




INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification ⁵ : G10L 3/02, 7/04, 9/18	A1	(11) International Publication Number: WO 94/07237 (43) International Publication Date: 31 March 1994 (31.03.94)
(21) International Application Number: PCT/US93/08572 (22) International Filing Date: 10 September 1993 (10.09.93) (30) Priority data: 07/948,147 21 September 1992 (21.09.92) US (71) Applicant: AWARE, INC. [US/US]; One Memorial Drive, 4th Floor, Cambridge, MA 02142-1301 (US). (72) Inventors: STAUTNER, John, P. ; 12 Pine Ridge Road, Wellesley Hills, MA 02181 (US). MORRELL, William, R. ; 4 Willow Avenue, Somerville, MA 02144 (US). JAY-ASIMHA, Sriram ; 1 Longfellow Place, Apt. 1020, Boston, MA 02114 (US).		(74) Agent: WARD, Calvin, B.; McCubbrey, Bartels, Meyer & Ward, Suite 2700, One Post Street, San Francisco, CA 94104-5231 (US). (81) Designated States: AU, CA, JP, KR, European patent (AT, BE, CH, DE, DK, ES, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE). Published <i>With international search report.</i>

(54) Title: AUDIO COMPRESSION SYSTEM EMPLOYING MULTI-RATE SIGNAL ANALYSIS

**(57) Abstract**

Audio input is broken into eight frequency bands by an initial filter bank (31). The first three frequency bands are sent to three filter banks (32, 33 and 34) which further divide the first three frequency bands into eight, five and four frequency bands, respectively. This results in a total of 22 frequency bands.

FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AT	Austria	FR	France	MR	Mauritania
AU	Australia	GA	Gabon	MW	Malawi
BB	Barbados	GB	United Kingdom	NE	Niger
BE	Belgium	GN	Guinea	NL	Netherlands
BF	Burkina Faso	GR	Greece	NO	Norway
BG	Bulgaria	HU	Hungary	NZ	New Zealand
BJ	Benin	IE	Ireland	PL	Poland
BR	Brazil	IT	Italy	PT	Portugal
BY	Belarus	JP	Japan	RO	Romania
CA	Canada	KP	Democratic People's Republic of Korea	RU	Russian Federation
CF	Central African Republic	KR	Republic of Korea	SD	Sudan
CG	Congo	KZ	Kazakhstan	SE	Sweden
CH	Switzerland	LI	Liechtenstein	SI	Slovenia
CI	Côte d'Ivoire	LK	Sri Lanka	SK	Slovak Republic
CM	Cameroon	LU	Luxembourg	SN	Senegal
CN	China	LV	Latvia	TD	Chad
CS	Czechoslovakia	MC	Monaco	TG	Togo
CZ	Czech Republic	MG	Madagascar	UA	Ukraine
DE	Germany	ML	Mali	US	United States of America
DK	Denmark	MN	Mongolia	UZ	Uzbekistan
ES	Spain			VN	Viet Nam
FI	Finland				

**AUDIO COMPRESSION SYSTEM EMPLOYING MULTI-RATE SIGNAL
ANALYSIS**

Field of the Invention

5

The present invention relates to audio compression and decompression systems.

Background of the Invention

10

While digital audio recordings provide many advantages over analog systems, the data storage requirements for high-fidelity recordings are substantial. A high fidelity recording typically requires more than one million bits per second of playback time. The total storage needed for even a short recording is too high
15 for many computer applications. In addition, the digital bit rates inherent in non-compressed high fidelity audio recordings makes the transmission of such audio tracks over limited bandwidth transmission systems difficult. Hence, systems for compressing audio sound tracks to reduce the storage and bandwidth requirements are in great demand.

20

One class of prior art audio compression systems divide the sound track into a series of segments. Over the time interval represented by each segment, the sound track is analyzed to determine the signal components in each of a plurality of frequency bands. The measured components are then replaced by
25 approximations requiring fewer bits to represent, but which preserve features of the sound track that are important to a human listener. At the receiver, an approximation to the original sound track is generated by reversing the analysis process with the approximations in place of the original signal components.

30

The analysis and synthesis operations are normally carried out with the aid of perfect, or near perfect, reconstruction filter banks. The systems in question include an analysis filter bank which generates a set of decimated subband outputs from a segment of the sound track. Each decimated subband output represents the signal in a predetermined frequency range. The inverse operation is carried out by
35 a synthesis filter bank which accepts a set of decimated subband outputs and generates therefrom a segment of audio sound track. In practice, the synthesis and

-2-

analysis filter banks are implemented on digital computers which may be general purpose computers or special computers designed to more efficiently carry out the operations. If the analysis and synthesis operations are carried out with sufficient precision, the segment of audio sound track generated by the synthesis filter bank will match the original segment of audio sound track that was inputted to the analysis filter bank. The differences between the reconstructed audio sound track and the original sound track can be made arbitrarily small. In this case, the specific filter bank characteristics such as the length of the segment analyzed, the number of filters in the filter bank, and the location and shape of filter response characteristics would be of little interest, since any set of filter banks satisfying the perfect, or near-perfect, reconstruction condition would exactly regenerate the audio segment.

Unfortunately, the replacement of the frequency components generated by the analysis filter bank with a quantized approximation thereto results in artifacts that do depend on the detail characteristics of the filter banks. There is no single segment length for which the artifacts in the reconstructed audio track can be minimized. Hence, the length of the segments analyzed in prior art systems is chosen to be a compromise. When the frequency components are replaced by approximations, an error is introduced in each component. An error in a given frequency component produces an acoustical effect which is equivalent to the introduction of a noise signal with frequency characteristics that depend on filter characteristics of the corresponding filter in the filter bank. The noise signal will be present over the entire segment of the reconstructed sound track. Hence, the length of the segments is reflected in the types of artifacts introduced by the approximations. If the segment is short, the artifacts are less noticeable. Hence, short segments are preferred. However, if the segment is too short, there is insufficient spectral resolution to acquire information needed to properly determine the minimum number of bits needed to represent each frequency component. On the other hand, if the segment is too long, temporal resolution of the human auditory system will detect artifacts.

Prior art systems also utilize filter banks in which the frequency bands are uniform in size. Systems with a few (16-32) sub-bands in a 0-22 kHz frequency range are generally called "subband coders" while those with a large number of sub-bands (≥ 64) are called "transform coders". It is known from psychophysical

studies of the human auditory system that there are critical bandwidths which vary with frequency. The information in a critical band may be approximated by a component representing the time averaged signal amplitude in the critical band.

5 In addition, the ear's sensitivity to a noise source in the presence of a localized frequency component such as a sine tone depends on the relative levels of the signals and on the relation of the noise spectral components to the tone. The errors introduced by approximating the frequency components may be viewed as "noise". The noise becomes significantly less audible if its spectral energy is
10 within one critical bandwidth of the tone. Hence, it is advantageous to use frequency decompositions which approximate the critical band structure of the auditory system.

 Systems which utilize uniform frequency bands are poorly suited for
15 systems designed to take advantage of this type of approximation. In principle, each audio segment can be analyzed to generate a large number of uniform frequency bands, and then, several bands at the higher frequencies could be merged to provide a decomposition into critical bands. This approach imposes the same temporal constraints on all frequency bands. That is, the time window over
20 which the low frequency data is generated for each band is the same as the time window over which each high-frequency band is generated. To provide accuracy in the low frequency ranges, the time window must be very long. This leads to temporal artifacts that become audible at higher frequencies. Hence, systems in which the audio segment is decomposed into uniform sub-bands with adequate
25 low-frequency frequency resolution cannot take full advantage of the critical band properties of the auditory system.

 Prior art systems that recognize this limitation have attempted to solve the problem by utilizing analysis and synthesis filter banks based on QMF filter banks
30 that analyze a segment of an audio sound track to generate frequency components in two frequency bands. To obtain a decomposition of the segment into frequency components representing the amplitudes of the signal in critical bands, these two frequency band QMF filters are arranged in a tree-structured configuration. That is, each of the outputs of the first level filter becomes the input to another filter
35 bank at least one of whose two outputs is fed to yet another level, and so on. The leaf nodes of this tree provide an approximation to a critical band analysis of the

-4-

input audio track. It can be shown that this type of filter bank used different length audio segments to generate the different frequency components. That is, a low frequency component represents the signal amplitude in an audio segment that is much longer than a high-frequency component. Hence, the need to choose a
5 single compromise audio segment length is eliminated.

While tree structured filter banks having many layers may be used to decompose the frequency spectrum into critical bands, such filter banks introduce significant aliasing artifacts that limit their utility. In a multilevel filter bank, the
10 aliasing artifacts are expected to increase exponentially with the number of levels. Hence, filter banks with large numbers of levels are to be avoided. Unfortunately, filter banks based on QMF filters which divide the signal into two bandlimited signals require large numbers of levels.

15 Prior art audio compression systems are also poorly suited to applications in which the playback of the material is to be carried out on a digital computer. The use of audio for computer applications is increasingly in demand. Audio is being integrated into multimedia applications such as computer based entertainment, training, and demonstration systems. Over the course of the next
20 few years, many new personal computers will be outfitted with audio playback and recording capability. In addition, existing computers will be upgraded for audio with the addition of plug-in peripherals.

Computer based audio and video systems have been limited to the use of
25 costly outboard equipment such as an analog laser disc player for playback of audio and video. This has limited the usefulness and applicability of such systems. With such systems it is necessary to provide a user with a highly specialized playback configuration, and there is no possibility of distributing the media electronically. However, personal computer based systems using compressed
30 audio and video data promise to provide inexpensive playback solutions and allow distribution of program material on digital disks or over a computer network.

Until recently, the use of high quality audio on computer platforms has been limited due to the enormous data rate required for storage and playback.
35 Quality has been compromised in order to store the audio data conveniently on disk. Although some increase in performance and some reduction in bandwidth

-5-

has been gained using conventional audio compression methods, these improvements have not been sufficient to allow playback of high fidelity recordings on the commonly used computer platforms without the addition of expensive special purpose hardware.

5

One solution to this problem would be to use lower quality playback on computer platforms that lack the computational resources to decode compressed audio material at high fidelity quality levels. Unfortunately, this solution requires that the audio material be coded at various quality levels. Hence, each audio
10 program would need to be stored in a plurality of formats. Different types of users would then be sent the format suited to their application. The cost and complexity of maintaining such multi-format libraries makes this solution unattractive. In addition, the storage requirements of the multiple formats partially defeats the basic goal of reducing the amount of storage needed to store the audio material.

15

Furthermore, the above discussion assumes that the computational resources of a particular playback platform are fixed. This assumption is not always true in practice. The computational resources of a computing system are often shared among a plurality of applications that are running in a time-shared
20 environment. Similarly, communication links between the playback platform and shared storage facilities also may be shared. As the playback resources change, the format of the audio material must change in systems utilizing a multi-format compression approach. This problem has not been adequately solved in prior art systems.

25

Broadly, it the object of the present invention to provide an improved audio compression system.

It is a further object of the present invention to provide an audio
30 compression system which utilizes a frequency decomposition system that has good frequency resolution at low frequencies and good temporal resolution at high frequencies without utilizing tree structured filter banks having large numbers of levels.

35

It is yet another object of the present invention to provide an audio compression system that allows the compressed material to played back on a

-6-

variety of playback platforms with different computational capabilities without maintaining multiple copies of the compressed material.

5 It is a still further object of the present invention to provide an audio compression system in which the bandwidth needed to transmit the audio material may be varied in response to changes in the available bandwidth.

10 These and other objects of the present invention will become apparent from the following detailed description of the invention and the accompanying drawings.

Summary of the Invention

15 The present invention comprises audio compression and decompression systems. An audio compression system according to the present invention converts an audio signal into a series of sets of frequency components. Each frequency component represents an approximation to the audio signal in a corresponding frequency band over a time interval that depends on the frequency band. The received audio signal is analyzed in a tree-structured sub-band analysis
20 filter. The sub-band analysis filter bank comprises a tree-structured array of sub-band filters, the audio signal forming the input of the root node of the tree-structured array and the frequency components being generated at the leaf nodes of the tree-structured array. Each of the sub-band filter banks comprises a plurality of FIR filters having a common input for receiving an input audio signal.
25 Each filter generates an output signal representing the input audio signal in a corresponding frequency band, the number of FIR filters in at least one of the sub-band filter bank is greater than two, and the number of said FIR filters in at least one of the sub-band filters is different than the number of FIR filters in another of the sub-band filters. The frequency components generated by the sub-band
30 analysis filter are then quantized using information about the masking features of the human auditory system.

35 A decompression system according to the present invention regenerates a time-domain audio signal from the sets of frequency components such as those generated by a compression system according to the present invention. The decompression system receives a compressed audio signal comprising sets of

-7-

frequency components, the number of frequency components in each set being M. The decompression apparatus synthesizes M time domain audio signal values from each of the received set of frequency components. The synthesis sub-system generates 2M polyphase components from the set of frequency components. Then
5 it generates a W entry array from the polyphase phase components and multiplies each entry in the array by a corresponding weight value derived from a prototype filter. The time domain audio samples are then generated from the weighted array. The generated samples are stored in a FIFO buffer and outputted to a D/A converter. The FIFO buffer generates a signal indicative of the number of time
10 domain audio signal values stored therein. The rate at which these sample values are outputted to the D/A converters is determined by clock. The preferred embodiment of the decompression system includes a controller that uses the level indicator in the FIFO buffer or other operating system loading parameter to adjust the computational complexity of the algorithm used to synthesize the time domain
15 samples. When the level indicator indicates that the number of time domain samples stored in the FIFO buffer is less than a first predetermined value, the normal synthesis operation is replaced by one that generates an approximation to the time domain samples. This approximation requires a smaller number of computations than would be required to generate the time domain audio signal
20 values. The approximation may be generated by substituting a truncated or shorter prototype filter or by eliminating the contributions of selected frequency components from the computation of the polyphase components. In stereophonic systems, the controller may also switch the synthesis system to a monaural mode based on average frequency components which are obtained by averaging
25 corresponding frequency components for the left and right channels.

Brief Description of the Drawings

Figure 1 is a block diagram of an audio compression system.
30

Figure 2 is a block diagram of a sub-band decomposition filter according to the present invention.

Figure 3 illustrates the relationship between the length of the segment of
35 the original audio signal used to generate the frequency of each sub-band and the bandwidth of each band.

Figure 4 illustrates the relationship between successive overlapping segments of an audio signal.

5 Figure 5(a) is a block diagram of an audio filter based on a low-frequency filter and a modulator.

Figure 5(b) is a block diagram of a sub-band analysis filter for generating a set of frequency components.

10

Figure 6 illustrates the manner in which a sub-band analysis filter may be utilized to obtain the frequency information needed for psycho-acoustical analysis of the audio signal prior to quantization.

15 Figure 7 is a block diagram of an audio decompression system for decompressing the compressed audio signals generated by a compression system.

Figure 8 is a block diagram of a synthesizer according to the present invention.

20

Figure 9 is a block diagram of an audio decompression system utilizing the variable computational load techniques of the present invention.

25 Figure 10 is a block diagram of a stereophonic decompression system according to the present invention.

Figure 11 is a block diagram of a stereophonic decompression system according to the present invention using a serial computation system.

30 Figure 12 is a block diagram of an audio compression apparatus utilizing variable computational complexity.

Detailed Description of the Invention

The manner in which the present invention obtains its advantages over prior art audio compression systems may be more easily understood with reference to the manner in which a conventional audio compression system operates. Figure 1 is a block diagram of an audio compression system 10 using a conventional sub-band analysis system. The audio compression system accepts an input signal 11 which is divided into a plurality of segments 19. Each segment is analyzed by a filter bank 12 which provides the frequency components for the segment. Each frequency component is a time average of the amplitude of the signal in a corresponding frequency band. The time average is, in general, a weighted average. The frequencies of the sub-bands are uniformly distributed between a minimum and maximum value which depend on the number of samples in each segment 19 and the rate at which samples are taken. The input signal is preferably digital in nature; however, it will be apparent to those skilled in the art that an analog signal may be used by including an analog-to-digital converter prior to filter bank 12.

The component waveforms generated by filter bank 12 are replaced by digital approximations by quantizer 14. The number of bits assigned to each amplitude is determined by a psycho-acoustic analyzer 16 which utilizes information about the auditory system to minimize the distortions introduced by the quantization. The quantized frequency components are then further coded by coder 18 which makes use of the redundancy in the quantized components to further reduce the number of bits needed to represent the coded coefficients. Coder 18 does not introduce further errors into the frequency components. Coding algorithms are well known to those skilled in the signal compression arts, and hence, will not be discussed in more detail here.

The quantization process introduces errors into the frequency coefficients. A quantization scheme replaces the amplitude of each frequency component by an integer having a finite precision. The number of bits used to represent the integers will be denoted by P. The integers in question are then transmitted in place of the individual frequency components. At the receiver, the inverse of the mapping used to assign the integer values to the frequency components is used to produce amplitudes that are used in place of the original amplitudes for the frequency

components. There are at most 2^P distinct values that can be represented; hence, if there are more than 2^P different frequency component values, at least some of the frequency components will not be exactly recovered. The goal of the quantization algorithm is to minimize the overall effect of the quantization errors on the listener.

The errors introduced by the quantization algorithm affect the reconstructed audio track for a time period equal to the length of the segment analyzed to calculate the frequency components. The artifacts introduced by these errors are particularly noticeable in regions of the audio track in which the sound increases or decreases in amplitude over a period of time which is short compared to the length of the segments being analyzed. Because of the rapid rise, the set of frequency components of audio track in the segment will have a number of high-frequency components of significant amplitude which are not present in the segments on either side of the segment in question. Consider a quantization error in one of these high-frequency components. The error is equivalent to adding noise to the original signal. The amplitude of the noise will be determined by the quantization error. This noise will be present for the entire length of the segment in the reconstructed audio track. The noise resulting from the quantization error commences at the boundary of the segment even though the attack begins in the middle of the segment. The amplitude of the noise in the early part of segment may be of the same order of magnitude as the reconstructed audio track; hence, the noise will be particularly noticeable. Since the noise precedes the actual rise in intensity of the audio track, it is perceived as a "pre-echo". If the segment duration is long compared to the rise time of the audio signal, the pre-echo is particularly noticeable. Hence, it would be advantageous to choose filter bands in which the high-frequency components are calculated from segments that are shorter than those used to calculate the low-frequency components. This arrangement avoids the situation in which the segment used to compute high-frequency components is long compared to the rate of change of the component being computed.

Low bit rate audio compression systems operate by distributing the noise introduced by quantization so that it is masked by the signal. The ear's sensitivity to a noise source in the presence of a localized frequency component such as a sine tone depends on the relative levels of the signals and on the relation of the

-11-

noise spectral components to the tone. The noise becomes significantly less audible if its spectral energy is within one critical bandwidth of the tone. Hence, it would be advantageous to choose filter bands that more closely match the critical bands of the human auditory system.

5

The present invention utilizes a filter bank in which different frequency bands utilize different segment lengths. In prior art systems, each segment is analyzed in a bank of finite impulse response filters. The number of samples in the input segment over which each frequency component is computed is the same.

10 The present invention uses different width segments for different frequency components. Ideally, an audio decomposition should exhibit a time and frequency dependency similar to that of human hearing. This may be accomplished by relating the frequency divisions or sub-bands of the decomposition to the critical bandwidths of human hearing. The resulting decomposition has fine frequency
15 resolution with relatively poor temporal resolution at low frequencies, and coarse frequency resolution with fine temporal resolution at high frequencies. As a result, the segment length corresponding to high-frequency components does not greatly exceed the rise time of attacks in the audio track. This reduces the pre-echo artifacts discussed above.

20

In one embodiment of the present invention, a tree structured decomposition which approximates the ear's time and frequency sensitivity is utilized. This filter may be used to replace sub-band analysis filter bank 12 shown in Figure 1. A block diagram of a sub-band decomposition filter for carrying out
25 this decomposition is shown at 30 in Figure 2. Filter 30 includes two levels of filter banks. The manner in which the filter banks are constructed will be discussed in more detail below. For the purposes of the present discussion, it is important to note that the decomposition is carried out with only two levels of filters, and hence, avoids the aliasing problems inherent in QMF filter banks that
30 require many levels. The aliasing problems encountered with QMF filter banks become significant when the number of levels exceeds 4.

The first level of filter 30 consists of a filter bank 31 which divides the input signal into eight sub-bands of equal size. The second level sub-divides the
35 lowest three frequency bands from filter bank 31 into finer sub-divisions. The second level consists of three filter banks 32-34. Filter bank 32 divides the lowest

-12-

sub-band from filter bank 31 into 8 equal sub-bands. Filter bank 33 and filter bank 34 divide the second and third sub-bands created by filter bank 31 into four sub-bands. The combination of the two levels generates 21 frequency sub-bands. The relationship between the length of the segment of the original audio signal
5 used to generate the frequency and phase of each sub-band and the bandwidth of each band is shown schematically in Figure 3. The lower frequencies, bands 1-8, have the finest frequency resolution, but the poorest temporal resolution. The highest frequencies, bands 17-21, have the poorest frequency resolution, but the finest time resolution. This arrangement more nearly approximates the ear's
10 sensitivity than systems utilizing filter banks in which all bands have the same temporal resolution, while avoiding the aliasing problems inherent in tree-structured filters having many levels of filters.

While quantization errors in each of the amplitudes still introduces noise,
15 the noise spectrum obtained with this embodiment of the present invention is less objectionable to a human listener than that obtained with prior art systems. As noted above, prior art systems tend to have a noise spectrum which changes abruptly at the segment boundaries. In the present invention, the amplitude of the quantization noise can switch more rapidly at higher frequencies. If the length of
20 the low frequency segments is denoted by T , then the medium frequencies are measured on segments that are $T/2$, and the highest frequencies are measured on segments that are $T/8$ in length. The quantization noise is the sum of all of the quantization errors in all of the frequency bands. As a result, the quantization noise changes every $T/8$. To obtain the same resolution in the low frequency
25 components, a conventional filter system would measure all of the frequency components on segments of length T . Hence, the prior art would introduce quantization noise which changes abruptly every T samples. The present invention introduces a more gradual change in the noise level in the $T/8$ interval for the high and medium sub-bands thus giving less perceptible distortion at
30 higher frequencies.

The manner in which the input signal is divided into segments can effect the quality of the regenerated audio signal. Consider the case in which the signal is analyzed on segments that do not overlap. This analysis is equivalent to
35 employing a model in which the regenerated signal is produced by summing the signals of a number of harmonic oscillators whose amplitudes remain constant

-13-

over the duration of the segment on which each amplitude was calculated. In general, this model is a poor approximation to an actual audio track. In general, the amplitudes of the various frequency components would be expected to change over the duration of the segments in question. Models that do not take this change into account will have significantly greater distortions than models in which the amplitudes can change over the duration of the segment, since there will be abrupt changes in the amplitudes of the frequency components at each segment boundary.

One method for reducing the discontinuities in the frequency component amplitudes at the segment boundaries is to employ a sub-band analysis filter that utilizes overlapping segments to generate successive frequency component amplitudes. The relationship of the segments is shown in Figure 4 for a signal 301. The sub-band analysis filter generates M frequency components for signal 301 for each M signal values. However, each frequency component is generated over a segment having a duration much greater than M. Each component is generated over a segment having a length of W sample values, where $W > M$. Typical segments are shown at 312 and 313. It should be noted that successive segments overlap by (W-M) samples.

In the preferred embodiment of the present invention, the various frequency bands in a sub-band analysis filter bank have the same shape but are shifted relative to one another. This arrangement guarantees that all frequency bands have the same aliasing properties. Such a filter bank can be constructed from a single low frequency band pass filter having the desired band shape. The manner in which the various filter bands are constructed is most easily understood with reference to Figure 5(a) which is a block diagram of a single filter constructed from a low-frequency bandpass filter 377 and a mixer 376. Assume that the low-pass filter 377 has a center frequency of F_c and that the desired center frequency of filter 350 is to be F. Then by shifting the input audio signal by a frequency of $F - F_c$ prior to analyzing the signal with low-frequency bandpass filter 377, the output of low-frequency bandpass filter 377 will be the amplitude of the audio signal in a band having a center frequency of F. Modulator 376 accomplishes this frequency shift.

-14-

A filter bank can then be constructed from a single prototype low-frequency bandpass filter by using different modulation frequencies to shift the incoming audio signal prior to analysis by the prototype filter. While such a filter bank can be constructed from analog circuit components, it is difficult to obtain
 5 filter performance of the type needed. Hence, the preferred embodiment of the present invention utilizes digital filter techniques.

A block diagram of a sub-band analysis filter 350 for generating a set of M frequency components, S_i , from a W sample window is shown in Figure 5(b).
 10 The M audio samples are clocked into a W-sample shift register 320 by controller 325. The oldest M samples in shift register 320 are shifted out the end of the shift register and discarded. The contents of the shift register are then used to generate 2M polyphase components P_k , for $k=0$ to $2M-1$. The polyphase components are generated by a windowing operation followed by partial summation. The
 15 windowing operation generates a W- component array Z_i from the contents of shift register 320 by multiplying each entry in the shift register by a corresponding weight, i.e.,

$$Z_i = h_i * x_i \quad (1)$$

20 where the x_i , for $i=0..W-1$ are the values stored in shift register 320, and the h_i are coefficients of a low pass prototype filter which are stored in controller 325. For those wishing a more detailed explanation of the process for generating sets of filter coefficients, see J. Rothweiler, "POLYPHASE QUADRATURE FILTERS -
 25 A NEW SUB-BAND CODING TECHNIQUE" *IEEE Proceedings of the 1983 ICASSP Conference*, pp 1280-1283. The polyphase components are then generated from the Z_i by the following summing operations:

$$P_k = \sum_{j=0}^{2M} Z_{i+2Mj} \quad (2)$$

30

The frequency components, S_i , are obtained via the following matrix multiplication from the polyphase components

$$S_i = \sum_{k=0}^{2M-1} P_k \cos\left[\frac{(2i+1)(k-M/2)\pi}{2M}\right] \quad (3a)$$

-15-

This operation is equivalent to passing the polyphase components through M finite impulse response filters of length 2M. The cosine modulation of the polyphase components shown in Eq. (3a) may be replaced by other such modulation terms. The form shown in Eq. (3a) leads to near-perfect reconstruction. An alternative modulation scheme which allows for perfect reconstruction is as follows:

$$S_i = \sum_{k=0}^{2M-1} P_k \cos\left[\frac{(2i+1)(2k+1-M)\pi}{4M}\right] \quad (3b)$$

It can be seen by comparison to Figure 5(a) that the matrix multiplication provides an operation analogous to the modulation of the incoming audio signal. The windowing operation performs the analysis with the prototype low-frequency filter.

As will be discussed in more detail below, the computational workload in analyzing and synthesizing audio tracks of a great importance in providing systems that can operate on general purpose computing platforms. It will be apparent from the above discussion that the computational workload inherent in generating M frequency components from a window of W audio sample values is approximately $(W+2M^2)$ multiplies and adds. In this regard, it should be noted that a two level filter bank of the type used in the present invention significantly reduces the overall computational workload even in situations in which the frequency spectrum is to be divided into uniform bands. For example, consider a system in which the frequency spectrum is to be divided into 64 bands utilizing a window of 512 samples. If a prior art one level filter bank is utilized, the workload will be approximately 8,704 multiplies and adds. If the filter bank is replaced by a two level filter bank according to the present invention, then the filter bank will consist of 9 filter banks, each dividing the frequency spectrum into 8 bands. The computational workload inherent in this arrangement is only 5,760 multiplies and adds. Hence, a filter bank according to the present invention typically requires less computational capability than a one level filter bank according to the prior art. In addition, a filter bank according to the present invention also provides a means for providing a non-uniform band structure.

The transformation of the audio signal into sets of frequency components as described above does not, in itself, result in a decrease in the number of bits needed to represent the audio signal. For each M audio samples received by a sub-band analysis filter, M frequency components are generated. The actual
 5 signal compression results from the quantization of the frequency components. As noted above, the number of bits that must be allocated to each frequency component is determined by a phenomena known as "masking". Consider a tone at a frequency f . The ability of the ear to detect a signal at frequency f' depends on the energy in the tone and difference in frequency between the signal and the tone,
 10 i.e., $(f-f')$. Research in human hearing has led to measurements of a threshold function $T(E,f,f')$ which measures the minimum energy at which the second frequency component can be detected in the presence to the first frequency component with energy E . In general, the threshold function will vary in shape with frequency.

15

The threshold function is used to construct a masking function as follows. Consider a segment of the incoming audio signal. Denote the energy as a function of frequency in this segment by $E(f)$. Then a mask level, $L(f)$, is constructed by convolving $E(f)$ and $T(f,f')$, i.e.,

20

$$L(f) = \int T(E(f')f,f') E(f') df' \quad (4)$$

Consider the filtered signal value in a band $f_0 \pm \Delta f$. Denote the minimum value of L in this frequency band by $L_{\min}$. It should be noted that $L_{\min}$ may depend on
 25 frequency components outside the band in question, since a peak in an adjacent band may mask a signal in the band in question.

According to the masking model, any noise in this frequency band that has an energy less than $L_{\min}$ will not be perceived by the listener. In particular, the
 30 noise introduced by replacing the measured signal amplitude in this band by a quantized approximation therefore will not be perceived if the quantization error is less than $L_{\min}$. The noise in question will be less than $L_{\min}$ if the signal amplitude is quantized to accuracy equal to $S/L_{\min}$, where S is the energy of the signal in the band in question.

35

The above-described quantization procedure requires a knowledge of frequency spectrum of the incoming audio signal at a resolution which is significantly greater than that of the sub-analysis of the incoming signal. In general, the minimum value of the mask function L will depend on the precise location of any peaks in the frequency spectrum of the audio signal. The signal amplitude provided by the sub-band analysis filter measures the average energy in the frequency band; however, it does not provide any information about the specific location of any spectral peaks within the band.

Hence, a more detailed frequency analysis of the incoming audio signal is required. This can be accomplished by defining a time window about each filtered signal component and performing a frequency analysis of the audio samples in this window to generate an approximation to $E(f)$. In prior art systems, the frequency analysis is typically performed by calculating a FFT of the audio samples in the time window.

In one embodiment of a quantization sub-component according to the present invention, this is accomplished by further subdividing each sub-band using another layer of filter banks. The output of each of the sub-band filters in the analysis filter bank is inputted to another sub-band analysis filter which splits the original sub-band into a plurality of finer sub-bands. These finer sub-bands provide a more detailed spectral measurement of the audio signal in the frequency band in question, and hence, can be used to compute the overall mask function L discussed above.

While a separate $L_{\min}$ value may be calculated for each filtered signal value from each sub-band filter, the preferred embodiment of the present invention operates on blocks of filtered signal values. If a separate quantization step size is used for each filtered value, then the step size would need to be communicated with each filtered value. The bits needed to specify the step size reduce the degree of compression. To reduce this "overhead", a block of samples is quantized using the same step size. This approach reduces the number of overhead bits/sample, since the step size need only communicated once. The blocks of filtered samples utilized consist of a sequential set of filtered signal values from one of the sub-band filters. As noted above, these values can be inputted to a second sub-band analysis filter to obtain a fine spectral measurement of the energy in the sub-band.

One embodiment of such a system is shown in Figure 6 at 400. The audio signal values are input to a sub-band analysis filter 402 which is similar to that shown in Figure 2. The filtered outputs are quantized by quantizer 404 in blocks of 8 values. Each set of 8 values leaving sub-band analysis filter 402 is processed by a sub-band analysis filter 408 to provide a finer spectral measurement of the audio signal. Sub-band analysis filters 408 divide each band into 8 uniform sub-bands. The outputs of sub-band analysis filters 408 are then used by psycho-acoustic analyzer 406 to determine the masking thresholds for each of the frequency components in the block. While the above embodiment splits each band into 8 subbands for the purpose of measuring the energy spectrum, it will be apparent to those skilled in the art that other numbers of subbands may be used. Furthermore, the number of subbands may be varied with the frequency band.

The manner in which an audio decompression system according to the present invention operates will now be explained with the aid of Figure 7 which is a block diagram of an audio decompression system 410 for decompressing the compressed audio signals generated by a compression system such as that shown in shown in Figure 6. The compressed signal is first decoded to recover the quantized signal values by a decoder 412. The quantized signal values are then used to generate approximations to the filtered signal values by de-quantizer 414. Since the present invention utilizes multi-rate sampling, the number of filtered signal values depends on the specific frequency bands. In the case in point, there are 21 such bands. As discussed above, the five highest bands are sampled at 8 times the rate of the lowest 8 frequency bands, and the intermediate frequency bands are sampled at twice the rate of the lowest frequency bands. The filtered signal values are indicated by kS_m , where m indicates the frequency band, and k indicates the number of the signal value relative to the lowest frequency bands, i.e., k runs from 1 to 8 for the highest frequency bands, and 1 to 2 for the intermediate frequency bands.

The filtered samples are inputted to an inverse sub-band filter 426 which generates an approximation to the original audio signal from the filtered signal values. Filter 402 shown in Figure 6 and filter 426 form a perfect, or near perfect, reconstruction filter bank. Hence, if the filtered samples had not been replaced by approximations thereto by quantizer 404, the decompressed signal generated by

-19-

filter bank 426 would exactly match the original audio signal input to filter 402 to a specified precision.

Inverse sub-band filter bank 426 also comprises a tree-structured filter bank. To distinguish the filters used in the inverse sub-band filters from those used in the sub-band filter banks which generated the filtered audio samples, the inverse filter banks will be referred to as synthesizers. The filtered signal values enter the tree at the leaf nodes thereof, and the reconstructed audio signal exits from the root node of the tree. The low and intermediate filtered samples pass through two levels of synthesizers. The first level of synthesizers are shown at 427 and 428. For each group of four filtered signal values accepted by synthesizers 427 and 428, four sequential values which represent filtered signal values in a frequency band which is four times wider are generated. Similarly, for each group of eight filtered signal values accepted by synthesizer 429, eight sequential values which represent filtered signal values in a frequency band which is eight times as wide are generated. Hence, the number of signal values entering synthesizer 430 on each input is now the same even though the number of signal values provided by de-quantizer 414 for each frequency band varied from band to band.

20

The synthesis of the audio signal from the sub-band components is carried out by analogous operations. Given M sub-band components that were obtained from 2M polyphase components P_i , the original polyphase components can be obtained from the following matrix multiplication:

25

$$P_i = \sum_{k=0}^{M-1} S_k \cos\left[\frac{(i+\frac{M}{2})(2k+1)\pi}{2M}\right] \quad (5a)$$

As noted above, there are a number of different cosine modulations that may be used. Eq. (5a) corresponds to modulation using the relationship shown in Eq. 3(a). If the modulation shown in Eq. 3(b) is utilized, then the polyphase components are obtained from the following matrix multiplication:

30

$$P_i = \sum_{k=0}^{M-1} S_k \cos\left[\frac{(2i+1+M)(2k+1)\pi}{4M}\right] \quad (5b)$$

35

-20-

The time domain samples x_k are computed from the polyphase components by the inverse of the windowing transform described above. A block diagram of a synthesizer according to the present invention is shown in Figure 8 at 500. The M frequency components are first transformed into the corresponding polyphase components by a matrix multiplication shown at 510. The resultant 2M polyphase components are then shifted into a 2W entry shift register 512 and the oldest 2M values in the shift register are shifted out and discarded. The contents in the shift register are inputted to array generator 513 which builds a W value array 514 by iterating the following loop 8 times: take the first M samples from shift register 512, ignore the next 2M samples, then take the next M samples. The contents of array 514 are then multiplied by W weight coefficients, h'_i which are related to the h_i used in the corresponding sub-band analysis filter to generate a set of weighted values $w_i = h'_i * u_i$, which are stored in array 516. Here the u_i are the contents of array 514. The M time domain samples, x_j for $j=0, \dots, M-1$, are then generated by summing circuit 518 which sums the appropriate w_i values, i.e.,

$$x_j = \sum_{i=0}^{W/M-1} w_{j+Mi} \quad (6)$$

While the above-described embodiments of synthesizers and sub-band analysis filters are described in terms of special purpose hardware for carrying out the various operations, it will be apparent to those skilled in the art that the entire operation may be carried out on a general purpose digital computer.

As pointed out above, it would be advantageous to provide a single high-quality compressed audio signal that could be played back on a variety of playback platforms having varying computational capacities. Each such playback platform would reproduce the audio material at a quality consistent with the computational resources of the platform.

Furthermore, the quality of the playback should be capable of being varied in real time as the computational capability of the platform varies. This last requirement is particularly important in playback systems comprising multi-tasking computers. In such systems, the available computational capacity for the audio material varies in response to the computational needs of tasks having equal

-21-

or higher priority. Prior art decompression systems do not provide this capability.

The present invention allows the quality of the playback to be varied in response to the computational capability of the playback platform without the use of multiple copies of the compressed material. Consider an audio signal that has been compressed using a sub-band analysis filter bank in which the window contains W audio samples. The computational workload required to decompress the audio signal is primarily determined by the computations carried out by the synthesizers. The computational workload inherent in a synthesizer is W multiplies and adds from the windowing operations and $2M^2$ multiplies and adds from the matrix multiplication. The extent to which the filters approximate an ideal band pass filter, in general, depends on the number samples in the window, i.e., W . As the number of samples increases, the discrepancy between the sub-band analysis filter performance and that of an ideal band pass filter decreases. For example, a filter utilizing 128 samples has a side lobe suppression in excess of 48 dB, while a filter utilizing 512 samples has a side lobe suppression in excess of 96 dB. Hence, synthesis quality can be traded for a reduction in computational workload if a smaller window is used for the synthesizers.

20

In the preferred embodiment of the present invention, the size of the window used to generate the sub-band analysis filters in the compression system is chosen to provide filters having 96 dB rejection of signal energy outside a filter band. This value is consistent with playback on a platform having 16 bit D/A converters. In the preferred embodiment of the present invention, this condition can be met by 512 samples. The prototype filter coefficients, h_i , viewed as a function of i have a more or less sinc-shaped appearance with tails extending from a maximum. The tails provide the corrections which result in the 96 dB rejection. If the tails are truncated, the filter bands would have substantially the same bandwidths and center frequencies as those obtained from the non-truncated coefficients. However, the rejection of signal energy outside a specific filter's band would be less than the 96 dB discussed above. As a result, a compression and decompression system based on the truncated filter would show significantly more aliasing than the non-truncated filter.

35

-22-

The present invention utilizes this observation to trade sound quality for a reduction in computational workload in the decompression apparatus. In the preferred embodiment of the present invention, the audio material is compressed using filters based on a non-truncated prototype filter. When the available
5 computational capacity of the playback platform is insufficient to provide decompression using synthesis filters based on the non-truncated prototype filter, synthesizers based on the truncated filters are utilized. Truncating the prototype filter leads to synthesizers which have the same size window as those based on the non-truncated prototype. However, many of the filter coefficients used in the
10 windowing operation are zero. Since the identity of the coefficients which are now zero is known, the multiplications and additions involving these coefficients can be eliminated. It is the elimination of these operations that provides the reduced computational workload.

15 It should be noted that many playback platforms use D/A converters with less than 16 bits. In these cases, the full 96dB rejection is beyond the capability of the platform; hence, the system performance will not be adversely effected by using the truncated filter. These platforms also tend to be the less expensive computing systems, and hence, have lower computational capacity. Thus, the
20 trade-off between computational capacity and audio quality is made at the filter level, and the resultant system provides an audio quality which is limited by its D/A converters rather than its computational capacity.

Another method for trading sound quality for a reduction in computational
25 workload is to eliminate the synthesis steps that involve specific high-frequency components. If the sampled values in one or more of the high-frequency bands are below some predetermined threshold value, then the values can be replaced by zero. Since the specific components for which the substitution is made are known, the multiplications and additions involving these components may be eliminated,
30 thereby reducing the computational workload. The magnitude of the distortion generated in the reconstructed audio signal will, of course, depend on the extent of the error made in replacing the sampled values by zeros. If the original values were small, then the degradation will be small. This is more often the case for the high-frequency filtered samples than for the low frequency filtered samples. In
35 addition, the human auditory system is less sensitive at high frequencies; hence, the distortion is less objectionable.

It should also be noted that the computational workload inherent in decompressing a particular piece of audio material varies during the material. For example, the high-frequency filtered sampled may only have a significant amplitude during parts of the sound track. When the high-frequency components are not present or sufficiently small to be replaced by zeros without introducing noticeable distortions, the computational workload can be reduced by not performing the corresponding multiplications and additions. When the high-frequency components are large, e.g., during attacks, the computational workload is much higher.

It should be noted that the computational work associated with generating the P_k values from the S_i values can be organized by S_i . That is, the contribution to each P_k from a given S_i is calculated, then the contribution to each P_k from S_{i+1} , and so on. Since there are $2M$ P values involved with each value of S , the overhead involved in testing each value of S before proceeding with the multiplications and additions is small compared to the computations saved if a particular S value is 0 or deemed to be negligible. In the preferred embodiment of the present invention, the computations associated with S_i are skipped if the absolute value of S_i is less than some predetermined value, ϵ .

Because of the variation in workload, the preferred embodiment of the present invention utilizes a buffering system to reduce the required computational capacity from that needed to accommodate the peak workload to that need to accommodate the average workload. In addition, this buffering facilitates the use of the above-described techniques for trading off the required computational capacity against sound quality. For example, when the computational workload is determined to be greater than that available, the value of ϵ can be increased which, in turn, reduces the number of calculations needed to generate the P_k values.

A block diagram of an audio decompression system utilizing the above-described variable computational load techniques is shown in Figure 9 at 600. The incoming compressed audio stream is decoded by decoder 602 and de-quantizer 604 to generate sets of frequency components $\{S_i\}$ which are used to reconstruct the time domain audio signal values. The output of synthesizer 606 is loaded into a FIFO buffer 608 which feeds a set of D/A converters 610 at a

-24-

constant rate determined by clock 609. The outputs of the D/A converters are used to drive speakers 612. Buffer 608 generates a signal that indicates the number of time domain samples stored therein. This signal is used by controller 614 to adjust the parameters that control the computational complexity of the synthesis operations in synthesizer 606. When this number falls below a predetermined minimum value, the computational algorithm used by synthesizer 606 is adjusted to reduce the computational complexity, thereby increasing the number of time domain samples generated per unit time. For example, controller 614 can increase the value of ϵ described above. Alternatively, controller 614 could force all of the high-frequency components from bands having frequencies above some predetermined frequency to be zero. In this case, controller 614 also instructs de-quantizer 604 not to unpack the high-frequency components that are not going to be used in the synthesis of the signal. This provides additional computational savings. Finally, controller 614 could change the windowing algorithm, i.e., use a truncated prototype filter.

If the number of stored values exceeds a second predetermined value, controller 614 adjusts the computational algorithm to regain audio quality if synthesizer 606 is not currently running in a manner that provides the highest audio quality. In this case, controller 614 reverses the approximations introduced into synthesizer 606 discussed above.

While audio decompression system 600 has been discussed in terms of individual computational elements, it will be apparent to those skilled in the art that the functions of decoder 602, de-quantizer 604, synthesizer 606, buffer 608 and controller 614 can be implemented on a general purpose digital computer. In this case, the functions provided by clock 609 may be provided by the computer's clock circuitry.

In stereophonic decompression systems having parallel computational capacity, two synthesizers may be utilized. A stereophonic decompression system according to the present invention is shown in Figure 10 at 700. The incoming compressed audio signal is decoded by a decoder 702 and de-quantized by de-quantizer 704 which generates two sets of frequency components 705 and 706. Set 705 is used to regenerate the time domain signal for the left channel with the aid of synthesizer 708, and set 706 is used to generate the time domain signal for

the right channel with the aid of synthesizer 709. The outputs of the synthesizers are stored in buffers 710 and 712 which feed time domain audio samples at regular intervals to D/A converters 714 and 715, respectively. The timing of the signal feed is determined by clock 720. The operation of decompression system
5 700 is controlled by a controller 713 which operates in a manner analogous to controller 614 described above.

If a stereophonic decompression system does not have parallel computational capacity, then the regeneration of the left and right audio channels
10 must be carried out by time-sharing a single synthesizer. When the computational workload exceeds the capacity of the decompression system, the trade-offs discussed above may be utilized to trade audio quality for a reduction in the computational workload. In addition, the computational workload may be reduced by switching to a monaural reproduction mode, thereby reducing the
15 computational workload imposed by the synthesis operations by a factor of two.

A stereophonic decompression system using this type of serial computation system is shown in Figure 11 at 800. The incoming compressed audio signal is decoded by a decoder 802 and de-quantized by de-quantizer 804
20 which generates sets of frequency components for use in synthesizing the left and right audio signals. When there is sufficient computational capacity available to synthesize both left and right channels, controller 813 time shares synthesizer 806 with the aid of switches 805 and 806. When there is insufficient computational capacity, controller 813 causes switch 805 to construct a single set of frequency
25 components by averaging the corresponding frequency components in the left and right channels. The resulting set of frequency components is then used to synthesize a single set of monaural time domain samples which is stored in buffers 810 and 812.

30 The techniques described above for varying the computational complexity required to synthesize a signal may also be applied to vary the computational complexity required to analyze a signal. This is particularly important in situations in which the audio signal must be compressed in real time prior to being distributed through a communication link having a capacity which is less than that
35 needed to transmit the uncompressed audio signal. If a computational platform

having sufficient capacity to compress the audio signal at full audio quality is available, the methods discussed above can be utilized.

However, there are situations in which the computational capacity of the compression platform may be limited. This can occur when the computational platform has insufficient computing power, or in cases in which the platform performing the compression may also include a general purpose computer that is time-sharing its capacity among a plurality of tasks. In the later case, the ability to trade-off computational workload against audio quality is particularly important.

A block diagram of an audio compression apparatus 850 utilizing variable computational complexity is shown in Figure 12 at 850. Compression apparatus 850 must provide a compressed signal to a communication link. For the purposes of this discussion, it will be assumed that the communication link requires a predetermined amount of data for regenerating the audio signal at the other end of the communication link. Incoming audio signal values from an audio source such as microphone 852 are digitized and stored in buffer 854. In the case of stereophonic systems, a second audio stream is provided by microphone 851. To simplify the following discussion, it will be assumed that apparatus 850 is operating in a monaural mode unless otherwise indicated. In this case, only one of the microphones provides signal values.

When M such signal values have been received, subband analysis filter bank 856 generates M signal components from these samples while the next M audio samples are being received. The signal components are then quantized and coded by quantizer 858 and stored in an output buffer 860. The compressed audio signal data is then transmitted to the communication link at a regular rate that is determined by clock 862 and controller 864.

Consider the case in which subband analysis filter 856 utilizes a computational platform that is shared with other applications running on the platform. When the computational capacity is restricted, sub-band analysis filter bank 856 will not be able to process incoming signal values at the same rate at which said signal values are received. As a result, the number of signal values stored in buffer 854 will increase. Controller 864 periodically senses the number of values stored in buffer 854. If the number of values exceeds a predetermined

-27-

number, controller 864 alters the operations of subband analysis filter bank 856 in a manner that decreases the computational workload of the analysis process. The audio signal synthesized from the resulting compressed audio signal will be of lesser quality than the original audio signal; however, compression apparatus 850
5 will be able to keep up with the incoming data rate. When controller 864 senses that the number of samples in buffer 854 returns to a safe operating level, it alters the operation of subband analysis filter bank 856 in such a manner that the computational workload and audio quality increases.

10 The many of the techniques described above may be used to vary the computational workload of the subband analysis filter. First, the prototype filter may be replaced by a shorter filter or a truncated filter thereby reducing the computational workload of the windowing operation. Second, the higher frequency signal components can be replaced by zero's. This has the effect of
15 reducing "M" and thereby reducing the computational workload.

Third, in stereophonic systems, the audio signals from each of the microphones 851 and 852 can be combined by circuitry in buffer 854 to form a monaural signal which is analyzed. The compressed monaural signal is then used
20 for both the left and right channel signals.

While the above embodiments of the present invention utilize the level in a buffer as an indicator of whether or not the computational capacity of the analysis or synthesis filters is sufficient, it will be apparent to those skilled in the art that
25 other signals may be used for this feedback loop. For example, in multi-tasking systems in which the synthesis or analysis filters are one of the tasks, the system can provide a signal indicating the level of computational resources available to the filters. The controller can then adjust the computational complexity of the filtering operation to match the available resources.

30

Accordingly, there has been described herein a novel audio compression system. Various modifications to the present invention will become apparent to those skilled in the art from the foregoing description and accompanying drawings. Hence, the present invention is to be limited solely by the scope of the
35 following claims.

WHAT IS CLAIMED IS:

1. An audio compression apparatus comprising: means for receiving an audio signal; sub-band analysis filter means for generating a M frequency
5 components from said received audio signal, each said frequency component representing the amplitude of said audio signal in a corresponding frequency band, said sub-band analysis filter means comprising a tree-structured array of sub-band filter banks, said audio signal forming the input of the root node of said tree-
10 structured array and said frequency components being generated by the leaf nodes of said tree-structured array, each said sub-band filter bank comprising a plurality of FIR filters having a common input for receiving an input audio signal, each said filter generating an output signal representing the amplitude of said input audio
15 signal in a corresponding frequency band, the number of said FIR filters in at least one of said sub-band filters being greater than two, wherein the number of said FIR filters in at least one of said sub-band filter banks is different than the number of said FIR filters in another of said sub-band filter banks; and quantization means for replacing each of said frequency components by one of a plurality of possible approximations thereto.
- 20 2. The audio compression apparatus of Claim 1 wherein said receiving means comprises shift register means for storing the most recently received W signal values and means for determining that M new audio signals values have been received, and wherein said sub-band analysis filter banks comprise: means
25 for generating 2M polyphase components from the contents of said shift register means; and means for multiplying a vector comprising said polyphase components by a matrix to generate said M frequency components.
3. The audio compression apparatus of Claim 1 wherein said quantization
30 means further comprises M sub-band filters, each said sub-band filter having one of said frequency components as the input thereof, each said sub-band filter generating a plurality of frequency components representing the amplitude of said input in a corresponding frequency band.
- 35 4. The audio compression apparatus of Claim 1 wherein at least one of said filters comprises a plurality of alternative filters, each said alternative filter requiring a different computational workload and wherein said audio compression

-29-

apparatus further comprises: means for receiving a signal specifying the computational resources available for operating said sub-band analysis filter banks; means for periodically testing said signal to determine if the available computational resources are greater than a predetermined value; and means for
5 selecting one of said alternative filters if said testing means determines that the computational resources available are less than the computational resources needed to filter said input audio signal at a predetermined rate.

5. An audio decompression apparatus for generating a first time domain
10 audio signal comprising a first set of time domain audio signal values from a plurality of sets of first frequency components, each said first frequency component representing the average amplitude of said time domain audio signal in a corresponding frequency band over a corresponding time period, said apparatus comprising: means for receiving a compressed audio signal comprising said sets
15 of first frequency components, the number of frequency components in each said set being M; synthesis means for generating M time domain audio signal values from each said received set of first frequency components, said synthesis means comprising means for generating 2M polyphase components from said set of first frequency components; means for generating a W entry array from said polyphase
20 phase components, and multiplying each entry in said array by a corresponding weight value derived from a prototype filter; FIFO buffer means for storing said time domain audio signal values and for outputting the same to a D/A converter, said FIFO buffer means further comprising level indicator means for generating a signal indicative of the number of time domain audio signal values stored therein;
25 clock means for controlling the rate at which said FIFO buffer means outputs said time domain audio signal values; and control means responsive to said level indicator means for causing said synthesis means to generate an approximation to said time domain audio signal values if said indicator means indicates that the number of time domain samples stored in said FIFO buffer means is less than a
30 first predetermined value, said approximation requiring a lower number of computations than would be required to generate said time domain audio signal values.

6. The audio decompression apparatus of Claim 5 wherein said control
35 comprises means for causing said synthesis means to generate an approximation to

-30-


said 2M polyphase components in place of said 2M polyphase components that would have been generated from each said set of frequency components.

7. The audio decompression apparatus of Claim 6 wherein said control
5 means further comprises means for storing a plurality of prototype filters; and means, responsive to said level indicator means, for selecting one of said stored prototype filters for use by said synthesis means.

8. The audio decompression apparatus of Claim 5 wherein said synthesis
10 means comprises: a tree-structured array of sub-band synthesizers, each said sub-band synthesizer combining a plurality of frequency components representing the amplitude of an audio signal in a corresponding frequency band to generate a plurality of time domain samples representing the amplitude of an audio signal in a frequency band comprising the corresponding frequency bands of said frequency
15 components, each said set of frequency components comprising the inputs to the leaf nodes of said tree-structured array and said time domain audio signal values comprising the output of the root node of said tree-structured array, the number of inputs of at least one of said sub-band synthesizers being greater than 2, wherein the number of said inputs of at least one of said sub-band synthesizers is different
20 from the number of said inputs of another of said sub-band synthesizers.

9. The audio decompression apparatus of Claim 5 wherein said received audio signal further comprises a second set of frequency components corresponding to each said set of frequency components, wherein said synthesis
25 means further comprises means for generating a second set of time domain audio signal values from said second set of frequency components, and wherein said control means further comprises means, responsive to said level indicator means, for generating an averaged set of frequency components by averaging corresponding frequency components in said first and second sets of frequency
30 components; and means for causing said first and second set of time domain audio signal values to be generated from said averaged frequency components.

1/11






FIGURE 2

3/11




FIGURE 3




FIGURE 4




FIGURE 5(a)




FIGURE 5(b)

5/11




FIGURE 6




FIGURE 7

7/11




FIGURE 8

8/11




FIGURE 9

9/11




FIGURE 10

10/11




FIGURE 11




FIGURE 12

INTERNATIONAL SEARCH REPORT

International application No.

PCT/US93/08572

A. CLASSIFICATION OF SUBJECT MATTER

IPC(5) : G10L 3/02, 7/04, 9/18

US CL : 395/2.14, 2.15, 2.39, 2.74

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

U.S. : 395/2.14, 2.15, 2.39, 2.74; 381/29, 37, 45, 51; 364/724.01, 724.1

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)

APS: Quantiz?, tree, frequenc?, polyphase, poly phase, buffer, FIR, finite impulse response, speech, voice, music, audio.

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X ----- Y	US, A, 4,896,356 (Millar) 23 January 1990, Fig. 4, col. 4, col. 7 and col. 10	1, 3, 4 ----- 2, 5-9
Y	IEEE ASSP Magazine, July 1987, P.P. Vaidyanathan, "Quadrature Mirror Filter Banks, M-Band Extensions and Perfect-Reconstruction Techniques", pages 4-20, especially pages 9, 10 and 15.	2, 5-9
Y	US, A, 4,815,023 (Arbeiter) 21 March 1989, figs. 1, 4, 5, abstract, col. 1.	1, 2, 4, 5, 6, 8
Y	US, A, 4,799,179 (Masson et al.) 17 January 1989, Figs. 4a, 4b, 5-8, col. 2.	2, 5, 6
Y	US, A, 4,949,383 (Koh et al.) 14 August 1990, Fig. 1.	2

☒ Further documents are listed in the continuation of Box C.
 ☐ See patent family annex.

* Special categories of cited documents:	*T* later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
A document defining the general state of the art which is not considered to be part of particular relevance	*X* document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
E earlier document published on or after the international filing date	*Y* document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art
L document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)	*Z* document member of the same patent family
O document referring to an oral disclosure, use, exhibition or other means	
P document published prior to the international filing date but later than the priority date claimed	

Date of the actual completion of the international search

08 November 1993

Date of mailing of the international search report

13 DEC 1993

 Name and mailing address of the ISA/US
 Commissioner of Patents and Trademarks
 Box PCT
 Washington, D.C. 20231

Authorized officer

DAVID D. KNEPPER

B. Nando

Facsimile No. NOT APPLICABLE

Telephone No. (703) 305-9600

for

INTERNATIONAL SEARCH REPORT

International application No.
PCT/US93/08572

C (Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT		
Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	US, A, 3,947,827 (Dautremont, Jr. et al.) 30 March 1976, figs. 2, 4, 5.	1, 8
Y	US, A, 4,216,354 (Esteban et al.) 5 August 1980, fig. 4.	1, 8
Y	US, A, 4,622,680 (Zinser) 11 November 1986, fig. 2.	1, 8
Y	US, A, 4,489,437 (Fukuichi et al.) 18 December 1984, col. 3.	5
A	US, A, 4,468,804 (Kates et al.) 28 August 1984, abstract.	5-9
A	US, A, 4,435,832 (Asada et al.) 6 March 1984, col. 2.	5-9

INTERNATIONAL SEARCH REPORT

International application No.
PCT/US93/08572**Box I Observations where certain claims were found unsearchable (Continuation of item 1 of first sheet)**

This international report has not been established in respect of certain claims under Article 17(2)(a) for the following reasons:

1. ☐ Claims Nos.:
because they relate to subject matter not required to be searched by this Authority, namely:
2. ☐ Claims Nos.:
because they relate to parts of the international application that do not comply with the prescribed requirements to such an extent that no meaningful international search can be carried out, specifically:
3. ☐ Claims Nos.:
because they are dependent claims and are not drafted in accordance with the second and third sentences of Rule 6.4(a).

Box II Observations where unity of invention is lacking (Continuation of item 2 of first sheet)

This International Searching Authority found multiple inventions in this international application, as follows:

Telephone Practice

- I. Claims 1-4, drawn to an "audio compression apparatus" and
- II. Claims 5-9, drawn to an "audio decompression apparatus".

1. ☒ As all required additional search fees were timely paid by the applicant, this international search report covers all searchable claims.
2. ☐ As all searchable claims could be searched without effort justifying an additional fee, this Authority did not invite payment of any additional fee.
3. ☐ As only some of the required additional search fees were timely paid by the applicant, this international search report covers only those claims for which fees were paid, specifically claims Nos.:
4. ☐ No required additional search fees were timely paid by the applicant. Consequently, this international search report is restricted to the invention first mentioned in the claims; it is covered by claims Nos.:

Remark on Protest

- ☐ The additional search fees were accompanied by the applicant's protest.
- ☒ No protest accompanied the payment of additional search fees.