




(86) Date de dépôt PCT/PCT Filing Date: 1998/10/28
(87) Date publication PCT/PCT Publication Date: 1999/05/06
(45) Date de délivrance/Issue Date: 2005/12/27
(85) Entrée phase nationale/National Entry: 2000/04/27
(86) N° demande PCT/PCT Application No.: US 1998/022870
(87) N° publication PCT/PCT Publication No.: 1999/022365
(30) Priorité/Priority: 1997/10/28 (08/958,567) US

(51) Cl.Int.⁷/Int.Cl.⁷ G10L 21/04
(72) Inventeurs/Inventors:
WU, SHUWU, US;
MANTEGNA, JOHN, US
(73) Propriétaire/Owner:
AMERICA ONLINE, INC., US
(74) Agent: MBM & CO.

(54) Titre : CODAGE AUDIO SOUS-BANDE PERCEPTEUR AU MOYEN D'UNE QUANTIFICATION VECTORIELLE CLAIRSEMEE ADAPTATIVE DE TYPE MULTIPLE, ET DISPOSITIF DE MISE A L'ECHELLE DE SIGNAUX PAR SATURATION

(54) Title: PERCEPTUAL SUBBAND AUDIO CODING USING ADAPTIVE MULTITYPE SPARSE VECTOR QUANTIZATION, AND SIGNAL SATURATION SCALER



(57) Abrégé/Abstract:

A reduced-complexity real-time audio coder/decoder ("codec") using adaptive sparse vector quantization (ASVQ) classifies signal vectors into one of a plurality of different types. The codec provides low bit-rate compression for music and speech, while being



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

applicable to higher bit-rate audio compression. It includes an in-path implementation of psychoacoustic spectral masking, and frequency-domain quantization using ASVQ and audio-compression algorithms. Critically-sampled subband filter banks maintain time-domain continuity across frame boundaries. In-path spectral masking usually results in sparse vectors. ASVQ is particularly effective for quantizing sparse signal vectors. In the preferred embodiment, ASVQ classifies signal vectors into six different types (504, 508, 512, 516, 520, and 522). The ASVQ technique applies to general-purpose data quantization as well as to quantization for audio compression. The invention also includes a "soft-clipping" algorithm in the decoder as a post-processing stage, which preserves the waveform shapes of the reconstructed time-domain audio signal in a frame- or block-oriented stateless manner while maintaining continuity across frame or block boundaries.

**PCT**WORLD INTELLECTUAL PROPERTY ORGANIZATION
International Bureau


INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification ⁶ : G10L 7/04, 3/02	A1	(11) International Publication Number: WO 99/22365 (43) International Publication Date: 6 May 1999 (06.05.99)
(21) International Application Number: PCT/US98/22870 (22) International Filing Date: 28 October 1998 (28.10.98) (30) Priority Data: 08/958,567 28 October 1997 (28.10.97) US (71) Applicant: AMERICA ONLINE, INC. [US/US]; 22000 AOL Way, Dulles, VA 20166 (US). (72) Inventors: WU, Shuwu; 18 Camarin Street, Foothill Ranch, CA 92610 (US). MANTEGNA, John; P.O. Box 6074, Irvine, CA 92616 (US). (74) Agent: LAND, John; Fish & Richardson P.C., Suite 1400, 4225 Executive Square, La Jolla, CA 92037 (US).	(81) Designated States: AU, BR, CA, JP, MX, European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE). Published <i>With international search report. Before the expiration of the time limit for amending the claims and to be republished in the event of the receipt of amendments.</i>	

(54) Title: **PERCEPTUAL SUBBAND AUDIO CODING USING ADAPTIVE MULTITYPE SPARSE VECTOR QUANTIZATION, AND SIGNAL SATURATION SCALER**

(57) Abstract

A reduced-complexity real-time audio coder/decoder ("codec") using adaptive sparse vector quantization (ASVQ) classifies signal vectors into one of a plurality of different types. The codec provides low bit-rate compression for music and speech, while being applicable to higher bit-rate audio compression. It includes an in-path implementation of psychoacoustic spectral masking, and frequency-domain quantization using ASVQ and audio-compression algorithms. Critically-sampled subband filter banks maintain time-domain continuity across frame boundaries. In-path spectral masking usually results in sparse vectors. ASVQ is particularly effective for quantizing sparse signal vectors. In the preferred embodiment, ASVQ classifies signal vectors into six different types (504, 508, 512, 516, 520, and 522). The ASVQ technique applies to general-purpose data quantization as well as to quantization for audio compression. The invention also includes a "soft-clipping" algorithm in the decoder as a post-processing stage, which preserves the waveform shapes of the reconstructed time-domain audio signal in a frame- or block-oriented stateless manner while maintaining continuity across frame or block boundaries.



- 1 -

PERCEPTUAL SUBBAND AUDIO CODING USING ADAPTIVE MULTITYPE SPARSE VECTOR QUANTIZATION, AND SIGNAL SATURATION SCALER

TECHNICAL FIELD

5 This invention relates to compression and decompression of audio signals, and more particularly to a method and apparatus for compression and decompression of audio signals using adaptive sparse vector quantization, and a novel adaptive sparse vector quantization technique for general purpose data compression.

BACKGROUND

10 Audio compression techniques have been developed to transmit audio signals in constrained bandwidth channels and store such signals on media with limited capacity. In audio compression, no assumptions can be made about the source or characteristics of the sound. Algorithms must be general enough to deal with arbitrary types of audio signals, which in turn poses a substantial constraint on viable approaches. (In this
15 document, the term "audio" refers to a signal that can be any sound in general, such as music of any type, speech, and a mixture of music and voice). General audio compression thus differs from speech coding in one significant aspect: in speech coding where the source is known *a priori*, model based algorithms are practical.

 Many audio compression techniques rely upon a "psychoacoustic model" to
20 achieve substantial compression. Psychoacoustics describes the relationship between acoustic events and the resulting perceived sounds. Thus, in a psychoacoustic model, the response of the human auditory system is taken into account in order to remove audio signal components that are imperceptible to human ears. Spectral "masking" is one of the most frequently exploited psychoacoustic phenomena. "Masking" describes the effect by
25 which a fainter, but distinctly audible, signal becomes inaudible when a louder signal occurs simultaneously with, or within a very short time of, the lower amplitude signal. Masking depends on the spectral composition of both the masking signal and the masked signal, and on their variations with time. For example, FIG. 1 is plot of the spectrum for

- 2 -

a typical signal (trumpet) 10 and of the human perceptual threshold 12. The perceptual threshold 12 varies with frequency and power. Note that a great deal of the signal 10 is below the perceptual threshold 12 and therefore redundant. Thus, this part of the audio signal may be discarded.

5 One well-known technique that utilizes a psychoacoustic model is embodied in the MPEG-Audio standard (ISO/IEC 11172-3; 1993(E)) (here, simply "MPEG"). FIG. 2 is a block diagram of a conventional MPEG audio encoder. A digitized audio signal (e.g., a 16-bit pulse code modulated – PCM – signal) is input into one or more filter banks 20 and into a psychoacoustic "model" 22. The filter banks 20 perform a time-
10 to-frequency mapping, generating multiple subbands (e.g., 32). The filter banks 20 are "critically" sampled so that there are as many samples in the analyzed domain as there are in the time domain. The filter banks 20 provide the primary frequency separation for the encoder; a similar set of filter banks 20 serves as the reconstruction filters for the corresponding decoder. The output samples of the filter banks 20 are then quantized by
15 a bit or noise allocation function 24.

The parallel psychoacoustic model 22 calculates a "just noticeable" noise level for each band of the filter banks 20, in the form of a "signal-to-mask" ratio. This noise level is used in the bit or noise allocation function 24 to determine the actual quantizer and quantizer levels. The quantized samples from the bit or noise allocation
20 function 24 are then applied to a bitstream formatting function 26, which outputs the final encoded (compressed) bitstream. The output of the psychoacoustic model 22 may be used to adjust bit allocations in the bitstream formatting function 26, in known fashion.

Most approaches to audio compression can be broadly divided into two major categories: time and frequency domain quantization. An MPEG coder/decoder ("codec")
25 is an example of an approach employing time domain scalar quantization. In particular, MPEG employs scalar quantization of the time domain signal in individual subbands (typically 32 subbands) while bit allocation in the scalar quantizer is based on a psychoacoustic model, which is implemented separately in the frequency domain (dual-path approach).

- 3 -

MPEG audio compression is limited to applications with higher bit-rates, 1.5 bits per sample and higher. At 1.5 bits per sample, MPEG audio does not preserve the full range of frequency content. Instead, frequency components at or near the Nyquist limit are thrown away in the compression process. In a sense, MPEG audio does not truly
5 achieve compression at the rate of 1.5 bits per sample.

Quantization is one of the most common and direct techniques to achieve data compression. There are two basic quantization types: scalar and vector. Scalar quantization encodes data points individually, while vector quantization groups input data into vectors, each of which is encoded as a whole. Vector quantization typically searches
10 a codebook (a collection of vectors) for the closest match to an input vector, yielding an output index. A de-quantizer simply performs a table lookup in an identical codebook to reconstruct the original vector. Other approaches that do not involve codebooks are known, such as closed form solutions.

It is well known that scalar quantization is not optimal with respect to
15 rate/distortion tradeoffs. Scalar quantization cannot exploit correlations among adjacent data points and thus scalar quantization yields higher distortion levels than vector quantization for a given bit rate. Vector quantization schemes usually can achieve far better compression ratios at a given distortion level. Thus, time domain scalar quantization limits the degree of compression, resulting in higher bit-rates. Further,
20 human ears are sensitive to the distortion associated with zeroing even a single time domain sample. This phenomenon makes direct application of traditional vector quantization techniques on a time domain audio signal an unattractive proposition, since vector quantization at the rate of 1 bit per sample or lower often leads to zeroing of some vector components (that is, time domain samples).

25 Frequency domain quantization based audio compression is an alternative to time domain quantization based audio compression. However, there is a significant difficulty that needs to be resolved in frequency domain quantization based audio compression. The input audio signal is continuous, with no practical limits on the total time duration. It is thus necessary to encode the audio signal in a piecewise manner. Each

piece is called an audio encode or decode frame. Performing quantization in the frequency domain on a per frame basis generally leads to discontinuities at the frame boundaries. Such discontinuities result in objectionable audible artifacts (*e.g.*, “clicks” and “pops”). One remedy to this discontinuity problem is to use overlapped frames, which results in proportionally lower compression ratios and higher computational complexity. A more popular approach is to use “critically filtered” subband filter banks, which employ a history buffer that maintains continuity at frame boundaries, but at a cost of latency in the codec-reconstructed audio signal. Another complex approach is to enforce boundary conditions as constraints in audio encode and decode processes.

The inventors have determined that it would be desirable to provide an audio compression technique suitable for real-time applications while having reduced computational complexity. The technique should provide low bit-rate compression (about 1-bit per sample) for music and speech, while being applicable to higher bit-rate audio compression. The present invention provides such a technique.

The invention includes an audio coder/decoder (“codec”) that is suitable for real-time applications due to reduced computational complexity. The invention provides low bit-rate compression for music and speech, while being applicable to higher bit-rate audio compression. The invention includes an in-path implementation of psychoacoustic spectral masking, and frequency domain quantization using a novel adaptive sparse vector quantization (ASVQ) scheme and algorithms specific to audio compression.

More particularly, the inventive audio codec employs frequency domain quantization with critically sampled subband filter banks to maintain time domain continuity across frame boundaries. The invention uses an in-path spectral masking algorithm which reduces computational complexity for the codec. The input audio signal is transformed into the frequency domain in which spectral masking can be directly applied. This in-path spectral masking usually results in sparse vectors. The sparse frequency domain signal is itself quantized and encoded in the output bit-stream.

- 5 -

The ASVQ scheme used by the invention is a vector quantization algorithm that is particularly effective for quantizing sparse signal vectors. In the preferred embodiment, ASVQ adaptively classifies signal vectors into six different types of sparse vector quantization, and performs quantization accordingly. ASVQ is most effective for sparse signals; however, it provides multiple types of vector quantization that deal with different types of occasionally non-sparse or dense signal vectors. Because of this ability to deal with dense vectors as well as sparse ones, ASVQ is a general-purpose vector quantization technique.

The invention also includes a "soft clipping" algorithm in the decoder as a post-processing stage. The soft clipping algorithm preserves the waveform shapes of the reconstructed time domain audio signal in a frame- or block-oriented stateless manner while maintaining continuity across frame or block boundaries. The soft clipping algorithm provides significant advantages over the conventional "hard clipping" methods and becomes highly desirable for low bit-rate audio compression. Although the soft clipping algorithm is applied to reconstructed time domain audio signals in the preferred audio codec, its applications extend to saturated signals in general, time domain or otherwise (frequency domain or any type of transformed domain).

One aspect of the invention includes a method for compressing a digitized time-domain audio input signal, including the steps of: filtering the input signal into a plurality of subbands sufficient to provide a frequency domain representation of the input signal; spectrally masking the plurality of subbands using an in-path psychoacoustic model to generate masked subbands; classifying the masked subbands into one of a plurality of quantization vector types; computing vector quantization indices for each quantization vector type; formatting the vector quantization indices for each quantization vector type as an output bit-stream. The invention further includes related apparatus and computer programs.

An advantage of the invention is that in-path spectral masking naturally prepares the frequency domain signal for ASVQ, a novel and yet general adaptive vector quantization technique for signal vectors that often contain a significant number of zero

- 6 -

elements. In-path spectral masking and ASVQ are a natural match in the context of audio compression: the former prepares for the latter and the latter requires the former for efficient quantization.

Other advantages of the invention include:

- 5 • A new general-purpose adaptive sparse vector quantization technique for data compression. Such data may include audio, image, and other types of data.
- 10 • Adaptive quantization type selection in accordance with the invention chooses an optimal quantization technique based on time-varying properties of the input. This approach avoids some problems of the prior art, such as varying the number of subbands which intrinsically cause discontinuities to which the human auditory system is quite sensitive. ASVQ simply searches for the best possible quantization for a given input vector, and does not directly cause any
- 15 discontinuities.
- Higher data compression ratio or lower bit-rate, ideal for applications like real-time or non-real-time audio transmission over the Internet with limited connection bandwidth.
- 20 • Ultra-low bit-rate compression of certain types of audio/music. For example, one embodiment achieves audio compression at variable low bit-rates in the neighborhood of 0.5 to 1.2 bits per sample. This audio compression system is extensible to audibly transparent sound coding and reproduction at higher bit-rates.
- 25 • Low computational complexity, which leads to fast real-time applications.

The details of one or more embodiments of the invention are set forth in the accompanying drawings and the description below. Other features, objects, and

- 7 -

advantages of the invention will be apparent from the description and drawings, and from the claims.

DESCRIPTION OF DRAWINGS

FIG. 1 is plot of the spectrum for a typical signal (trumpet) and of the human
5 perceptual threshold, as is known in the prior art.

FIG. 2 is a block diagram of a conventional MPEG audio encoder, as is known in the prior art..

FIG. 3 is a block diagram of a preferred audio encoding system in accordance with the invention.

10 FIG. 4 is a block diagram of a preferred audio decoding system in accordance with the invention.

FIG. 5 is a flowchart describing a preferred embodiment of a type classifier in accordance with the invention.

FIG. 6 shows a block diagram of a programmable processing system that may
15 be used in conjunction with the invention.

Like reference numbers and designations in the various drawings indicate like elements.

DETAILED DESCRIPTION

Audio Encoding

20 FIG. 3 is a block diagram of a preferred audio encoding system in accordance with the invention. The audio encoder 300 may be implemented in software or hardware, and has five basic components: subband filter bank analysis 302; in-path spectral masking 304; adaptive sparse vector quantization 306; bit-stream formatting for output 308; and an optional rate control 310 as a feed back loop to the spectral masking
25 component 304. Each of these components is described below.

- 8 -

Subband Filter Bank Analysis

The audio encoder 300 preferably receives an input audio signal in the form of a pulse-coded modulation (PCM) 16-bit sampled time-series signal 312. Generation of PCM coded audio is well-known in the art. The input signal 312 is applied to the
5 subband filter bank analysis component 302 which generates a number of channels, N_c , from an input frame which is critically filtered to yield N_c subband samples. With N_c sufficiently high (no less than 64, and preferably 256 or 512), the output subband samples can be regarded as a frequency domain representation of the input time domain signal.

The preferred implementation of the subband filter bank analysis component
10 302 is similar to the filter banks 20 of FIG. 2 for an MPEG audio encoder, with the following parameter changes:

- The number of subbands should be no less than 64 (versus a typical 32 for MPEG), and preferably 256 or 512 for an 11.025 KHz input signal sample
15 rate.
- More aggressive windowing is used (e.g., a Kaiser-Bessel window with beta parameter exceeding 20).
- A shorter history buffer is used to reduce codec latency, typically 6 or 8 times the number of subbands (versus a typical multiplier of 16 for MPEG).
- 20 • Each encode frame consists of 1 subband sample per subband (versus typically 12 or 36 for MPEG, layer dependent).
- The well-known Fast Discrete Cosine Transform (Fast DCT) is used in performing the Modified DCT algorithm of the MPEG Audio standard.

The output of the subband filter bank analysis component 302 is a set of
25 subband samples 314 for each frame of input signals. As shown in the illustrated embodiment, much of the energy in the input signal 312 is in several lower frequencies, especially near 25 Hz, 50 Hz, and 100 Hz.

- 9 -

Spectral Masking: In-Path Implementation of Psychoacoustic Model

As noted above, spectral masking entails the idea that relatively weak spectral content in the vicinity of a strong or a group of strong spectral components may not be perceptible by human ears. Consequently, a psychoacoustic model is employed to throw
 5 away such imperceptible frequency content, an extremely useful step towards audio data compression.

The audio codec of the invention differs from conventional implementations of spectral masking by using an in-path implementation. Conventional schemes involve encoding the audio signal in one signal path while carrying out spectral masking in a
 10 separate and parallel signal path. The result is total flexibility in implementing spectral masking but at a higher cost of computational complexity. The in-path implementation of the invention actually performs spectral masking on the signal to be encoded. Thus, there is only one signal path for both encoding and spectral masking. Advantages of this approach are reduced computational complexity and natural compatibility with ASVQ
 15 (discussed below).

In-path implementation also simplifies rate control that enables ultra-low bit-rate compression with good reproductive quality of certain types of music or sound. In some cases the bit-rate can be as low as 0.5 bits per sample with acceptable quality, a feat that has not been achieved by any state-of-the-art audio compression algorithm to the best
 20 knowledge of the inventors. The preferred implementation is described as follows.

At encode initialization:

- (1) Calculate a "linear frequency-to-Bark" lookup table, F2B, based on the following equations ("Barks" are units of a frequency scale derived by mapping frequencies to critical-band numbers; critical
 25 band numbers are based on empirically derived data describing the frequency response of the human auditory system):

$$f = 0 \rightarrow f_n;$$

$$F2B = 6 * \sinh^{-1}(f / 600);$$

where f_n is the Nyquist frequency (half of the sample frequency) in Hz.

- 10 -

- (2) Calculate a "Bark-to-linear frequency" lookup table, B2F, based on the following equations:

$$B = 0 \rightarrow B_n,$$

$$B2F = 600 * \sinh(B/6);$$

where B_n is the Nyquist frequency in Barks and B2F is given in Hz.

For each audio encode frame:

- 5 (3) Determine N_{sm} as the number of strongest spectral components, where N_{sm} can be either the number of spectral components that are greater than a threshold value N_t , or a fraction of the number of subbands N_{cf} or the minimum value of N_t and N_{cf} .

- 10 (4) Repeat step 5 through 8 for the N_{sm} strongest spectral components, *i.e.*, for

$$j = 0 \rightarrow N_{sm} - 1,$$

- (5) Determine the j^{th} masker (spectral component) to be tonal (sinusoid-like) or non-tonal based on the following equations:

$$X(j) - X(j+k) \geq 7dB \Rightarrow \text{tonal},$$

$$\text{otherwise} \Rightarrow \text{non-tonal},$$

$$k = -\max_k \rightarrow +\max_k;$$

where: $X(j)$ is the spectral level in dB; $\max_k$ is the maximum k value, which depends on the sample rate and the number of subbands.

- 15 (6) Calculate a masking index av based on the following equations:

$$B(j) = F2B[f(j)],$$

$$\text{tonal} \Rightarrow av = -1.525 - 0.275 * B(j) - 4.5dB,$$

$$\text{non-tonal} \Rightarrow av = -1.525 - 0.175 * B(j) - 0.5dB;$$

where $B(j)$ is the frequency of the j^{th} masker in Barks.

- 11 -

- (7) Calculate a differential frequency in a Bark-to-masking factor lookup table, dB2MF, based on the following equations:

$$dB = -3 \rightarrow 8,$$

$$vf = vf(dB, X[B(j)]);$$

where dB is the differential frequency in Barks; vf is the MPEG Audio masking function which depends on dB and $X[B(j)]$; and $X[B(j)]$ is the level of the j^{th} masker.

- (8) Calculate an individual masking threshold $LT(j,i)$:

$$LT[B(j), B(i)] = X[B(j)] + av + vf,$$

$$LT(j,i) = LT\{B2F[B(j)], B2F[B(i)]\};$$

- (9) Calculate:

$$LTg(i) = \sum_{j=0}^{Nsm-1} 10^{LT(j,i)/10}.$$

- (10) For each spectral component, set the component to zero if it is less than the global masking threshold:

$$i = 0 \rightarrow Nc - 1,$$

$$SBS(i) \leq LTg(i) \Rightarrow SBS(i) = 0.$$

A simplified approach can be obtained in the case of low bit-rate audio encoding. The simplification is based on the following approximations:

$$av \approx av(\text{tonality}),$$

$$vf \approx vf(dB).$$

In other words, av is approximated to be independent of $B(j)$ for the j^{th} masker and vf is approximated to be independent of $X[B(j)]$. Both approximations are of zero'th order in nature. For low bit rate non-transparent audio encoding, such approximations

- 12 -

yield good and reasonable re-constructed audio output while the computational complexity is greatly reduced.

The output of the spectral masking component 304 is a set of spectrally masked subband samples 316 for each frame of input signals. As shown in the illustrated embodiment, a number of frequencies have been reduced to zero amplitude, as being inaudible.

Adaptive Sparse Vector Quantization Encoding

Adaptive sparse vector quantization is a general-purpose vector quantization technique that applies to arbitrary input vectors. However, it is most efficient in achieving a high degree of compression if the input vectors are *mostly* sparse. The basic idea in sparse vector quantization (SVQ) is to encode the locations of non-zero elements in a sparse vector and subsequently collapse the sparse vector into a reduced vector of all non-zero elements. This reduced vector, whose dimensionality is called sparse dimensionality, is then quantized by a conventional vector quantization technique, such as product lattice-pyramid vector quantization or split-vector quantization. In accordance with the invention, adaptive SVQ (ASVQ) adaptively classifies an input vector into one of six types of vectors and applies SVQ encoding.

More particularly, in operation, the output from the spectral masking component 304 is treated as a vector input to the adaptive sparse vector quantization component 306. If desired, input data can be normalized to reduce dynamic range of subsequent vector quantization. This proves to be very useful in audio encoding because of the intrinsic large audio dynamic range. In the preferred embodiment, the ASVQ component 306 classifies each vector into one of six vector types and then SVQ encodes the vector. The output of the ASVQ component 306 are sets of ASVQ indices 318.

The preferred method for quantization of arbitrary input data by adaptive sparse vector quantization comprises the steps of:

- (1) grouping consecutive points of the original data into vectors;

- 13 -

- (2) adaptively classifying the vectors into one of a plurality of vector types, including at least one sparse vector type;
- (3) collapsing each sparse vector into a corresponding compact form;
- (4) computing a plurality of vector quantization indices for each compact
5 vector by conventional vector quantization techniques; and
- (5) formatting the vector quantization indices for each vector type as an output bit-stream.

The method of adaptively classifying vector types is preferably accomplished by categorizing each vector as follows:

- 10 (1) vectors with all zero elements (Type 0);
- (2) vectors with local clustering (Type I);
- (3) vectors with amplitude similarity in non-zero elements (Type II);
- (4) dense vectors (Type III);
- (5) vectors to which a pre-vector splitting scheme should be applied
15 (Type IV); and
- (6) vectors to which a post-vector splitting scheme should be applied (Type V).

The method of collapsing sparse vectors is preferably accomplished as follows:

- 20 (1) determining locations of non-zero elements for each sparse vector;
- (2) computing lengths of regions consisting of consecutive zero elements for each sparse vector;
- (3) computing an index representation for each such computed length of region;
- 25 (4) deriving a compact vector from the sparse vector by removing all zero elements.

- 14 -

The method of computing the index representation preferably employs recursive enumeration of vectors containing non-negative integer components.

ASVQ is very flexible in the sense that the input vectors can have either low or high dimensionalities. One way to deal with input vectors with high dimensionalities
5 in ASVQ is to pre-split the input down to smaller and more manageable dimensions. This is the classical "divide-and-conquer" approach. However, this fixed mechanism of partitioning may not always make sense in practical situations. ASVQ offers a better alternative in such scenarios. The ASVQ vector-splitting mechanism can internally post-split the input vector, preserving its physical properties. For example, the subband
10 samples for a voiced frame in speech usually consists of several locally clustered spectral components. The exact location for each cluster is data-dependent, which requires an adaptive solution for optimal compression. ASVQ Type V quantization (discussed below) can be employed to achieve this end. ASVQ generally results in variable bit allocations. The variations stem from the adaptive classification of quantization types and
15 potentially from underlying variable vector quantization schemes that support various ASVQ quantization types. ASVQ thus supports differing bit allocations which enable different quality settings for data compression.

Each of the quantization types are described below, followed by an output summary table which identifies preferred output codes; the vector type classification
20 mechanism is then described in greater detail. The preferred output codes are defined as follows:

- 15 -

<i>Code</i>	<i>Description</i>
QTI	Quantization Type Index: 0-5
SDI	Sparse Dimensionality Index: number of non-zero elements in sparse input vector
ELI	Element Location Index: index to non-zero element locations
SAI	Signal Amplitude Index: index to signal amplitude codebook (Type II only)
SBV	Sign Bit Vector: represents sign of non-zero elements (Type II only)
VQI	Vector Quantization Indices: indices to the vector quantization codebooks. In a product lattice-pyramid vector quantization implementation, VQI consists of a hyper-pyramid index (HPI) and a lattice-vector index (LVI). In a split-vector full-search VQ approach, VQI consists of a codebook index for each split-vector.
VPI	Vector Partition Index: index to partitioning schemes (described below in Type V)

Type 0 SVQ: This is the trivial case among SVQ types, where the input vector is quantized as a vector of all zero elements. This type uses the least bits for quantization, hence its usefulness.

<i>Type 0 Quantization Output Summary</i>		
<i>Code</i>	<i>Bit Allocation</i>	<i>Name</i>
QTI	fixed	Quantization Type Index

Type I SVQ: In a sense, this is the original or generic case of sparse vector quantization. A lossless process is used to determine the location of non-zeros elements in order to generate an Element Location Index (ELI), and a Sparse Dimensionality Index (*i.e.*, the number of non-zero elements in the sparse input vector). The original sparse vector is then collapsed into a vector of all non-zero elements with reduced dimensionality. This reduced vector can then be vector quantized employing any one of conventional vector quantization schemes to produce Vector Quantization Indices (VQI). For example, the product lattice pyramid vector quantization algorithm could be used for this purpose. Type I SVQ does not require a particular range for input vector dimensionality. However, practical implementations may require the input vector to be pre-split down to smaller and more manageable chunks before being sent to the ASVQ quantizer. The technique of using the ELI is perfectly applicable in quantization of binary tree codes and of the best bases in wavelet and cosine packet transforms.

- 16 -

The following describes the lossless process of encoding the Element Location Index. Consider a sparse vector of dimension N with D non-zero elements. The D non-zero elements divide the $(N-D)$ zero elements into $D+1$ regions. If the number of zero elements in each of the $D+1$ regions is known, the location of the D non-zero elements can be found by:

```

location[0] = n[0]
for    i = 1 → D-1
    location[i] = location[i-1] + 1 + n[i]
end

```

where $n[i]$ is the number of zero-elements in the i^{th} region, and $\text{location}[i]$ is the location of i^{th} non-zero element.

Conversely, if the locations of D non-zero elements are known, the number of zero elements in each of the $D+1$ regions can be found by:

```

n[0] = location[0]
for    i = 1 → D-1
    n[i] = location[i] - location[i-1] - 1
end
n[D] = N - D - ∑i=0D-1 n[i]

```

10 Therefore the problem is reduced to encoding $n[i]$, $i = 0 \rightarrow D$. One can see that the $n[i]$ array obeys the following constraints:

$$n[i] \geq 0, \quad i = 0 \rightarrow D$$

$$\sum_{i=0}^D n[i] = N - D.$$

Consequently, the encoding problem becomes the indexing problem for a $D+1$ -dimensional vector with non-negative integer components and L1-norm of $N-D$, where L1-norm is the sum of the absolute values of vector components. This indexing problem can be solved as follows:

- 17 -

```

index = 0
i = 0
k = N - D
l = D + 1
while k ≠ 0
  if n[i] > 0
    index = index +  $\sum_{j=0}^{n[i]-1} N(l-1, k-j)$ 
    k = k - n[i]
  end
  l = l - 1
  i = i + 1
end

```

where $N(l, k)$ is given by the following recursive relationships:

$$\begin{aligned}
 N(0, 0) &= 1 \\
 N(0, k) &= 0, \quad k > 0 \\
 N(1, k) &= 1, \quad k > 0 \\
 N(d, 0) &= 1, \quad d > 0 \\
 N(d, 1) &= d, \quad d > 0 \\
 N(d, k) &= N(d, k-1) + N(d-1, k), \quad d, k > 1
 \end{aligned}$$

5

<i>Type I Quantization Output Summary</i>		
<i>Code</i>	<i>Bit Allocation</i>	<i>Name</i>
QTI	fixed	Quantization Type Index
SDI	fixed	Sparse Dimensionality Index
ELI	variable	Element Location Index
VQIs	variable	Vector Quantization Indices

Type II SVQ: This can be considered a very special case of Type I SVQ. In Type II SVQ, all non-zero elements have, based on some thresholding or selection criteria, close

- 18 -

or similar magnitudes. In such a scenario, only the element location index, magnitude, and sign bits of non-zero elements need to be encoded. This type of SVQ achieves significant reduction in required bits when compared to the Type I SVQ.

5 10	<i>Type II Quantization Output Summary</i>		
	<i>Code</i>	<i>Bit Allocation</i>	<i>Name</i>
	QTI	fixed	Quantization Type Index
	SDI	fixed	Sparse Dimensionality Index
	ELI	variable	Element Location Index
	SAI	fixed	Signal Amplitude Index
	SBV	variable	Sign Bit Vector

Type III SVQ: This is the case of non-sparse or dense vectors. In such cases, it is too expensive in terms of required encode bits to treat the input vectors as Type I SVQ. Thus, a conventional vector quantization technique or split vector quantization scheme may be used. Examples of suitable algorithms may be found in "*Vector Quantization and Signal*
 15 *Compression*" by A. Gersho and R. Gray (1991), which includes a discussion on various vector quantization techniques including split vector quantization (product coding).

20	<i>Type III Quantization Output Summary</i>		
	<i>Code</i>	<i>Bit Allocation</i>	<i>Name</i>
	QTI	fixed	Quantization Type Index
	VQIs	variable	Vector Quantization Indices

Type IV SVQ: This is the case where the input vectors are fairly sparse when considered as a whole (globally sparse), but non-zero elements are concentrated or clustered locally inside the input vector. Such clustered cases result in higher dimensionality in the reduced vector (by collapsing; see Type I SVQ), which requires a subsequent split vector
 25 quantization technique. Notice that the dimensionality of the reduced vector may not be lowered by simply pre-splitting the input vector before submitting to the ASVQ quantizer, as in the case of Type I SVQ, due to local clustering. However if the definition of Type I SVQ is broadened to allow for subsequent split vector quantization, then Type

- 19 -

IV SVQ can be absorbed into Type I SVQ. However, there is good reason to treat Type IV SVQ as a separate type from the Type I SVQ: locally clustered input vectors, time domain or otherwise, usually imply perceptually significant transient signals, like short audio bursts or voiced frames in speech. As such, Type IV SVQ preferably is classified as a separate type that requires more encoding bits.

<i>Type IV Quantization Output Summary</i>		
<i>Code</i>	<i>Bit Allocation</i>	<i>Name</i>
QTI	fixed	Quantization Type Index
SDI	fixed	Sparse Dimensionality Index
VQIs	variable	Vector Quantization Indices

Type V SVQ: This is an extension of Type I SVQ. Type V SVQ deals with input vectors with higher vector dimensionality, in which quantization requires pre-splitting of the input vector for practical reasons. Type I SVQ covers such input vectors if the pre-splitting is performed before quantization. However, in scenarios where pre-splitting is inappropriate, the system has to quantize the input vector as a whole. Such scenarios lead to Type V SVQ. In contrast to Type I SVQ, Type V SVQ performs post-splitting of an input vector, which breaks the input vector into several separate sparse vectors. The number of non-zero elements in each sparse vector is encoded (losslessly) in a so-called vector partition index (VPI). The subsequent quantization of each sparse vector then becomes Type I SVQ without any pre-splitting. The mechanism of encoding VPI is identical to that of ELI.

<i>Type V Quantization Output Summary</i>		
<i>Code</i>	<i>Bit Allocation</i>	<i>Name</i>
VPI	variable	Vector Partition Index
QTI	fixed	Quantization Type Index
SDIs	fixed	Sparse Dimensionality Indices
ELIs	variable	Element Location Indices
VQIs	variable	Vector Quantization Indices

- 20 -

Type Classifier: The type classifier adaptively classifies input vectors into the above mentioned six types of sparse vector quantization. The classification rules are based on sparseness of the input frame, the presence of clusters, and the similarity in amplitudes of non-zero components. There are different approaches to implementing such a type classifier. FIG. 5 is a flowchart describing a preferred embodiment of a type classifier in accordance with the invention. The process includes the following steps, which need not necessarily be performed in the stated order:

- Scan the input vector (STEP 500).
- If the input vector consists of all zero elements (STEP 502), classify the input vector as Type 0 (STEP 504).
- Otherwise, test for local clustering in the input vector based on three criteria:
 - (1) the maximum amplitude of the unnormalized input vector should be greater than a threshold value, which ensures that the input vector contains strong signal components;
 - (2) the number of strong normalized non-zero elements, determined by thresholding, should exceed a threshold value, which ensures a high number of strong non-zero elements; and
 - (3) the weighted and normalized standard deviation of non-zero element positions should be smaller than a threshold value, which ensures local clustering.

If all three criteria are met (STEP 506), the input vector is classified as Type IV (STEP 508).

- Otherwise, test whether the maximum magnitude of the input vector is less than the mean of non-zero amplitudes of the input vector times a factor K

- 21 -

(e.g., 2.0) (STEP 510). If so, then the input vector is classified as Type II (STEP 512).

- Otherwise, if the number of non-zero elements in the input vector is greater than a threshold value T (STEP 514), the input vector is classified as Type III (STEP 516).
- Otherwise, based on whether pre-splitting or post-splitting makes more sense for a particular application (the criteria are dependent on the physical properties of the input) (STEP 518), determine whether to use Type V (STEP 520) or Type I (STEP 522).

10 *Bit-stream Formatting*

The ASVQ indices 318 output by the ASVQ component 306 are then formatted into a suitable bit-stream form 320 by the bit-stream formatting component 308. In the preferred embodiment, the format is the "ART" multimedia format used by America Online and further described in U.S. Patent Application Serial No. 08/866,857, filed 5/30/97, entitled "Encapsulated Document and Format System", assigned to the assignee of the present invention and hereby incorporated by reference. However, other formats may be used, in known fashion. Formatting may include such information as identification fields, field definitions, error detection and correction data, version information, etc.

20 The formatted bit stream represents a compressed audio file that may then be transmitted over a channel, such as the Internet, or stored on a medium, such as a magnetic or optical data storage disk.

- 22 -

Rate Control

The optional rate control component 310 serves as a feed back loop to the spectral masking component 304 to control the allocation of bits. Rate control is a known technique for keeping the bit-rate within a user-specified range. This is accomplished by
5 adapting spectral-masking threshold parameters and/or bit-allocations in the quantizer. In the preferred embodiment, rate control affects two components in the encoder 300. In the spectral masking component 304, varying spectral masking thresholds determines the sparsity of the spectrum to be encoded downstream by the ASVQ component 306. Higher spectral masking thresholds yield a sparser spectrum which requires fewer bits to encode.
10 In the ASVQ component 306, the bit-rate can be further controlled via adaptive bit allocation. The rate control process yields higher quality at higher bit rates. Thus, rate control is a natural mechanism to achieve quality variation.

Audio Decoding

FIG. 4 is a block diagram of a preferred audio decoding system in accordance
15 with the invention. The audio decoder 400 may be implemented in software or hardware, and has four basic components: bit-stream decoding 402; adaptive sparse vector quantization 404; subband filter bank synthesis 406; and soft clipping 408 before outputting the reconstructed waveform.

Bit-stream Decoding

20 An incoming bit-stream 410 previously generated by an audio encoder 300 in accordance with the invention is coupled to a bit-stream decoding component 402. The decoding component simply disassembles the received binary data into the original audio data, separating out the ASVQ indices 412 in known fashion.

Adaptive Sparse Vector Quantization Decoding

25 As noted above, "de-quantizing" generally involves performing a table lookup in a codebook to reconstruct the original vector. If the reconstructed vector is in

- 23 -

compacted form, then the compacted form is expanded to a sparse vector form. More particularly, the preferred method for de-quantization of compressed bitstream data by adaptive sparse vector de-quantization comprises the steps of:

- 5 (1) decoding the input bitstream into a plurality of vector quantization indices;
- (2) reconstructing compact vectors from the vector quantization indices by conventional vector de-quantization techniques;
- (3) expanding compact vectors into sparse form for each sparse vector type;
- 10 (4) assembling sparse vectors into transcoded data.

The method of expanding compact vectors is preferably accomplished by:

- (1) computing lengths of regions consisting of consecutive zero elements from the index representation;
- 15 (2) determining locations of non-zero elements from the computed lengths of regions;
- (3) creating a corresponding sparse vector consisting of all zero elements; and
- (4) reconstructing each sparse vector by inserting compact vector components in respective determined locations.

20 The method of computing the lengths of regions preferably employs recursive reconstruction of vectors containing non-negative integer components from the index representation.

As one example, decoding the ASVQ indices 412 involves computing $n[i]$, $i = 0 \rightarrow D$ (where D is as defined above for Type I ASVQ) for an input index value. The
25 preferred algorithm is:

- 24 -

```

 $n[i] = 0, \quad i = 0 \rightarrow D$ 
 $ind = 0$ 
 $i = 0$ 
 $k = N - D$ 
 $l = D + 1$ 
while  $k > 0$ 
    if  $index == ind$ 
         $n[i] = 0$ 
        break
    end
     $j = 0$ 
    forever
         $ind = ind + N(l - 1, k - j)$ 
        if  $index < ind$ 
             $n[i] = j$ 
            break
        else
             $j++$ 
        end
    end
     $k = k - n[i]$ 
     $l--$ 
     $i++$ 
end
if  $k > 0$ 
     $n[D] = k - n[i]$ 
end

```

Application of this algorithm to the ASVQ indices 412 will result in generation of reconstructed subband samples 414.

- 25 -

Subband Filter Bank Synthesis

The subband filter bank synthesis component 406 in the decoder 400 performs the inverse operation of subband filter bank analysis component 302 in the encoder 300. The reconstructed subband samples 414 are critically transformed to generate a reconstructed time domain audio sequence 416.

The preferred implementation of the subband filter bank synthesis component 406 is essentially similar to the corresponding filter banks of an MPEG audio decoder, with the following parameter changes:

- The number of subbands should be no less than 64 (versus a typical 32 for MPEG), and preferably 256 or 512 for an 11.025 KHz input signal sample rate.
 - More aggressive windowing is used, as in the encoder (*e.g.*, a Kaiser-Bessel window with beta parameter exceeding 20).
 - A shorter history buffer is used to reduce codec latency, typically 12 or 16 times the number of subbands (versus a typical multiplier of 32 for MPEG).
- More aggressive windowing (as in the encoder) is used;
- Each re-constructed audio frame consists of 1 subband sample per subband (versus typically 12 or 36 for MPEG, layer dependent).
 - The well-known Fast Discrete Cosine Transform (Fast DCT) is used in performing the Inverse Modified DCT algorithm of the MPEG Audio standard.

Soft Clipping

Signal saturation occurs when a signal exceeds the dynamic range of the sound generation system, and is a frequent by-product of low bit-rate audio compression due to lossy algorithms. An example of such a signal is shown in enlargement 420 in FIG. 4. If a simple and naive "hard clipping" mechanism is used to cut off the excess signal, as shown by the solid horizontal line in enlargement 420, audible distortion will

- 26 -

occur. In the preferred embodiment, an optional soft clipping component 408 is used to reduce such spectral distortion.

Soft clipping in accordance with the invention detects the presence of saturation in an input frame or block. If no saturation is found in the input frame or block, the signal is passed through without any modifications. If saturation is detected, the signal is divided into regions of saturation. Each region is considered to be a single saturation even though the region may consist of multiple or disconnected saturated samples. Each region is then processed to remove saturation while preserving waveform shapes or characteristics. The algorithm also takes care of continuity constraints at frame or block boundaries in a stateless manner, so no history buffers or states are needed. The results are more natural "looking" and sounding reproduced audio, even at lower quality settings with higher compression ratios. Further, for over-modulated original material, the inventive algorithm reduces associated distortion. The preferred implementation is described as follows:

- 15 (1) *Saturation detection*: Perform frame-oriented or block-oriented saturation detection as follows:

- 27 -

```

i = 0,
while i < N - 1,
  if S(i) < MIN_VALUE || S(i) > MAX_VALUE
    j = i
    while j > 0 & abs(S(j)) > min_amp
      j --
    end
    ilo = j
    j = i
    while j < N - 1 & abs(S(j)) > min_amp
      j ++
    end
    ihi = j
    saturationFound
    leftEdge = ilo
    rightEdge = ihi
    i = ihi
  end
  i ++
end

```

where N is the number of samples in a signal frame or block, MIN_VALUE and MAX_VALUE are minimum and maximum signal values for a given signal dynamic range, respectively, and min_amp is the amplitude threshold.

- 5 (2) *Scaling saturated regions:* Soft clipping for each saturation region is achieved through point-wise multiplication of the signal sequence by a set of scaling factors. All of the individual multiplication factors constitute an attenuation curve for the saturated region. A requirement for the attenuation curve is that it should yield identity at each end.
- 10 Each saturation region can be divided into contiguous left, center, and

- 28 -

right sub-regions. The center region contains all the saturated samples. The required loss factors for the center region can be simply determined by a factor that is just sufficient to bring all saturated samples within the signal dynamic range. The attenuation factors for the remaining two sub-regions can be determined through the constraint that the resulting attenuation curve should be continuous and, ideally, smooth. Further, it is preferable to maintain the relative order of the absolute sample values, *i.e.*, a larger absolute sample value in the original signal should yield a larger clipped absolute sample value.

The final output results in an uncompressed, soft-clipped signal 418 that is a version of the reconstructed time domain audio sequence 416. The peak amplitude characteristics of the soft-clipped signal 418 are similar to that shown in enlargement 422, where the approximate shape – and thus spectral characteristics – of the saturated input signal are preserved while reducing the amplitude of the signal below the saturation threshold; compare enlargement 420 with enlargement 422.

Computer Implementation

The invention may be implemented in hardware or software, or a combination of both. However, preferably, the invention is implemented in computer programs executing on programmable computers each including at least one processor, at least one data storage system (including volatile and non-volatile memory and/or storage elements), at least one input device, and at least one output device. Program code is applied to input data to perform the functions described herein and generate output information. The output information is applied to one or more output devices, in known fashion.

By way of example only, FIG. 6 shows a block diagram of a programmable processing system 60 that may be used in conjunction with the invention. The processing

- 29 -

system 60 preferably includes a CPU 60, a RAM 61, a ROM 62 (preferably writeable, such as a flash ROM) and an I/O controller 63 coupled by a CPU bus. The I/O controller 63 is coupled by means of an I/O bus to an I/O Interface 64. The I/O Interface 64 is for receiving and transmitting data in analog or digital form over a communications link, such as a serial link, local area network, wireless link, parallel link, *etc.* Also coupled to the I/O bus is a display 65 and a keyboard 66. Other connections may be used, such as separate busses, for the I/O Interface 64, display 65, and keyboard 66. The programmable processing system 60 may be pre-programmed, or may be programmed (and reprogrammed) by downloading a program from another source (*e.g.*, another computer).

Each program is preferably implemented in a high level procedural or object oriented programming language to communicate with a computer system. However, the programs can be implemented in assembly or machine language, if desired. In any case, the language may be a compiled or interpreted language.

Each such computer program is preferably stored on a storage media or device (*e.g.*, CDROM or magnetic diskette) readable by a general or special purpose programmable computer, for configuring and operating the computer when the storage media or device is read by the computer to perform the procedures described herein. The inventive system may also be considered to be implemented as a computer-readable storage medium, configured with a computer program, where the storage medium so configured causes a computer to operate in a specific and predefined manner to perform the functions described herein.

A number of embodiments of the present invention have been described. Nevertheless, it will be understood that various modifications may be made without departing from the spirit and scope of the invention. Accordingly, other embodiments are within the scope of the following claims.

THE EMBODIMENTS OF THE INVENTION FOR WHICH AN EXCLUSIVE PROPERTY OR PRIVILEGE IS CLAIMED ARE DEFINED AS FOLLOWS:

1. A method for compressing a digitized time-domain audio input signal, including the steps of:
 - (a) filtering the input signal into a plurality of subbands sufficient to provide a frequency domain representation of the input signal;
 - (b) spectrally masking the plurality of subbands using an in-path psychoacoustic model to generate masked subbands;
 - (c) classifying the masked subbands into one of a plurality of quantization vector types;
 - (d) computing vector quantization indices for each quantization vector type;
 - (e) formatting the vector quantization indices for each quantization vector type as an output bit-stream.
2. The method of claim 1, wherein at least one quantization vector type is a sparse vector quantization type.
3. The method of claim 1, further including decompressing the output bit-stream by the steps of:
 - (a) decoding the output bit stream into vector quantization indices;
 - (b) reconstructing the masked subbands from the vector quantization indices;
 - (c) synthesizing a digitized time-domain audio output signal from the reconstructed masked subbands.
4. The method of claim 3, further including the step of soft clipping the output signal to be within a specified dynamic range.
5. The method of claim 4, wherein the output signal is formatted in frames, and the step of soft clipping includes the steps of:
 - (a) detecting if any part of the output signal within a frame is saturated;

- (b) if saturation is detected, then dividing the output signal within the frame into regions of saturation;
 - (c) scaling each region of saturation while maintaining continuity across frame boundaries to produce a clipped output signal.
- 6. The method of claim 1, wherein the number of subbands is greater than or equal to 64.
- 7. A computer-readable medium carrying a computer program for compressing a digitized time-domain audio input signal, including instructions for causing a computer to:
 - (a) filter the input signal into a plurality of subbands sufficient to provide a frequency domain representation of the input signal;
 - (b) spectrally mask the plurality of subbands using an in-path psychoacoustic model to generate masked subbands;
 - (c) classify the masked subbands into one of a plurality of quantization vector types;
 - (d) compute vector quantization indices for each quantization vector type;
 - (e) format the vector quantization indices for each quantization vector type as an output bit-stream.
- 8. The computer-readable medium of claim 7, wherein at least one quantization vector type is a sparse vector quantization type.
- 9. The computer-readable medium of claim 7, wherein the computer program further includes instructions for decompressing the output bit-stream by causing the computer to:
 - (a) decode the output bit stream into vector quantization indices;
 - (b) reconstruct the masked subbands from the vector quantization indices;
 - (c) synthesize a digitized time-domain audio output signal from the reconstructed masked subbands.

10. The computer-readable medium of claim 9, wherein the computer program further includes instructions for causing the computer to soft clip the output signal to be within a specified dynamic range.
11. The computer-readable medium of claim 10, wherein the instructions for causing the computer to soft clip the output signal include instructions for causing the computer to:
 - (a) detect if any part of the output signal within a frame is saturated;
 - (b) if saturation is detected, then divide the output signal within the frame into regions of saturation;
 - (c) scale each region of saturation while maintaining continuity across frame boundaries to produce a clipped output signal.
12. The computer-readable medium of claim 7, wherein the number of subbands is greater than or equal to 64.
13. An apparatus for compressing a digitized time-domain audio input signal, including:
 - (a) means for filtering the input signal into a plurality of subbands sufficient to provide a frequency domain representation of the input signal;
 - (b) means for spectrally masking the plurality of subbands using an in-path psychoacoustic model to generate masked subbands;
 - (c) means for classifying the masked subbands into one of a plurality of quantization vector types;
 - (d) means for computing vector quantization indices for each quantization vector type;
 - (e) means for formatting the vector quantization indices for each quantization vector type as an output bit stream.
14. The apparatus of claim 13, wherein at least one quantization vector type is a sparse vector quantization type.

15. The apparatus of claim 13, further including means for decompressing the output bit stream by:
 - (a) decoding the output bit stream into vector quantization indices;
 - (b) reconstructing the masked subbands from the vector quantization indices;
 - (c) synthesizing a digitized time-domain audio output signal from the reconstructed masked subbands.
16. The apparatus of claim 15, further including means for soft clipping the output signal to be within a specified dynamic range.
17. The apparatus of claim 16, wherein the output signal is formatted in frames, and further including soft clipping means for:
 - (a) detecting if any part of the output signal within a frame is saturated;
 - (b) if saturation is detected, then dividing the output signal within the frame into regions of saturation;
 - (c) scaling each region of saturation while maintaining continuity across frame boundaries to produce a clipped output signal.
18. The apparatus of claim 13, wherein the number of subbands is greater than or equal to 64.
19. A method of decompressing a bit stream including vector quantization indices for a plurality of vector types, the vector quantization indices representing a digitized time-domain audio input signal compressed using adaptive sparse vector quantization applied to masked subbands generated from the digitized time-domain audio input signal, including the steps of:
 - (a) decoding the output bit stream into vector quantization indices;
 - (b) reconstructing masked subbands from the vector quantization indices;
 - (c) synthesizing the digitized time-domain audio output signal from the reconstructed masked subbands.

20. The method of claim 19, wherein the step of reconstructing masked subbands includes the step of reconstructing sparse vectors from at least some of the vector quantization indices.
21. A computer-readable medium carrying a computer program, for decompressing a bitstream including vector quantization indices for a plurality of vector types, the vector quantization indices representing a digitized time-domain audio input signal compressed using adaptive sparse vector quantization applied to masked subbands generated from the digitized time-domain audio input signal, including instructions for causing a computer to:
 - (a) decode the output bit stream into vector quantization indices;
 - (b) reconstruct masked subbands from the vector quantization indices;
 - (c) synthesize the digitized time-domain audio output signal from the reconstructed masked subbands.
22. The computer-readable medium of claim 21, wherein the instructions for causing a computer to reconstruct masked subbands further include instructions for causing the computer to reconstruct sparse vectors from at least some of the vector quantization indices.
23. An apparatus for decompressing a bit stream including vector quantization indices for a plurality of vector types, the vector quantization indices representing a digitized time-domain audio input signal compressed using adaptive sparse vector quantization applied to masked subbands generated from the digitized time-domain audio input signal, including:
 - (a) means for decoding the output bit stream into vector quantization indices;
 - (b) means for reconstructing masked subbands from the vector quantization indices;
 - (c) means for synthesizing the digitized time-domain audio output signal from the reconstructed masked subbands.

24. The apparatus of claim 23, wherein the means for reconstructing masked subbands includes means for reconstructing sparse vectors from at least some of the vector quantization indices.
25. A method for compressing a digitized time-domain input signal, including the steps of:
 - (a) filtering the input signal into a plurality of subbands sufficient to provide a frequency domain representation of the input signal;
 - (b) classifying the subbands into one of a plurality of quantization vector types, at least one of such quantization vector types being a sparse vector type;
 - (c) computing vector quantization indices for each quantization vector type;
 - (d) formatting the vector quantization indices for each vector type as an output bit stream.
26. A method for transforming and compressing signals representing a digitized timedomain input signal, the input signal being filtered into a plurality of subbands sufficient to provide a frequency domain representation of the input signal, including the steps of:
 - (a) classifying the subbands into one of a plurality of quantization vector types, at least one of such quantization vector types being a sparse vector type;
 - (b) computing vector quantization indices for each quantization vector type;
 - (c) outputting vector quantization indices for each vector type as a bit-stream representing a transformed and compressed version of the digitized time domain input signal.
27. The method of claims 25 or 26, wherein the step of computing vector quantization indices includes computing vector quantization indices for a quantization vector type based on the degree of sparseness of such quantization vector type.
28. The method of claims 25 or 26, wherein the input signal is an audio signal.

29. The method of claim 28, further including the step of spectrally masking the subbands using an in-path psychoacoustic model to generate masked subbands before computing the vector quantization indices.
30. A method for decompressing a bitstream including vector quantization indices for a plurality of vector types, the vector quantization indices representing a digitized time-domain input signal compressed using adaptive sparse vector quantization applied to subbands generated from the digitized time-domain input signal, including the steps of:
 - (a) decoding the output bit stream into vector quantization indices;
 - (b) reconstructing subbands from the vector quantization indices;
 - (c) synthesizing the digitized time-domain output signal from the reconstructed subbands.
31. A computer-readable medium carrying a computer program, for compressing a digitized time-domain input signal, including instructions for causing a computer to:
 - (a) filter the input signal into a plurality of subbands sufficient to provide a frequency domain representation of the input signal;
 - (b) classify the subbands into one of a plurality of quantization vector types, at least one of such quantization vector types being a sparse vector type;
 - (c) compute vector quantization indices for each quantization vector type;
 - (d) format the vector quantization indices for each vector type as an output bit stream.
32. A computer-readable medium carrying a computer program, for transforming and compressing signals representing a digitized time-domain input signal, the input signal being filtered into a plurality of subbands sufficient to provide a frequency domain representation of the input signal, including instructions for causing a computer to:
 - (a) classify the subbands into one of a plurality of quantization vector types, at

least one of such quantization vector types being a sparse vector type;
(b) compute vector quantization indices for each quantization vector type;
(c) output vector quantization indices for each vector type as a bit-stream representing a transformed and compressed version of the digitized time domain input signal.

33. The computer-readable medium of claims 31 or 32, wherein the instructions for causing a computer to compute vector quantization indices includes instructions for causing the computer to compute vector quantization indices for a quantization vector type based on the degree of sparseness of such quantization vector type.
34. The computer-readable medium of claims 31 or 32, wherein the input signal is an audio signal.
35. The computer-readable medium of claim 34, further including instructions for causing the computer to spectrally mask the subbands using an in-path psychoacoustic model to generate masked subbands before computing the vector quantization indices.
36. A computer-readable medium carrying a computer program, for decompressing a bitstream including vector quantization indices for a plurality of vector types, the vector quantization indices representing a digitized time-domain input signal compressed using adaptive sparse vector quantization applied to subbands generated from the digitized time-domain input signal, including instructions for causing a computer to:
 - (a) decode the output bit stream into vector quantization indices;
 - (b) reconstruct subbands from the vector quantization indices;
 - (c) synthesize the digitized time-domain output signal from the reconstructed subbands.

37. An apparatus for compressing a digitized time-domain input signal, including:
- (a) means for filtering the input signal into a plurality of subbands sufficient to provide a frequency domain representation of the input signal;
 - (b) means for classifying the subbands into one of a plurality of quantization vector types, at least one of such quantization vector types being a sparse vector type;
 - (c) means for computing vector quantization indices for each quantization vector type;
 - (d) means for formatting the vector quantization indices for each vector type as an output bit-stream.
38. An apparatus for transforming and compressing signals representing a digitized time-domain input signal, the input signal being filtered into a plurality of subbands sufficient to provide a frequency domain representation of the input signal, including:
- (a) means for classifying the subbands into one of a plurality of quantization vector types, at least one of such quantization vector types being a sparse vector type;
 - (b) means for computing vector quantization indices for each quantization vector type;
 - (c) means for outputting vector quantization indices for each vector type as a bit stream representing a transformed and compressed version of the digitized time-domain input signal.
39. The apparatus of claims 37 or 38, wherein the means for computing vector quantization indices includes means for computing vector quantization indices for a quantization vector type based on the degree of sparseness of such quantization vector type.
40. The apparatus of claims 37 or 38, wherein the input signal is an audio signal.
41. The apparatus of claim 40, further including means for spectrally masking the subbands using an in-path psychoacoustic model to generate masked subbands before computing the vector quantization indices.

42. An apparatus for decompressing a bitstream including vector quantization indices for a plurality of vector types, the vector quantization indices representing a digitized time-domain input signal compressed using adaptive sparse vector quantization applied to subbands generated from the digitized time-domain input signal, including:
- (a) means for decoding the output bit stream into vector quantization indices;
 - (b) means for reconstructing subbands from the vector quantization indices;
 - (c) means for synthesizing the digitized time-domain output signal from the reconstructed subbands.
43. A method for quantization of arbitrary data, input into a computer, by adaptive sparse vector quantization including the steps of:
- (a) grouping consecutive points of the original data into vectors;
 - (b) adaptively classifying the vectors into one of a plurality of vector types, including at least one sparse vector type;
 - (c) collapsing each sparse vector into a corresponding compact form;
 - (d) computing a plurality of vector quantization indices for each compact vector;
- and
- (e) formatting the vector quantization indices for each vector type as an output bit-stream
44. The method of claim 43, wherein the step of adaptively classifying vectors includes the steps of:
- (a) analyzing each vector;
 - (b) classifying each analyzed vector with all zero elements as a first vector type;
 - (c) classifying each analyzed vector with local clustering as a second vector type;
 - (d) classifying each analyzed vector with amplitude similarity in non-zero elements as a third vector type;
 - (e) classifying each analyzed vector with dense vectors as a fourth vector type;
 - (f) classifying each analyzed vector with vectors to which a pre-vector splitting scheme should be applied as a fifth vector type;

- (g) classifying each analyzed vector with vectors to which a post-vector splitting scheme should be applied as a sixth vector type.
45. The method of claim 43, wherein the step of collapsing sparse vectors includes the steps of:
- (a) determining locations of non-zero elements in each sparse vector;
 - (b) computing lengths of regions consisting of consecutive zero elements in each sparse vector;
 - (c) computing an index representation for the computed lengths of regions for each sparse vector;
 - (d) deriving a compact vector from each sparse vector by removing all zero elements.
46. The method of claim 45, wherein the step of computing an index representation includes the step of applying recursive enumeration to each vector containing non-negative integer components.
47. A computer-readable medium carrying a computer program, for quantization of arbitrary data, input into a computer, by adaptive sparse vector quantization, including instructions for causing the computer to:
- (a) group consecutive points of the original data into vectors;
 - (b) adaptively classify the vectors into one of a plurality of vector types, including at least one sparse vector type;
 - (c) collapse each sparse vector into a corresponding compact form;
 - (d) compute a plurality of vector quantization indices for each compact vector; and
 - (e) format the vector quantization indices for each vector type as an output bit-stream.
48. The computer-readable medium of claim 47, wherein the instructions for causing a computer to adaptively classify vectors includes instructions for causing a computer to:

- (a) analyze each vector;
 - (b) classify each analyzed vector with all zero elements as a first vector type;
 - (c) classify each analyzed vector with local clustering as a second vector type;
 - (d) classify each analyzed vector with amplitude similarity in non-zero elements as a third vector type;
 - (e) classify each analyzed vector with dense vectors as a fourth vector type;
 - (f) classify each analyzed vector with vectors to which a pre-vector splitting scheme should be applied as a fifth vector type;
 - (g) classify each analyzed vector with vectors to which a post-vector splitting scheme should be applied as a sixth vector type.
49. The computer-readable medium of claim 47, wherein the instructions for causing a computer to collapse sparse vectors includes the steps of:
- (a) determine locations of non-zero elements in each sparse vector;
 - (b) compute lengths of regions consisting of consecutive zero elements in each sparse vector;
 - (c) compute an index representation for the computed lengths of regions for each sparse vector;
 - (d) derive a compact vector from each sparse vector by removing all zero elements.
50. The computer-readable medium of claim 49, wherein the instructions for causing a computer to compute an index representation includes instructions for causing a computer to apply recursive enumeration to each vector containing non-negative integer components.
51. An apparatus for quantization of arbitrary data, input into a computer, by adaptive sparse vector quantization including:
- (a) means for grouping consecutive points of the original data into vectors;
 - (b) means for adaptively classifying the vectors into one of a plurality of vector types, including at least one sparse vector type;
 - (c) means for collapsing each sparse vector into a corresponding compact form;

- (d) means for computing a plurality of vector quantization indices for each compact vector; and
 - (e) means for formatting the vector quantization indices for each vector type as an output bit-stream.
52. The apparatus of claim 51, wherein the means for adaptively classifying vectors includes:
- (a) means for analyzing each vector;
 - (b) means for classifying each analyzed vector with all zero elements as a first vector type;
 - (c) means for classifying each analyzed vector with local clustering as a second vector type;
 - (d) means for classifying each analyzed vector with amplitude similarity in non-zero elements as a third vector type;
 - (e) means for classifying each analyzed vector with dense vectors as a fourth vector type;
 - (f) means for classifying each analyzed vector with vectors to which a pre-vector splitting scheme should be applied as a fifth vector type;
 - (g) means for classifying each analyzed vector with vectors to which a post-vector splitting scheme should be applied as a sixth vector type.
53. The apparatus of claim 51, wherein the means for collapsing sparse vectors includes:
- (a) means for determining locations of non-zero elements in each sparse vector;
 - (b) means for computing lengths of regions consisting of consecutive zero elements in each sparse vector;
 - (c) means for computing an index representation for the computed lengths of regions for each sparse vector;
 - (d) means for deriving a compact vector from each sparse vector by removing all zero elements.

54. The apparatus of claim 53, wherein the means for computing an index representation includes means for applying recursive enumeration to each vector containing non-negative-integer components.
55. A method for de-quantization of compressed input bitstream data, input into a computer, by adaptive sparse vector de-quantization, including the steps of:
- (a) decoding the input bitstream data into a plurality of vector quantization indices, at least one type of such vector quantization indices defining a sparse vector type;
 - (b) reconstructing compact vectors from the vector quantization indices;
 - (c) expanding each compact vector into sparse vector form for each sparse vector type;
 - (d) assembling sparse vectors into transcoded data.
56. The method of claim 55, wherein the step of expanding compact vectors includes the steps of:
- (a) computing lengths of regions consisting of consecutive zero elements from the vector quantization indices;
 - (b) determining locations of non-zero elements from the computed lengths of regions;
 - (c) creating a corresponding sparse vector consisting of all zero elements; and
 - (d) reconstructing each sparse vector by inserting compact vector components in the determined locations.
57. The method of claim 56, wherein the step of computing lengths of regions includes the step of applying recursive reconstruction of vectors containing non-negative integer components from the vector quantization indices.
58. A computer-readable medium carrying a computer program, for de-quantization of compressed input bitstream data, input into a computer, by adaptive sparse vector de-quantization, including instructions for causing the computer to:

- (a) decode the input bitstream data into a plurality of vector quantization indices, at least one type of such vector quantization indices defining a sparse vector type;
 - (b) reconstruct compact vectors from the vector quantization indices;
 - (c) expand each compact vector into sparse vector form for each sparse vector type;
 - (d) assemble sparse vectors into transcoded data.
59. The computer-readable medium of claim 58, wherein the instructions for causing a computer to expand compact vectors includes instructions for causing the computer to:
- (a) compute lengths of regions consisting of consecutive zero elements from the vector quantization indices;
 - (b) determine locations of non-zero elements from the computed lengths of regions;
 - (c) create a corresponding sparse vector consisting of all zero elements; and
 - (d) reconstruct each sparse vector by inserting compact vector components in the determined locations.
60. The computer-readable medium of claim 59, wherein the instructions for causing a computer to compute lengths of regions includes instructions for causing the computer to apply recursive reconstruction of vectors containing non-negative integer components from the vector quantization indices.
61. An apparatus for de-quantization of compressed input bitstream data, input into a computer, by adaptive sparse vector de-quantization, including:
- (a) means for decoding the input bitstream data into a plurality of vector quantization indices, at least one type of such vector quantization indices defining a sparse vector type;
 - (b) means for reconstructing compact vectors from the vector quantization indices;
 - (c) means for expanding each compact vector into sparse vector form for each sparse vector type;
 - (d) means for assembling sparse vectors into transcoded data.

62. The apparatus of claim 61, wherein the means for expanding compact vectors includes:
- (a) means for computing lengths of regions consisting of consecutive zero elements from the vector quantization indices;
 - (b) means for determining locations of non-zero elements from the computed lengths of regions;
 - (c) means for creating a corresponding sparse vector consisting of all zero elements; and
 - (d) means for reconstructing each sparse vector by inserting compact vector components in the determined locations.
63. The apparatus of claim 62, wherein the means for computing lengths of regions includes means for applying recursive reconstruction of vectors containing non-negative integer components from the vector quantization indices.







Fig. 1 (Prior Art)

Fig. 2
(Prior Art)

Audio Encoder





4/5




FIG. 5




FIG. 6

