be determined by a formula, as the coefficients are the result of an optimization algorithm. Fig. 9 shows a plot of the window shape for frame size 512.

In case the low delay SBR (LD-SBR) tool is used in conjunction with the AAC-ELD coder, the filter banks of the LD-SBR module are downscaled as well. This ensures that the SBR

module operates with the same frequency resolution and therefore, no more adaptations are required.

Thus, the above description reveals that there is a need for downscaling decoding operations such as, for example, downscaling a decoding at an AAC-ELD. It would be feasible to find out the coefficients for the downsampled synthesis window function anew, but this is a cumbersome task, necessitates additional storage for storing the downsampled version and renders a conformity check between the non-downsampled decoding and the downsampled decoding more complicated or, from another perspective, does not comply with the manner of downscaling requested in the AAC-ELD, for example. Depending on the downscale ratio, i.e. the ratio between the original sampling rate and the downsampled sampling rate, one could derive the downsampled synthesis window function simply by downsampling, i.e. picking out every second, third, ... window coefficient of the original synthesis window function, but this procedure does not result in a sufficient conformity of the non-downsampled decoding and downsampled decoding, respectively. Using more sophisticated decimating procedures applied to the synthesis window function, lead to unacceptable deviations from the original synthesis window function shape. Therefore, there is a need in the art for an improved downsampled decoding concept.

Accordingly, it is an object of the present invention to provide an audio decoding scheme which allows for such an improved downsampled decoding.

This object is achieved by the subject matter of the independent claims.

The present invention is based on the finding that a downsampled version of an audio decoding procedure may more effectively and/or at improved compliance maintenance be achieved if the synthesis window used for downsampled audio decoding is a downsampled version of a reference synthesis window involved in the non-downsampled audio decoding procedure by downsampling by the downsampling factor by which the downsampled sampling rate and the original sampling rate deviate, and downsampled using a segmental interpolation in segments of $1/4$ of the frame length.

Advantageous aspects of the present application are the subject of dependent claims. Preferred embodiments of the present application are described below with respect to the figures, among which:

- Fig. 1 shows a schematic diagram illustrating perfect reconstruction requirements needed to be obeyed when downscaling decoding in order to preserve perfect reconstruction;
- 5 Fig. 2 shows a block diagram of an audio decoder for downsampled decoding according to an embodiment;
- Fig. 3 shows a schematic diagram illustrating in the upper half the manner in which an audio signal has been coded at an original sampling rate into a data stream and, in the lower half separated from the upper half by a dashed horizontal line, a downsampled decoding operation for reconstructing the audio signal from the data stream at a reduced or downsampled sampling rate, so as to illustrate the mode of operation of the audio decoder of Fig. 2;
- 10
- Fig. 4 shows a schematic diagram illustrating the cooperation of the windower and time domain aliasing canceler of Fig. 2;
- 15
- Fig. 5 illustrates a possible implementation for achieving the reconstruction according to Fig. 4 using a special treatment of the zero-weighted portions of the spectral-to-time modulated time portions;
- 20
- Fig. 6 shows a schematic diagram illustrating the downsampling to obtain the downsampled synthesis window;
- 25
- Fig. 7 shows a block diagram illustrating a downsampled operation of AAC-ELD including the low delay SBR tool;
- Fig. 8 shows a block diagram of an audio decoder for downsampled decoding according to an embodiment where modulator, windower and canceller are implemented according to a lifting implementation; and
- 30
- Fig. 9 shows a graph of the window coefficients of a low delay window according to AAC-ELD for 512 sample frame size as an example of a reference synthesis window to be downsampled.
- 35

The following description starts with an illustration of an embodiment for downsampled decoding with respect to the AAC-ELD codec. That is, the following description starts with an embodiment which could form a downsampled mode for AAC-ELD. This description concurrently forms a kind of explanation of the motivation underlying the embodiments of the present application. Later on, this description is generalized, thereby leading to a description of an audio decoder and audio decoding method in accordance with an embodiment of the present application.

As described in the introductory portion of the specification of the present application, AAC-ELD uses low delay MDCT windows. In order to generate downsampled versions thereof, i.e. downsampled low delay windows, the subsequently explained proposal for forming a downsampled mode for AAC-ELD uses a segmental spline interpolation algorithm which maintains the perfect reconstruction property (PR) of the LD-MDCT window with a very high precision. Therefore, the algorithm allows the generation of window coefficients in the direct form, as described in ISO/IEC 14496-3:2009, as well as in the lifting form, as described in [2], in a compatible way. This means both implementations generate 16bit-conform output.

The interpolation of Low Delay MDCT window is performed as follows.

In general a spline interpolation is to be used for generating the downsampled window coefficients to maintain the frequency response and mostly the perfect reconstruction property (around 170dB SNR). The interpolation needs to be constraint in certain segments to maintain the perfect reconstruction property. For the window coefficients c covering the DCT kernel of the transformation (see also Figure 1, $c(1024)..c(2048)$), the following constraint is required,

$$1 = \left| \left(\text{sgn} \cdot c(i) \cdot c(2N - 1 - i) + c(N + i) \cdot c(N - 1 - i) \right) \right| \quad \text{for } i = 0 \dots N/2 - 1 \quad (1)$$

where N denotes the frame size. Some implementation may use different signs to optimize the complexity, here, denoted by sgn . The requirement in (1) can be illustrated by Fig. 1. It should be recalled that simply in even in case of $F=2$, i.e. halvening the sample rate, leaving-out every second window coefficient of the reference synthesis window to obtain the downsampled synthesis window does not fulfil the requirement.

The coefficients $c(0) \dots c(2N - 1)$ are listed along the diamond shape. The $N/4$ zeros in the window coefficients, which are responsible for the delay reduction of the filter bank, are marked using a bold arrow. Fig. 1 shows the dependencies of the coefficients caused by the folding involved in the MDCT and also the points where the interpolation needs to be constraint in order to avoid any undesired dependencies.

- Every $N/2$ coefficient, the interpolation needs to stop to maintain (1)
- Additionally, the interpolation algorithm needs to stop every $N/4$ coefficients due to the inserted zeros. This ensures that the zeros are maintained and the interpolation error is not spread which maintains the PR.

The second constraint is not only required for the segment containing the zeros but also for the other segments. Knowing that some coefficients in the DCT kernel were not determined by the optimization algorithm but were determined by formula (1) to enable PR, several discontinuities in the window shape can be explained, e.g. around $c(1536+128)$ in Figure 1. In order to minimize the PR error, the interpolation needs to stop at such points, which appear in a $N/4$ grid.

Due to that reason, the segment size of $N/4$ is chosen for the segmental spline interpolation to generate the downsampled window coefficients. The source window coefficients are always given by the coefficients used for $N = 512$, also for downscaling operations resulting in frame sizes of $N = 240$ or $N = 120$. The basic algorithm is outlined very briefly in the following as MATLAB code:

```

25  FAC = Downscaling factor    % e.g. 0.5
    sb  = 128;                 % segment size of source window
    w_down = [];               % downsampled window
    nSegments = length(W)/(sb); % number of segments; W=LD window
                                coefficients for N=512
30
    xn = ((0:(FAC*sb-1))+0.5)/FAC-0.5; % spline init
    for i=1:nSegments,
        w_down=[w_down, spline([0:(sb-1)], W((i-1)*sb+(1:(sb))), xn)];
35  end;
```

As the spline function may not be fully deterministic, the complete algorithm is exactly specified in the following section, which may be included into ISO/IEC 14496-3:2009, in order to form an improved downsampled mode in AAC-ELD.

In other words, the following section provides a proposal as to how the above-outlined idea could be applied to ER AAC ELD, i.e. as to how a low-complex decoder could decode a ER AAC ELD bitstream coded at a first data rate at a second data rate lower than the first data rate. It is emphasized however, that the definition of N as used in the following adheres to the standard. Here, N corresponds to the length of the DCT kernel whereas hereinabove, in the claims, and the subsequently described generalized embodiments, N corresponds to the frame length, namely the mutual overlap length of the DCT kernels, i.e. the half of the DCT kernel length. Accordingly, while N was indicated to be 512 hereinabove, for example, it is indicated to be 1024 in the following.

The following paragraphs are proposed for inclusion to 14496-3:2009 via Amendment.

A.0 Adaptation to systems using lower sampling rates

For certain applications, ER AAC LD can change the playout sample rate in order to avoid additional resampling steps (see 4.6.17.2.7). ER AAC ELD can apply similar downscaling steps using the Low Delay MDCT window and the LD-SBR tool. In case AAC-ELD operates with the LD-SBR tool, the downscaling factor is limited to multiples of 2. Without LD-SBR, the downsampled frame size needs to be an integer number.

A.1 Downscaling of Low Delay MDCT window

The LD-MDCT window w_{LD} for $N=1024$ is downsampled by a factor F using a segmental spline interpolation. The number of leading zeros in the window coefficients, i.e. $N/8$, determines the segment size. The downsampled window coefficients $w_{LD,d}$ are used for the inverse MDCT as described in 4.6.20.2 but with a downsampled window length $N_d = N / F$. Please note that the algorithm is also able to generate downsampled lifting coefficients of the LD-MDCT.

```

fs_window_size = 2048; /* Number of fullscale window coefficients. According to ISO/IEC 14496-3:2009,
                        use 2048. For lifting implemenations, please adjust this variable accordingly */
ds_window_size = N * fs_window_size / (1024 * F); /* downsampled window coefficients; N determines the
                                                    transformation length according to 4.6.20.2 */
fs_segment_size = 128;
num_segments    = fs_window_size / fs_segment_size;
ds_segment_size = ds_window_size / num_segments;
tmp[128], y[128]; /* temporary buffers */

/* loop over segments */
for (b = 0; b < num_segments; b++) {
    /* copy current segment to tmp */
    copy(&W_LD[b * fs_segment_size], tmp, fs_segment_size);

    /* apply cubic spline interpolation for downscaling */

```

```

/* calculate interpolating phase */
phase = (fs_window_size - ds_window_size) / (2 * ds_window_size);

/* calculate the coefficients c of the cubic spline given tmp */
5 /* array of precalculated constants */
m = {0.166666672, 0.25, 0.266666681, 0.267857134,
     0.267942578, 0.267948717, 0.267949164};
n = fs_segment_size; /* for simplicity */

10 /* calculate vector r needed to calculate the coefficients c */
for (i = n - 3; i >= 0; i--)
    r[i] = 3 * ((tmp[i + 2] - tmp[i + 1]) - (tmp[i + 1] - tmp[i]));
for (i = 1; i < 7; i++)
    r[i] -= m[i - 1] * r[i - 1];
15 for (i = 7; i < n - 4; i++)
    r[i] -= 0.267949194 * r[i - 1];

/* calculate coefficients c */
c[n - 2] = r[n - 3] / 6;
20 c[n - 3] = (r[n - 4] - c[n - 2]) * 0.25;
for (i = n - 4; i > 7; i--)
    c[i] = (r[i - 1] - c[i + 1]) * 0.267949194;
for (i = 7; i > 1; i--)
    c[i] = (r[i - 1] - c[i + 1]) * m[i - 1];
25 c[1] = r[0] * m[0];
c[0] = 2 * c[1] - c[2];
c[n - 1] = 2 * c[n - 2] - c[n - 3];

/* keep original samples in temp buffer y because samples of
30 tmp will be replaced with interpolated samples */
copy(tmp, y, fs_segment_size);

/* generate downscaled points and do interpolation */
for (k = 0; k < ds_segment_size; k++) {
35     step = phase + k * fs_segment_size / ds_segment_size;
    idx = floor(step);
    diff = step - idx;
    di = (c[idx + 1] - c[idx]) / 3;
    bi = (y[idx + 1] - y[idx]) - (c[idx + 1] + 2 * c[idx]) / 3;
40     /* calculate downscaled values and store in tmp */
    tmp[k] = y[idx] + diff * (bi + diff * (c[idx] + diff * di));
}

/* assemble downscaled window */
45 copy(tmp, &W_LD_d[b * ds_segment_size], ds_segment_size);
}

```

A.2 Downscaling of Low Delay SBR tool

In case the Low Delay SBR tool is used in conjunction with ELD, this tool can be
50 downscaled to lower sample rates, at least for downscaling factors of a multiple of 2. The
downscale factor F controls the number of bands used for the CLDFB analysis and
synthesis filter bank. The following two paragraphs describe a downscaled CLDFB
analysis and synthesis filter bank, see also 4.6.19.4.

4.6.20.5.2.1 Downscaled analyses CLDFB filter bank

- Define number of downscaled CLDFB bands $B = 32/F$.
- Shift the samples in the array x by B positions. The oldest B samples are discarded and B new samples are stored in positions 0 to $B - 1$.

- 5 • Multiply the samples of array x by the coefficient of window ci to get array z . The window coefficients ci are obtained by linear interpolation of the coefficients c , i.e. through the equation

$$ci(i) = \frac{1}{2} [c(2F \cdot i + 1 + p) + c(2F \cdot i + p)], \quad 0 \leq i < (10B), \quad p = \text{int} \left(\frac{64}{2B} - 0.5 \right).$$

The window coefficients of c can be found in Table 4.A.90.

- 10 • Sum the samples to create the $2B$ -element array u :
- $$u(n) = z(n) + z(n + 2B) + z(n + 4B) + z(n + 6B) + z(n + 8B), \quad 0 \leq n < (2B).$$
- Calculate B new subband samples by the matrix operation Mu , where

$$M(k, n) = 2 \cdot \exp \left(\frac{j \cdot \pi \cdot (k+0.5) \cdot (2n - (3B-1))}{2B} \right), \quad \begin{cases} 0 \leq k < B \\ 0 \leq n < 2B \end{cases}$$

15

In the equation, $\exp()$ denotes the complex exponential function and j is the imaginary unit.

4.6.20.5.2.2 Downscaled synthesis CLDFB filter bank

- 20 • Define number of downscaled CLDFB bands $B = 64/F$.
- Shift the samples in the array v by $2B$ positions. The oldest $2B$ samples are discarded.
 - The B new complex-valued subband samples are multiplied by the matrix N , where

$$N(k, n) = \frac{1}{64} \cdot \exp \left(\frac{j \cdot \pi \cdot (k+0.5) \cdot (2 \cdot n - (B-1))}{2B} \right), \quad \begin{cases} 0 \leq k < B \\ 0 \leq n < 2B \end{cases}$$

25

In the equation, $\exp()$ denotes the complex exponential function and j is the imaginary unit. The real part of the output from this operation is stored in the positions 0 to $2B - 1$ of array v .

- Extract samples from v to create the $10B$ -element array g .

30

$$\begin{aligned} g(2B \cdot n + k) &= v(4B \cdot n + k) \\ g(2B \cdot n + B + k) &= v(4B \cdot n + 3B + k) \end{aligned} \quad \begin{cases} 0 \leq n \leq 4 \\ 0 \leq k < B \end{cases}$$

- Multiply the samples of array g by the coefficient of window ci to produce array w . The window coefficients ci are obtained by linear interpolation of the coefficients c , i.e. through the equation

$$ci(i) = \frac{1}{2} [c(2F \cdot i + 1 + p) + c(2F \cdot i + p)], \quad 0 \leq i < (10B), p = \text{int} \left(\frac{64}{2B} - 0.5 \right).$$

The window coefficients of c can be found in Table 4.A.90.

- Calculate B new output samples by summation of samples from array w according to $\text{output}(n) = \sum_{i=0}^{i \leq 9} w(Bi + n)$, $0 \leq n < B$.

10

Please note that setting $F = 2$ provides the downsampled synthesis filter bank according to 4.6.19.4.3. Therefore, to process a downsampled LD-SBR bit stream with an additional downscale factor F , F needs to be multiplied by 2.

15 4.6.20.5.2.3 Downscaled real-valued CLDFB filter bank

The downscaling of the CLDFB can be applied for the real valued versions of the low power SBR mode as well. For illustration, please also consider 4.6.19.5.

For the downscaled real-valued analysis and synthesis filter bank, follow the description in 4.6.20.5.2.1 and 4.6.20.2.2 and exchange the $\exp()$ modulator in M by a $\cos()$ modulator.

20

A.3 Low Delay MDCT Analysis

This subclause describes the Low Delay MDCT filter bank utilized in the AAC ELD encoder. The core MDCT algorithm is mostly unchanged, but with a longer window, such that n is now running from $-N$ to $N-1$ (rather than from 0 to $N-1$)

- 25 The spectral coefficient, $X_{i,k}$, are defined as follows:

$$X_{i,k} = -2 \cdot \sum_{n=-N}^{N-1} z_{i,n} \cos \left(\frac{2\pi}{N} (n + n_0) \left(k + \frac{1}{2} \right) \right) \quad \text{for } 0 \leq k < N/2$$

where:

- z_{in} = windowed input sequence
- N = sample index
- K = spectral coefficient index
- I = block index
- N = window length
- n_0 = $(-N / 2 + 1) / 2$

30

The window length N (based on the sine window) is 1024 or 960.

The window length of the low-delay window is 2*N. The windowing is extended to the past in the following way:

$$z_{i,n} = w_{LD}(N-1-n) \cdot x_{i,n}^1$$

for $n=-N, \dots, N-1$, with the synthesis window w used as the analysis window by inverting the order.

10 **A.4 Low Delay MDCT Synthesis**

The synthesis filter bank is modified compared to the standard IMDCT algorithm using a sine window in order to adopt a low-delay filter bank. The core IMDCT algorithm is mostly unchanged, but with a longer window, such that n is now running up to $2N-1$ (rather than up to $N-1$).

$$x_{i,n} = -\frac{2}{N} \sum_{k=0}^{\frac{N}{2}-1} \text{spec}[i][k] \cos\left(\frac{2\pi}{N}(n+n_0)\left(k+\frac{1}{2}\right)\right) \quad \text{for } 0 \leq n < 2N$$

where:

- n = sample index
- i = window index
- k = spectral coefficient index
- N = window length / twice the frame length
- n_0 = $(-N/2 + 1)/2$

20 with $N = 960$ or 1024 .

The windowing and overlap-add is conducted in the following way:

25 The length N window is replaced by a length $2N$ window with more overlap in the past, and less overlap to the future ($N/8$ values are actually zero).

Windowing for the Low Delay Window:

$$z_{i,n} = w_{LD}(n) \cdot x_{i,n}$$

30

Where the window now has a length of $2N$, hence $n=0, \dots, 2N-1$.

Overlap and add:

$$out_{i,n} = z_{i,n} + z_{i-1,n+\frac{N}{2}} + z_{i-2,n+N} + z_{i-3,n+N+\frac{N}{2}}$$

for $0 \leq n < N/2$

Here, the paragraphs proposed for being included into 14496-3:2009 via amendment end.

10

Naturally, the above description of a possible downsampled mode for AAC-ELD merely represents one embodiment of the present application and several modifications are feasible. Generally, embodiments of the present application are not restricted to an audio decoder performing a downsampled version of AAC-ELD decoding. In other words, 15
embodiments of the present application may, for instance, be derived by forming an audio decoder capable of performing the inverse transformation process in a downsampled manner only without supporting or using the various AAC-ELD specific further tasks such as, for instance, the scale factor-based transmission of the spectral envelope, TNS (temporal noise shaping) filtering, spectral band replication (SBR) or the like.

20

Subsequently, a more general embodiment for an audio decoder is described. The above-outlined example for an AAC-ELD audio decoder supporting the described downsampled mode could thus represent an implementation of the subsequently described audio decoder. In particular, the subsequently explained decoder is shown in Fig. 2 while Fig. 3 25
illustrates the steps performed by the decoder of Fig. 2.

30

The audio decoder of Fig. 2, which is generally indicated using reference sign 10, comprises a receiver 12, a grabber 14, a spectral-to-time modulator 16, a windower 18 and a time domain aliasing canceler 20, all of which are connected in series to each other in the order of their mentioning. The interaction and functionality of blocks 12 to 20 of audio decoder 10 are described in the following with respect to Fig. 3. As described at the end of the description of the present application, blocks 12 to 20 may be implemented in software, programmable hardware or hardware such as in the form of a computer program, an FPGA or appropriately programmed computer, programmed microprocessor

or application specific integrated circuit with the blocks 12 to 20 representing respective subroutines, circuit paths or the like.

In a manner outlined in more details below, the audio decoder 10 of Fig. 2 is configured to, – and the elements of the audio decoder 10 are configured to appropriately cooperate – in order to decode an audio signal 22 from a data stream 24 with a noteworthiness that audio decoder 10 decodes signal 22 at a sampling rate being $1/F^{\text{th}}$ of the sampling rate at which the audio signal 22 has been transform coded into data stream 24 at the encoding side. F may, for instance, be any rational number greater than one. The audio decoder may be configured to operate at different or varying downscaling factors F or at a fixed one. Alternatives are described in more detail below.

The manner in which the audio signal 22 is transform coded at the encoding or original sampling rate into the data stream is illustrated in Fig. 3 in the upper half. At 26 Fig. 3 illustrates the spectral coefficients using small boxes or squares 28 arranged in a spectrotemporal manner along a time axis 30 which runs horizontally in Fig. 3, and a frequency axis 32 which runs vertically in Fig. 3, respectively. The spectral coefficients 28 are transmitted within data stream 24. The manner in which the spectral coefficients 28 have been obtained, and thus the manner via which the spectral coefficients 28 represent the audio signal 22, is illustrated in Fig. 3 at 34, which illustrates for a portion of time axis 30 how the spectral coefficients 28 belonging to, or representing the respective time portion, have been obtained from the audio signal.

In particular, coefficients 28 as transmitted within data stream 24 are coefficients of a lapped transform of the audio signal 22 so that the audio signal 22, sampled at the original or encoding sampling rate, is partitioned into immediately temporally consecutive and non-overlapping frames of a predetermined length N, wherein N spectral coefficients are transmitted in data stream 24 for each frame 36. That is, transform coefficients 28 are obtained from the audio signal 22 using a critically sampled lapped transform. In the spectrotemporal spectrogram representation 26, each column of the temporal sequence of columns of spectral coefficients 28 corresponds to a respective one of frames 36 of the sequence of frames. The N spectral coefficients 28 are obtained for the corresponding frame 36 by a spectrally decomposing transform or time-to-spectral modulation, the modulation functions of which temporally extend, however, not only across the frame 36 to which the resulting spectral coefficients 28 belong, but also across E + 1 previous frames, wherein E may be any integer or any even numbered integer greater than zero. That is,

the spectral coefficients 28 of one column of the spectrogram at 26 which belonged to a certain frame 36 are obtained by applying a transform onto a transform window, which in addition the respective frame comprises $E + 1$ frames lying in the past relative to the current frame. The spectral decomposition of the samples of the audio signal within this transform window 38, which is illustrated in Fig. 3 for the column of transform coefficients 28 belonging to the middle frame 36 of the portion shown at 34 is achieved using a low delay unimodal analysis window function 40 using which the spectral samples within the transform window 38 are weighted prior to subjecting same to an MDCT or MDST or other spectral decomposition transform. In order to lower the encoder-side delay, the analysis window 40 comprises a zero-interval 42 at the temporal leading end thereof so that the encoder does not need to await the corresponding portion of newest samples within the current frame 36 so as to compute the spectral coefficients 28 for this current frame 36. That is, within the zero-interval 42 the low delay window function 40 is zero or has zero window coefficients so that the co-located audio samples of the current frame 36 do not, owing to the window weighting 40, contribute to the transform coefficients 28 transmitted for that frame and a data stream 24. That is, summarizing the above, transform coefficients 28 belonging to a current frame 36 are obtained by windowing and spectral decomposition of samples of the audio signal within a transform window 38 which comprises the current frame as well as temporally preceding frames and which temporally overlaps with the corresponding transform windows used for determining the spectral coefficients 28 belonging to temporally neighboring frames.

Before resuming the description of the audio decoder 10, it should be noted that the description of the transmission of the spectral coefficients 28 within the data stream 24 as provided so far has been simplified with respect to the manner in which the spectral coefficients 28 are quantized or coded into data stream 24 and/or the manner in which the audio signal 22 has been pre-processed before subjecting the audio signal to the lapped transform. For example, the audio encoder having transform coded audio signal 22 into data stream 24 may be controlled via a psychoacoustic model or may use a psychoacoustic model to keep the quantization noise and quantizing the spectral coefficients 28 unperceivable for the hearer and/or below a masking threshold function, thereby determining scale factors for spectral bands using which the quantized and transmitted spectral coefficients 28 are scaled. The scale factors would also be signaled in data stream 24. Alternatively, the audio encoder may have been a TCX (transform coded excitation) type of encoder. Then, the audio signal would have had subject to a linear prediction analysis filtering before forming the spectrotemporal representation 26 of

spectral coefficients 28 by applying the lapped transform onto the excitation signal, i.e. the linear prediction residual signal. For example, the linear prediction coefficients could be signaled in data stream 24 as well, and a spectral uniform quantization could be applied in order to obtain the spectral coefficients 28.

5

Furthermore, the description brought forward so far has also been simplified with respect to the frame length of frames 36 and/or with respect to the low delay window function 40. In fact, the audio signal 22 may have been coded into data stream 24 in a manner using varying frame sizes and/or different windows 40. However, the description brought
10 forward in the following concentrates on one window 40 and one frame length, although the subsequent description may easily be extended to a case where the entropy encoder changes these parameters during coding the audio signal into the data stream.

Returning back to the audio decoder 10 of Fig. 2 and its description, receiver 12 receives
15 data stream 24 and receives thereby, for each frame 36, N spectral coefficients 28, i.e. a respective column of coefficients 28 shown in Fig. 3. It should be recalled that the temporal length of the frames 36, measured in samples of the original or encoding sampling rate, is N as indicated in Fig. 3 at 34, but the audio decoder 10 of Fig. 2 is configured to decode the audio signal 22 at a reduced sampling rate. The audio decoder
20 10 supports, for example, merely this downscaled decoding functionality described in the following. Alternatively, audio decoder 10 would be able to reconstruct the audio signal at the original or encoding sampling rate, but may be switched between the downscaled decoding mode and a non-downscaled decoding mode with the downscaled decoding mode coinciding with the audio decoder's 10 mode of operation as subsequently
25 explained. For example, audio encoder 10 could be switched to a downscaled decoding mode in the case of a low battery level, reduced reproduction environment capabilities or the like. Whenever the situation changes the audio decoder 10 could, for instance, switch back from the downscaled decoding mode to the non-downscaled one. In any case, in accordance with the downscaled decoding process of decoder 10 as described in the
30 following, the audio signal 22 is reconstructed at a sampling rate at which frames 36 have, at the reduced sampling rate, a lower length measured in samples of this reduced sampling rate, namely a length of N/F samples at the reduced sampling rate.

The output of receiver 12 is the sequence of N spectral coefficients, namely one set of N
35 spectral coefficients, i.e. one column in Fig. 3, per frame 36. It already turned out from the above brief description of the transform coding process for forming data stream 24 that

receiver 12 may apply various tasks in obtaining the N spectral coefficients per frame 36. For example, receiver 12 may use entropy decoding in order to read the spectral coefficients 28 from the data stream 24. Receiver 12 may also spectrally shape the spectral coefficients read from the data stream with scale factors provided in the data stream and/or scale factors derived by linear prediction coefficients conveyed within data stream 24. For example, receiver 12 may obtain scale factors from the data stream 24, namely on a per frame and per subband basis, and use these scale factors in order to scale the scale factors conveyed within the data stream 24. Alternatively, receiver 12 may derive scale factors from linear prediction coefficients conveyed within the data stream 24, for each frame 36, and use these scale factors in order to scale the transmitted spectral coefficients 28. Optionally, receiver 12 may perform gap filling in order to synthetically fill zero-quantized portions within the sets of N spectral coefficients 18 per frame. Additionally or alternatively, receiver 12 may apply a TNS-synthesis filter onto a transmitted TNS filter coefficient per frame to assist the reconstruction of the spectral coefficients 28 from the data stream with the TNS coefficients also being transmitted within the data stream 24. The just outlined possible tasks of receiver 12 shall be understood as a non-exclusive list of possible measures and receiver 12 may perform further or other tasks in connection with the reading of the spectral coefficients 28 from data stream 24.

Grabber 14 thus receives from receiver 12 the spectrogram 26 of spectral coefficients 28 and grabs, for each frame 36, a low frequency fraction 44 of the N spectral coefficients of the respective frame 36, namely the N/F lowest-frequency spectral coefficients.

That is, spectral-to-time modulator 16 receives from grabber 14 a stream or sequence 46 of N/F spectral coefficients 28 per frame 36, corresponding to a low-frequency slice out of the spectrogram 26, spectrally registered to the lowest frequency spectral coefficients illustrated using index "0" in Fig. 3, and extending till the spectral coefficients of index $N/F - 1$.

The spectral-to-time modulator 16 subjects, for each frame 36, the corresponding low-frequency fraction 44 of spectral coefficients 28 to an inverse transform 48 having modulation functions of length $(E + 2) \cdot N/F$ temporally extending over the respective frame and $E + 1$ previous frames as illustrated at 50 in Fig. 3, thereby obtaining a temporal portion of length $(E + 2) \cdot N/F$, i.e. a not-yet windowed time segment 52. That is, the spectral-to-time modulator may obtain a temporal time segment of $(E + 2) \cdot N/F$ samples of reduced sampling rate by weighting and summing modulation functions of the

same length using, for instance, the first formulae of the proposed replacement section A.4 indicated above. The newest N/F samples of time segment 52 belong to the current frame 36. The modulation functions may, as indicated, be cosine functions in case of the inverse transform being an inverse MDCT, or sine functions in case of the inverse transform being an inverse MDCT, for instance.

Thus, windower 52 receives, for each frame, a temporal portion 52, the N/F samples at the leading end thereof temporally corresponding to the respective frame while the other samples of the respective temporal portion 52 belong to the corresponding temporally preceding frames. Windower 18 windows, for each frame 36, the temporal portion 52 using a unimodal synthesis window 54 of length $(E + 2) \cdot N/F$ comprising a zero-portion 56 of length $1/4 \cdot N/F$ at a leading end thereof, i.e. $1/F \cdot N/F$ zero-valued window coefficients, and having a peak 58 within its temporal interval succeeding, temporally, the zero-portion 56, i.e. the temporal interval of temporal portion 52 not covered by the zero-portion 52. The latter temporal interval may be called the non-zero portion of window 58 and has a length of $7/4 \cdot N/F$ measured in samples of the reduced sampling rate, i.e. $7/4 \cdot N/F$ window coefficients. The windower 18 weights, for instance, the temporal portion 52 using window 58. This weighting or multiplying 58 of each temporal portion 52 with window 54 results in a windowed temporal portion 60, one for each frame 36, and coinciding with the respective temporal portion 52 as far as the temporal coverage is concerned. In the above proposed section A.4, the windowing processing which may be used by window 18 is described by the formulae relating $z_{i,n}$ to $x_{i,n}$, where $x_{i,n}$ corresponds to the aforementioned temporal portions 52 not yet windowed and $z_{i,n}$ corresponds to the windowed temporal portions 60 with i indexing the sequence of frames/windows, and n indexing, within each temporal portion 52/60, the samples or values of the respective portions 52/60 in accordance with a reduced sampling rate.

Thus, the time domain aliasing canceler 20 receives from windower 18 a sequence of windowed temporal portions 60, namely one per frame 36. Canceler 20 subjects the windowed temporal portions 60 of frames 36 to an overlap-add process 62 by registering each windowed temporal portion 60 with its leading N/F values to coincide with the corresponding frame 36. By this measure, a trailing-end fraction of length $(E + 1)/(E + 2)$ of the windowed temporal portion 60 of a current frame, i.e. the remainder having length $(E + 1) \cdot N/F$, overlaps with a corresponding equally long leading end of the temporal portion of the immediately preceding frame. In formulae, the time domain aliasing canceler 20 may operate as shown in the last formula of the above proposed version of section

A.4, where $out_{i,n}$ corresponds to the audio samples of the reconstructed audio signal 22 at the reduced sampling rate.

The processes of windowing 58 and overlap-adding 62 as performed by windower 18 and time domain aliasing canceler 20 are illustrated in more detail below with respect to Fig. 4. Fig. 4 uses both the nomenclature applied in the above-proposed section A.4 and the reference signs applied in Figs. 3 and 4. $x_{0,0}$ to $x_{0,(E+2) \cdot N/F-1}$ represents the 0th temporal portion 52 obtained by the spatial-to-temporal-modulator 16 for the 0th frame 36. The first index of x indexes the frames 36 along the temporal order, and the second index of x orders the samples of the temporal along the temporal order, the inter-sample pitch belonging to the reduced sample rate. Then, in Fig. 4, w_0 to $w_{(E+2) \cdot N/F-1}$ indicate the window coefficients of window 54. Like the second index of x , i.e. the temporal portion 52 as output by modulator 16, the index of w is such that index 0 corresponds to the oldest and index $(E + 2) \cdot N/F - 1$ corresponds to the newest sample value when the window 54 is applied to the respective temporal portion 52. Windower 18 windows the temporal portion 52 using window 54 to obtain the windowed temporal portion 60 so that $z_{0,0}$ to $z_{0,(E+2) \cdot N/F-1}$, which denotes the windowed temporal portion 60 for the 0th frame, is obtained according to $z_{0,0} = x_{0,0} \cdot w_0, \dots, z_{0,(E+2) \cdot N/F-1} = x_{0,(E+2) \cdot N/F-1} \cdot w_{(E+2) \cdot N/F-1}$. The indices of z have the same meaning as for x . In this manner, modulator 16 and windower 18 act for each frame indexed by the first index of x and z . Canceler 20 sums up $E + 2$ windowed temporal portions 60 of $E + 2$ immediately consecutive frames with offsetting the samples of the windowed temporal portions 60 relative to each other by one frame, i.e. by the number of samples per frame 36, namely N/F , so as to obtain the samples u of one current frame, here $u_{-(E+1),0} \dots u_{-(E+1),N/F-1}$. Here, again, the first index of u indicates the frame number and the second index orders the samples of this frame along the temporal order. The canceller joins the reconstructed frames thus obtained so that the samples of the reconstructed audio signal 22 within the consecutive frames 36 follow each other according to $u_{-(E+1),0} \dots u_{-(E+1),N/F-1}, u_{-E,0}, \dots u_{-E,N/F-1}, u_{-(E-1),0}, \dots$. the canceller 22 computes each sample of the audio signal 22 within the $-(E+1)^{th}$ frame according to $u_{-(E+1),0} = z_{0,0} + z_{-1,N/F} + \dots z_{-(E+1),(E+1) \cdot N/F}, \dots, u_{-(E+1),N/F-1} = z_{0,N/F-1} + z_{-1,2 \cdot N/F-1} + \dots + z_{-(E+1),(E+2) \cdot N/F-1}$, i.e. summing up $(e+2)$ addends per samples u of the current frame.

Fig. 5 illustrates a possible exploitation of the fact that, among the just windowed samples contributing to the audio samples u of frame $-(E + 1)$, the ones corresponding to, or having been windowed using, the zero-portion 56 of window 54, namely $z_{-(E+1),(E+7/4) \cdot N/F} \dots z_{-(E+1),(E+2) \cdot N/F-1}$ are zero valued. Thus, instead of obtaining all N/F samples within the -

(E+1)th frame 36 of the audio signal u using E+2 addends, canceler 20 may compute the leading end quarter thereof, namely $u_{-(E+1),(E+7/4) \cdot N/F} \dots u_{-(E+1),(E+2) \cdot N/F-1}$ merely using E+1 addends according to $u_{-(E+1),(E+7/4) \cdot N/F} = Z_{0,3/4 \cdot N/F} + Z_{-1,7/4 \cdot N/F} + \dots + Z_{-E,(E+3/4) \cdot N/F}$, ..., $u_{-(E+1),(E+2) \cdot N/F-1} = Z_{0,N/F-1} + Z_{-1,2 \cdot N/F-1} + \dots + Z_{-E,(E+1) \cdot N/F-1}$. In this manner, the windower could even leave out, effectively, the performance of the weighting 58 with respect to the zero-portion 56. Samples $u_{-(E+1),(E+7/4) \cdot N/F} \dots u_{-(E+1),(E+2) \cdot N/F-1}$ of current $-(E+1)^{\text{th}}$ frame would, thus, be obtained using E+1 addends only, while $u_{-(E+1),(E+1) \cdot N/F} \dots u_{-(E+1),(E+7/4) \cdot N/F-1}$ would be obtained using E+2 addends.

Thus, in the manner outlined above, the audio decoder 10 of Fig. 2 reproduces, in a downsampled manner, the audio signal coded into data stream 24. To this end, the audio decoder 10 uses a window function 54 which is itself a downsampled version of a reference synthesis window of length $(E+2) \cdot N$. As explained with respect to Fig. 6, this downsampled version, i.e. window 54, is obtained by downsampling the reference synthesis window by a factor of F, i.e. the downsampling factor, using a segmental interpolation, namely in segments of length $1/4 \cdot N$ when measured in the not yet downsampled regime, in segments of length $1/4 \cdot N/F$ in the downsampled regime, in segments of quarters of a frame length of frames 36, measured temporally and expressed independently from the sampling rate. In $4 \cdot (E+2)$ the interpolation is, thus, performed, thus yielding $4 \cdot (E+2)$ times $1/4 \cdot N/F$ long segments which, concatenated, represent the downsampled version of the reference synthesis window of length $(E+2) \cdot N$. See Fig. 6 for illustration. Fig. 6 shows the synthesis window 54 which is unimodal and used by the audio decoder 10 in accordance with a downsampled audio decoding procedure underneath the reference synthesis window 70 which has of length $(E+2) \cdot N$. That is, by the downsampling procedure 72 leading from the reference synthesis window 70 to the synthesis window 54 actually used by the audio decoder 10 for downsampled decoding, the number of window coefficients is reduced by a factor of F. In Fig. 6, the nomenclature of Figs. 5 and 6 has been adhered to, i.e. w is used in order to denote the downsampled version window 54, while w' has been used to denote the window coefficients of the reference synthesis window 70.

As just mentioned, in order to perform the downsampling 72, the reference synthesis window 70 is processed in segments 74 of equal length. In number, there are $(E+2) \cdot 4$ such segments 74. Measured in the original sampling rate, i.e. in the number of window coefficients of the reference synthesis window 70, each segment 74 is $1/4 \cdot N$ window

coefficients w' long, and measured in the reduced or downsampled sampling rate, each segment 74 is $1/4 \cdot N/F$ window coefficients w long.

5 Naturally, it would be possible to perform the downsampling 72 for each downsampled window coefficient w_i coinciding accidentally with any of the window coefficients w'_j of the reference synthesis window 70 by simply setting $w_i = w'_j$ with the sample time of w_i coinciding with that of w'_j , and/or by linearly interpolating any window coefficients w_i residing, temporally, between two window coefficients w'_j and w'_{j+2} by linear interpolation, but this procedure would result in a poor approximation of the reference synthesis window 10 70, i.e. the synthesis window 54 used by audio decoder 10 for the downsampled decoding would represent a poor approximation of the reference synthesis window 70, thereby not fulfilling the request for guaranteeing conformance testing of the downsampled decoding relative to the non-downsampled decoding of the audio signal from data stream 24. Thus, the downsampling 72 involves an interpolation procedure according to which the majority 15 of the window coefficients w_i of the downsampled window 54, namely the ones positioned offset from the borders of segments 74, depend by way of the downsampling procedure 72 on more than two window coefficients w' of the reference window 70. In particular, while the majority of the window coefficients w_i of the downsampled window 54 depend on more than two window coefficients w'_j of the reference window 70 in order to increase the 20 quality of the interpolation/downsampling result, i.e. the approximation quality, for every window coefficient w_i of the downsampled version 54 it holds true that same does not depend in window coefficients w'_j belonging to different segments 74. Rather, the downsampling procedure 72 is a segmental interpolation procedure.

25 For example, the synthesis window 54 may be a concatenation of spline functions of length $1/4 \cdot N/F$. Cubic spline functions may be used. Such an example has been outlined above in section A.1 where the outer for-next loop sequentially looped over segments 74 wherein, in each segment 74, the downsampling or interpolation 72 involved a mathematical combination of consecutive window coefficients w' within the current 30 segment 74 at, for example, the first for next clause in the section "calculate vector r needed to calculate the coefficients c ". The interpolation applied in segments, may, however, also be chosen differently. That is, the interpolation is not restricted to splines or cubic splines. Rather, linear interpolation or any other interpolation method may be used as well. In any case, the segmental implementation of the interpolation would cause the 35 computation of samples of the downsampled synthesis window, i.e. the outmost samples of

the segments of the downscaled synthesis window, neighboring another segment, to not depend on window coefficients of the reference synthesis window residing in different segments.

It may be that windower 18 obtains the downsampled synthesis window 54 from a storage where the window coefficients w_i of this downsampled synthesis window 54 have been stored after having been obtained using the downsampling 72. Alternatively, as illustrated in Fig. 2, the audio decoder 10 may comprise a segmental downsampler 76 performing the downsampling 72 of Fig. 6 on the basis of the reference synthesis window 70.

It should be noted that the audio decoder 10 of Fig. 2 may be configured to support merely one fixed downsampling factor F or may support different values. In that case, the audio decoder 10 may be responsive to an input value for F as illustrated in Fig. 2 at 78. The grabber 14, for instance, may be responsive to this value F in order to grab, as mentioned above, the N/F spectral values per frame spectrum. In a like manner, the optional segmental downsampler 76 may also be responsive to this value of F and operate as indicated above. The S/T modulator 16 may be responsive to F either in order to, for example, computationally derive downscaled/downsampled versions of the modulation functions, downscaled/downsampled relative to the ones used in not-downscaled operation mode where the reconstruction leads to the full audio sample rate.

Naturally, the modulator 16 would also be responsive to F input 78, as modulator 16 would use appropriately downsampled versions of the modulation functions and the same holds true for the windower 18 and canceler 20 with respect to an adaptation of the actual length of the frames in the reduced or downsampled sampling rate.

For example, F may lie between 1.5 and 10, both inclusively.

It should be noted that the decoder of Fig. 2 and 3 or any modification thereof outlined herein, may be implemented so as to perform the spectral-to-time transition using a lifting implementation of the Low Delay MDCT as taught in, for example, EP 2 378 516 B1.

Fig. 8 illustrates an implementation of the decoder using the lifting concept. The S/T modulator 16 performs exemplarily an inverse DCT-IV and is shown as followed by a block representing the concatenation of the windower 18 and the time domain aliasing canceller 20. In the example of Fig. 8 E is 2, i.e. $E=2$.

The modulator 16 comprises an inverse type-iv discrete cosine transform frequency/time converter. Instead of outputting sequences of $(E+2)N/F$ long temporal portions 52, it merely outputs temporal portions 52 of length $2 \cdot N/F$, all derived from the sequence of N/F long spectra 46, these shortened portions 52 corresponding to the DCT kernel, i.e. the $2 \cdot N/F$ newest samples of the erstwhile described portions.

The windower 18 acts as described previously and generates a windowed temporal portion 60 for each temporal portion 52, but it operates merely on the DCT kernel. To this end, windower 18 uses window function ω_i with $i=0 \dots 2N/F-1$, having the kernel size. The relationship between w_i with $i=0 \dots (E+2) \cdot N/F-1$ is described later, just as the relationship between the subsequently mentioned lifting coefficients and w_i with $i=0 \dots (E+2) \cdot N/F-1$ is.

Using the nomenclature applied above, the process described so far yields:

$$Z_{k,n} = \omega_n \cdot X_{k,n} \quad \text{for } n = 0, \dots, 2M-1$$

with redefining $M = N/F$, so that M corresponds to the frame size expressed in the downscaled domain and using the nomenclature of Fig. 2-6, wherein, however, $z_{k,n}$ and $x_{k,n}$ shall contain merely the samples of the windowed temporal portion and the not-yet windowed temporal portion within the DCT kernel having size $2 \cdot M$ and temporally corresponding to samples $E \cdot N/F \dots (E+2) \cdot N/F-1$ in Fig. 4. That is, n is an integer indicating a sample index and ω_n is a real-valued window function coefficient corresponding to the sample index n .

The overlap/add process of the canceller 20 operates in a manner different compared to the above description. It generates intermediate temporal portions $m_k(0), \dots, m_k(M-1)$ based on the equation or expression

$$m_{k,n} = Z_{k,n} + Z_{k-1,n+M} \quad \text{for } n = 0, \dots, M-1$$

In the implementation of Fig. 8, the apparatus further comprises a lifter 80 which may be interpreted as a part of the modulator 16 and windower 18 since the lifter 80 compensates the fact the modulator and the windower restricted their processing to the DCT kernel instead of processing the extension of the modulation functions and the synthesis window beyond the kernel towards the past which extension was introduced to compensate for the

zero portion 56. The lifter 80 produces, using a framework of the delayers and multipliers 82 and adders 84, the finally reconstructed temporal portions or frames of length M in pairs of immediately consecutive frames based on the equation or expression

$$5 \quad u_{k,n} = m_{k,n} + l_{n-M/2} \cdot m_{k-1,M-1-n} \quad \text{for } n = M/2, \dots, M-1 \quad ,$$

and

$$10 \quad u_{k,n} = m_{k,n} + l_{M-1-n} \cdot out_{k-1,M-1-n} \quad \text{for } n=0, \dots, M/2-1 \quad ,$$

wherein l_n with $n = 0 \dots M-1$ are real-valued lifting coefficients related to the downscaled synthesis window in a manner described in more detail below.

In other words, for the extended overlap of E frames into the past, only M additional multiplier-add operations are required, as can be seen in the framework of the lifter 80. These additional operations are sometimes also referred to as “zero-delay matrices”. Sometimes these operations are also known as “lifting steps”. The efficient implementation shown in Fig. 8 may under some circumstances be more efficient as a straightforward implementation. To be more precise, depending on the concrete implementation, such a more efficient implementation might result in saving M operations, as in the case of a straightforward implementation for M operations, it might be advisable to implement, as the implementation shown in Fig. 19, requires in principle, 2M operations in the framework of the module 820 and M operations in the framework of the lifter 830.

As to the dependency of ω_n with $n=0 \dots 2M-1$ and l_n with $n = 0 \dots M-1$ on the synthesis window w_i with $i = 0 \dots (E+2)M-1$ (it is recalled that here $E=2$), the following formulae describe the relationship between them with displacing, however, the subscript indices used so far into the parenthesis following the respective variable:

$$\begin{aligned}
 w(i) &= l\left(\frac{M}{2} - 1 - n\right) \cdot l(M-1-n) \cdot \omega(M+n) \\
 w(M/2 + i) &= l(n) \cdot l(M/2 + n) \cdot \omega(3M/2 + n) \\
 w(M + i) &= l\left(\frac{M}{2} - 1 - n\right) \cdot \omega(M+n) \\
 w(3M/2 + i) &= -l(n) \cdot \omega(3M/2 + n) \\
 w(2M + i) &= -\omega(M+n) - l(M-1-n) \cdot \omega(n) \\
 w(5M/2 + i) &= -\omega(3M/2 + n) - l(M/2 + n) \cdot \omega(M/2 + n) \\
 w(3M + i) &= -\omega(n)
 \end{aligned}$$

$$w(7M/2 + i) = \omega(M + n)$$

$$\text{for } i, n = 0 \dots \frac{M}{2} - 1$$

Please note that the window w_i contains the peak values on the right side in this
 5 formulation, i.e. between the indices $2M$ and $4M - 1$. The above formulae relate
 coefficients l_n with $n = 0 \dots M-1$ and ω_n $n = 0, \dots, 2M-1$ to the coefficients w_n with $n =$
 $0 \dots (E+2)M-1$ of the downscaled synthesis window. As can be seen, l_n with $n = 0 \dots M-1$
 actually merely depend on $\frac{3}{4}$ of the coefficients of the downsampled synthesis window,
 namely on w_n with $n = 0 \dots (E+1)M-1$, while ω_n $n = 0, \dots, 2M-1$ depend on all w_n with $n =$
 10 $0 \dots (E+2)M-1$.

As stated above, it might be that windower 18 obtains the downsampled synthesis window
 54 w_n with $n = 0 \dots (E+2)M-1$ from a storage where the window coefficients w_i of this
 downsampled synthesis window 54 have been stored after having been obtained using
 15 the downsampling 72, and from where same are read to compute coefficients l_n with $n =$
 $0 \dots M-1$ and ω_n $n = 0, \dots, 2M-1$ using the above relation, but alternatively, windower 18 may
 retrieve the coefficients l_n with $n = 0 \dots M-1$ and ω_n $n = 0, \dots, 2M-1$, thus computed from the
 pre-downsampled synthesis window, from the storage directly. Alternatively, as stated
 above, the audio decoder 10 may comprise the segmental downsampler 76 performing
 20 the downsampling 72 of Fig. 6 on the basis of the reference synthesis window 70, thereby
 yielding w_n with $n = 0 \dots (E+2)M-1$ on the basis of which the windower 18 computes
 coefficients l_n with $n = 0 \dots M-1$ and ω_n $n = 0, \dots, 2M-1$ using above relation/formulae. Even
 using the lifting implementation, more than one value for F may be supported.

Briefly summarizing the lifting implementation, same results in an audio decoder 10
 25 configured to decode an audio signal 22 at a first sampling rate from a data stream 24 into
 which the audio signal is transform coded at a second sampling rate, the first sampling
 rate being $1/F^{\text{th}}$ of the second sampling rate, the audio decoder 10 comprising the receiver
 12 which receives, per frame of length N of the audio signal, N spectral coefficients 28,
 30 the grabber 14 which grabs-out for each frame, a low-frequency fraction of length N/F out
 of the N spectral coefficients 28, a spectral-to-time modulator 16 configured to subject, for
 each frame 36, the low-frequency fraction to an inverse transform having modulation
 functions of length $2 \cdot N/F$ temporally extending over the respective frame and a previous
 frame so as to obtain a temporal portion of length $2 \cdot N/F$, and a windower 18 which
 35 windows, for each frame 36, the temporal portion $x_{k,n}$ according to $Z_{k,n} = \omega_n \cdot x_{k,n}$ for $n =$

0,...,2M-1 so as to obtain a windowed temporal portion $z_{k,n}$ with $n = 0 \dots 2M-1$. The time domain aliasing canceler 20 generates intermediate temporal portions $m_k(0), \dots, m_k(M-1)$ according to $m_{k,n} = z_{k,n} + z_{k-1,n+M}$ for $n = 0, \dots, M-1$. Finally, the lifter 80 computes frames $u_{k,n}$ of the audio signal with $n = 0 \dots M-1$ according to $u_{k,n} = m_{k,n} + l_{n-M/2} \cdot m_{k-1,M-1-n}$ for $n = M/2, \dots, M-1$, and $u_{k,n} = m_{k,n} + l_{M-1-n} \cdot m_{k-1,M-1-n}$ for $n=0, \dots, M/2-1$, wherein l_n with $n = 0 \dots M-1$ are lifting coefficients, wherein the inverse transform is an inverse MDCT or inverse MDST, and wherein l_n with $n = 0 \dots M-1$ and ω_n with $n = 0, \dots, 2M-1$ depend on coefficients w_n with $n = 0 \dots (E+2)M-1$ of a synthesis window, and the synthesis window is a downsampled version of a reference synthesis window of length $4 \cdot N$, downsampled by a factor of F by a segmental interpolation in segments of length $1/4 \cdot N$.

It already turned out from the above discussion of a proposal for an extension of AAC-ELD with respect to a downsampled decoding mode that the audio decoder of Fig. 2 may be accompanied with a low delay SBR tool. The following outlines, for instance, how the AAC-ELD coder extended to support the above-proposed downsampled operating mode, would operate when using the low delay SBR tool. As already mentioned in the introductory portion of the specification of the present application, in case the low delay SBR tool is used in connection with the AAC-ELD coder, the filter banks of the low delay SBR module are downsampled as well. This ensures that the SBR module operates with the same frequency resolution and therefore no more adaptations are required. Fig. 7 outlines the signal path of the AAC-ELD decoder operating at 96 kHz, with frame size of 480 samples, in down-sampled SBR mode and with a downscaling factor F of 2.

In Fig. 7, the bitstream arriving as processed by a sequence of blocks, namely an AAC decoder, an inverse LD-MDCT block, a CLDFB analysis block, an SBR decoder and a CLDFB synthesis block (CLDFB = complex low delay filter bank). The bitstream equals the data stream 24 discussed previously with respect to Figs. 3 to 6, but is additionally accompanied by parametric SBR data assisting the spectral shaping of a spectral replicate of a spectral extension band extending the spectra frequency of the audio signal obtained by the downsampled audio decoding at the output of the inverse low delay MDCT block, the spectral shaping being performed by the SBR decoder. In particular, the AAC decoder retrieves all of the necessary syntax elements by appropriate parsing and entropy decoding. The AAC decoder may partially coincide with the receiver 12 of the audio decoder 10 which, in Fig. 7, is embodied by the inverse low delay MDCT block. In Fig. 7, F is exemplarily equal to 2. That is, the inverse low delay MDCT block of Fig. 7 outputs, as an example for the reconstructed audio signal 22 of Fig. 2, a 48 kHz time signal

downsampled at half the rate at which the audio signal was originally coded into the arriving bitstream. The CLDFB analysis block subdivides this 48 kHz time signal, i.e. the audio signal obtained by downsampled audio decoding, into N bands, here $N = 16$, and the SBR decoder computes re-shaping coefficients for these bands, re-shapes the N bands accordingly – controlled via the SBR data in the input bitstream arriving at the input of the AAC decoder, and the CLDFB synthesis block re-transitions from spectral domain to time domain with obtaining, thereby, a high frequency extension signal to be added to the original decoded audio signals output by the inverse low delay MDCT block.

- 10 Please note, that the standard operation of SBR utilizes a 32 band CLDFB. The interpolation algorithm for the 32 band CLDFB window coefficients ci_{32} is already given in 4.6.19.4.1 in [1],

$$ci_{32}(i) = \frac{1}{2} [c_{64}(2i + 1) + c_{64}(2i)], \quad 0 \leq i < 320,$$

15

where c_{64} are the window coefficients of the 64 band window given in Table 4.A.90 in [1]. This formula can be further generalized to define window coefficients for a lower number of bands B as well

$$20 \quad ci_B(i) = \frac{1}{2} [c_{64}(2F \cdot i + 1 + p) + c_{64}(2F \cdot i + p)], \quad 0 \leq i < (10B), p = \text{int} \left(\frac{64}{2B} - 0.5 \right)$$

where F denotes the downscaling factor being $F = 32/B$. With this definition of the window coefficients, the CLDFB analysis and synthesis filter bank can be completely described as outlined in the above example of section A.2.

25

Thus, above examples provided some missing definitions for the AAC-ELD codec in order to adapt the codec to systems with lower sample rates. These definitions may be included in the ISO/IEC 14496-3:2009 standard.

- 30 Thus, in the above discussion it has, inter alias, been described:

An audio decoder may be configured to decode an audio signal at a first sampling rate from a data stream into which the audio signal is transform coded at a second sampling rate, the first sampling rate being $1/F^{\text{th}}$ of the second sampling rate, the audio decoder comprising: a receiver configured to receive, per frame of length N of the audio signal, N

35

- spectral coefficients; a grabber configured to grab-out for each frame, a low-frequency fraction of length N/F out of the N spectral coefficients; a spectral-to-time modulator configured to subject, for each frame, the low-frequency fraction to an inverse transform having modulation functions of length $(E + 2) \cdot N/F$ temporally extending over the
- 5 respective frame and $E+1$ previous frames so as to obtain a temporal portion of length $(E + 2) \cdot N/F$; a windower configured to window, for each frame, the temporal portion using a unimodal synthesis window of length $(E + 2) \cdot N/F$ comprising a zero-portion of length $1/4 \cdot N/F$ at a leading end thereof and having a peak within a temporal interval of the unimodal synthesis window, the temporal interval succeeding the zero-portion and having length $7/4$
- 10 $\cdot N/F$ so that the windower obtains a windowed temporal portion of length $(E + 2) \cdot N/F$; and a time domain aliasing canceler configured to subject the windowed temporal portion of the frames to an overlap-add process so that a trailing-end fraction of length $(E + 1)/(E + 2)$ of the windowed temporal portion of a current frame overlaps a leading end of length $(E + 1)/(E + 2)$ of the windowed temporal portion of a preceding frame, wherein the
- 15 inverse transform is an inverse MDCT or inverse MDST, and wherein the unimodal synthesis window is a downsampled version of a reference unimodal synthesis window of length $(E + 2) \cdot N$, downsampled by a factor of F by a segmental interpolation in segments of length $1/4 \cdot N/F$.
- 20 Audio decoder according to an embodiment, wherein the unimodal synthesis window is a concatenation of spline functions of length $1/4 \cdot N/F$.
- Audio decoder according to an embodiment, wherein the unimodal synthesis window is a concatenation of cubic spline functions of length $1/4 \cdot N/F$.
- 25 Audio decoder according to any of the previous embodiments, wherein $E = 2$.
- Audio decoder according to any of the previous embodiments, wherein the inverse transform is an inverse MDCT.
- 30 Audio decoder according to any of the previous embodiments, wherein more than 80% of a mass of the unimodal synthesis window is comprised within the temporal interval succeeding the zero-portion and having length $7/4 \cdot N/F$.

Audio decoder according to any of the previous embodiments, wherein the audio decoder is configured to perform the interpolation or to derive the unimodal synthesis window from a storage.

- 5 Audio decoder according to any of the previous embodiments, wherein the audio decoder is configured to support different values for F.

Audio decoder according to any of the previous embodiments, wherein F is between 1.5 and 10, both inclusively.

10

A method performed by an audio decoder according to any of the previous embodiments.

A computer program having a program code for performing, when running on a computer, a method according to an embodiment.

15

As far as the term "of ...length" is concerned it should be noted that this term is to be interpreted as measuring the length in samples. As far as the length of the zero portion and the segments is concerned it should be noted that same may be integer valued. Alternatively, same may be non-integer valued.

20

As to the temporal interval within which the peak is positioned it is noted that Fig. 1 shows this peak as well as the temporal interval illustratively for an example of the reference unimodal synthesis window with $E = 2$ and $N = 512$: The peak has its maximum at approximately sample No. 1408 and the temporal interval extends from sample No. 1024 to sample No. 1920. The temporal interval is, thus, $7/8$ of the DCT kernel long.

25

As to the term "downsampled version" it is noted that in the above specification, instead of this term, "downscaled version" has synonymously been used.

30

As to the term "mass of a function within a certain interval" it is noted that same shall denote the definite integral of the respective function within the respective interval.

35

In case of the audio decoder supporting different values for F, same may comprise a storage having accordingly segmentally interpolated versions of the reference unimodal synthesis window or may perform the segmental interpolation for a currently active value of F. The different segmentally interpolated versions have in common that the

interpolation does not negatively affect the discontinuities at the segment boundaries. They may, as described above, spline functions.

By deriving the unimodal synthesis window by a segmental interpolation from the reference unimodal synthesis window such as the one shown in Fig. 1 above, the $4 \cdot (E + 2)$ segments may be formed by spline approximation such as by cubic splines and despite the interpolation, the discontinuities which are to be present in the unimodal synthesis window at a pitch of $1/4 \cdot N/F$ owing to the synthetically introduced zero-portion as a means for lowering the delay are conserved.

References

[1] ISO/IEC 14496-3:2009

[2] M13958, "Proposal for an Enhanced Low Delay Coding Mode", October 2006, Hangzhou, China

Claims

1. Audio decoder configured to decode an audio signal at a first sampling rate from a data stream into which the audio signal is transform coded at a second sampling rate, the first sampling rate being $1/F^{\text{th}}$ of the second sampling rate, the audio decoder comprising:
 - 5 a receiver configured to receive, per frame of length N of the audio signal, N spectral coefficients;
 - 10 a grabber configured to grab-out for each frame, a low-frequency fraction of length N/F out of the N spectral coefficients;
 - 15 a spectral-to-time modulator configured to subject, for each frame, the low-frequency fraction to an inverse transform having modulation functions of length $(E + 2) \cdot N/F$ temporally extending over the respective frame and $E + 1$ previous frames so as to obtain a temporal portion of length $(E + 2) \cdot N/F$;
 - 20 a windower configured to window, for each frame, the temporal portion using a synthesis window of length $(E + 2) \cdot N/F$ comprising a zero-portion of length $1/4 \cdot N/F$ at a leading end thereof and having a peak within a temporal interval of the synthesis window, the temporal interval succeeding the zero-portion and having length $7/4 \cdot N/F$ so that the windower obtains a windowed temporal portion of length $(E + 2) \cdot N/F$; and
 - 25 a time domain aliasing canceler configured to subject the windowed temporal portion of the frames to an overlap-add process so that a trailing-end fraction of length $(E + 1)/(E + 2)$ of the windowed temporal portion of a current frame overlaps a leading end of length $(E + 1)/(E + 2)$ of the windowed temporal portion of a preceding frame,
 - 30 wherein the inverse transform is an inverse MDCT or inverse MDST, and
 - 35 wherein the synthesis window is a downsampled version of a reference synthesis window of length $(E + 2) \cdot N$, downsampled by a factor of F by a segmental interpolation in segments of length $1/4 \cdot N$.

2. Audio decoder according to claim 1, wherein the synthesis window is a concatenation of spline functions of length $1/4 \cdot N/F$.
- 5 3. Audio decoder according to claim 1 or 2, wherein the synthesis window is a concatenation of cubic spline functions of length $1/4 \cdot N/F$.
4. Audio decoder according to any of the previous claims, wherein $E = 2$.
- 10 5. Audio decoder according to any of the previous claims, wherein the inverse transform is an inverse MDCT.
6. Audio decoder according to any of the previous claims, wherein more than 80% of a mass of the synthesis window is comprised within the temporal interval
15 succeeding the zero-portion and having length $7/4 \cdot N/F$.
7. Audio decoder according to any of the previous claims, wherein the audio decoder is configured to perform the interpolation or to derive the synthesis window from a storage.
20
8. Audio decoder according to any of the previous claims, wherein the audio decoder is configured to support different values for F .
9. Audio decoder according to any of the previous claims, wherein F is between 1.5
25 and 10, both inclusively.
10. Audio decoder according to any of the previous claims, wherein the reference synthesis window is unimodal.
- 30 11. Audio decoder according to any of the previous claims, wherein the audio decoder is configured to perform the interpolation in such a manner that a majority of the coefficients of the synthesis window depends on more than two coefficients of the reference synthesis window.
- 35 12. Audio decoder according to any of the previous claims, wherein the audio decoder is configured to perform the interpolation in such a manner that each coefficient of

the synthesis window separated by more than two coefficient from segment borders depend on more than two coefficients of the reference synthesis window.

13. Audio decoder according to any of the previous claims, wherein the windower and the time domain aliasing canceller cooperate so that the windower skips the zero-portion in weighting the temporal portion using the synthesis window and the time domain aliasing canceller disregards a corresponding non-weighted portion of the windowed temporal portion in the overlap-add process so that merely $E+1$ windowed temporal portions are summed-up so as to result in the corresponding non-weighted portion of a corresponding frame and $E+2$ windowed portions are summed-up within a remainder of the corresponding frame.
14. Audio decoder for generating a downscaled version of a synthesis window of an audio decoder according to any of the previous claims, wherein $E=2$ so that the synthesis window function comprises a kernel related half of length $2 \cdot N/F$ preceded by a remainder half of length $2 \cdot N/F$ and wherein the spectral-to-time modulator, the windower and the time domain aliasing canceller are implemented so as to cooperate in a lifting implementation according to which

the spectral-to-time modulator confines the subjecting, for each frame, the low-frequency fraction to the inverse transform having modulation functions of length $(E + 2) \cdot N/F$ temporally extending over the respective frame and $E + 1$ previous frames, to a transform kernel coinciding with the respective frame and one previous frame so as to obtain the temporal portion $x_{k,n}$ with $n = 0 \dots 2M-1$ with $M=N/F$ being a sample index and k being a frame index;

the windower windowing, for each frame, the temporal portion $x_{k,n}$ according to $z_{k,n} = \omega_n \cdot x_{k,n}$ for $n = 0, \dots, 2M-1$ so as to obtain the windowed temporal portion $z_{k,n}$ with $n = 0 \dots 2M-1$;

the time domain aliasing canceller generates intermediate temporal portions $m_k(0), \dots, m_k(M-1)$ according to $m_{k,n} = z_{k,n} + z_{k-1,n+M}$ for $n = 0, \dots, M-1$, and

the audio decoder comprises a lifter configured to obtain the frames $u_{k,n}$ with $n = 0 \dots M-1$ according to

$$u_{k,n} = m_{k,n} + l_{n-M/2} \cdot m_{k-1,M-1-n} \quad \text{for } n = M/2, \dots, M-1,$$

and

$$u_{k,n} = m_{k,n} + l_{M-1-n} \cdot m_{k-1,M-1-n} \quad \text{for } n=0, \dots, M/2-1,$$

wherein l_n with $n = 0 \dots M-1$ are lifting coefficients, and wherein l_n with $n = 0 \dots M-1$ and ω_n with $n = 0, \dots, 2M-1$ depend on coefficients w_n with $n = 0 \dots (E+2)M-1$ of the synthesis window.

10

15. Audio decoder configured to decode an audio signal at a first sampling rate from a data stream into which the audio signal is transform coded at a second sampling rate, the first sampling rate being $1/F^{\text{th}}$ of the second sampling rate, the audio decoder comprising:

15

a receiver configured to receive, per frame of length N of the audio signal, N spectral coefficients;

20

a grabber configured to grab-out for each frame, a low-frequency fraction of length N/F out of the N spectral coefficients;

25

a spectral-to-time modulator configured to subject, for each frame, the low-frequency fraction to an inverse transform having modulation functions of length $2 \cdot N/F$ temporally extending over the respective frame and a previous frame so as to obtain a temporal portion of length $2 \cdot N/F$;

30

a windower configured to window, for each frame, the temporal portion $x_{k,n}$ according to $z_{k,n} = \omega_n \cdot x_{k,n}$ for $n = 0, \dots, 2M-1$ so as to obtain a windowed temporal portion $z_{k,n}$ with $n = 0 \dots 2M-1$;

35

a time domain aliasing canceler configured to generate intermediate temporal portions $m_k(0), \dots, m_k(M-1)$ according to $m_{k,n} = z_{k,n} + z_{k-1,n+M}$ for $n = 0, \dots, M-1$, and

a lifter configured to obtain frames $u_{k,n}$ of the audio signal with $n = 0 \dots M-1$ according to

$$u_{k,n} = m_{k,n} + l_{n-M/2} \cdot m_{k-1,M-1-n} \quad \text{for } n = M/2, \dots, M-1,$$

and

$$u_{k,n} = m_{k,n} + l_{M-1-n} \cdot m_{k-1,M-1-n} \quad \text{for } n=0, \dots, M/2-1,$$

wherein l_n with $n = 0 \dots M-1$ are lifting coefficients,

wherein the inverse transform is an inverse MDCT or inverse MDST, and

10

wherein l_n with $n = 0 \dots M-1$ and ω_n with $n = 0, \dots, 2M-1$ depend on coefficients w_n with $n = 0 \dots (E+2)M-1$ of a synthesis window, with E being an integer, and the synthesis window is a downsampled version of a reference synthesis window of length $4 \cdot N$, downsampled by a factor of F by a segmental interpolation in segments of length $1/4 \cdot N$.

15

16. Apparatus for generating a downsampled version of a synthesis window of an audio decoder according to any of the previous claims, wherein the synthesis window of length $(E + 2) \cdot N$ and the apparatus is configured to downsample the synthesis window by a factor of F by a segmental interpolation in $4 \cdot (E + 2)$ segments of equal length.

20

17. Method for generating a downsampled version of a synthesis window of an audio decoder according to any of claims 1 to 15, wherein the synthesis window of length $(E + 2) \cdot N$ and the method comprises downsampling the synthesis window by a factor of F by a segmental interpolation in $4 \cdot (E + 2)$ segments of equal length.

25

18. Method for decoding an audio signal at a first sampling rate from a data stream into which the audio signal is transform coded at a second sampling rate, the first sampling rate being $1/F^{\text{th}}$ of the second sampling rate, the method comprising:

30

receiving, per frame of length N of the audio signal, N spectral coefficients;

grabbing-out for each frame, a low-frequency fraction of length N/F out of the N spectral coefficients;

35

performing a spectral-to-time modulation by subjecting, for each frame, the low-frequency fraction to an inverse transform having modulation functions of length $(E + 2) \cdot N/F$ temporally extending over the respective frame and $E + 1$ previous frames so as to obtain a temporal portion of length $(E + 2) \cdot N/F$;

5

windowing, for each frame, the temporal portion using a synthesis window of length $(E + 2) \cdot N/F$ comprising a zero-portion of length $1/4 \cdot N/F$ at a leading end thereof and having a peak within a temporal interval of the synthesis window, the temporal interval succeeding the zero-portion and having length $7/4 \cdot N/F$ so that the windower obtains a windowed temporal portion of length $(E + 2) \cdot N/F$; and

10

performing a time domain aliasing cancellation by subjecting the windowed temporal portion of the frames to an overlap-add process so that a trailing-end fraction of length $(E + 1)/(E + 2)$ of the windowed temporal portion of a current frame overlaps a leading end of length $(E + 1)/(E + 2)$ of the windowed temporal portion of a preceding frame,

15

wherein the inverse transform is an inverse MDCT or inverse MDST, and

20

wherein the synthesis window is a downsampled version of a reference synthesis window of length $(E + 2) \cdot N$, downsampled by a factor of F by a segmental interpolation in segments of length $1/4 \cdot N$.

19. Computer program having a program code for performing, when running on a computer, a method according to claim 17 or 18.

25