

UK Patent Application

(19) GB (11) 2 431 837 (13) A

(43) Date of A Publication 02.05.2007

(21) Application No:	0522051.2	(51) INT CL: G10L 19/00 (2006.01) G06F 17/15 (2006.01)
(22) Date of Filing:	28.10.2005	(52) UK CL (Edition X): H4R RPX
(71) Applicant(s): Sony United Kingdom Limited (Incorporated in the United Kingdom) The Heights, Brooklands, WEYBRIDGE, Surrey, KT13 0XW, United Kingdom		(56) Documents Cited: US 5940429 A US 5768426 A
(72) Inventor(s): William Edmund Kentish Nicolas John Haynes		(58) Field of Search: UK CL (Edition X) G5R, H4F, H4R INT CL G06T, G10L, G11B Other: On-line: WPI, EPODOC, INSPEC
(74) Agent and/or Address for Service: D Young & Co 120 Holborn, LONDON, EC1N 2DY, United Kingdom		

(54) Abstract Title: **Audio processing**

(57) A digital watermark, comprising a noise like signal is inserted into a primary audio signal, e.g. the audio track of a video presentation, to aid forensic examination to establish provenance of a suspect recording. The watermark signal may contain information as to origin and authorisation, but is designed to be almost imperceptible to the listener. The invention relates to embedding the watermark to preserve its level to ensure good detection, whilst remaining hidden. To achieve this, the level of the watermark signal is increased or decreased in response to the level of the primary audio signal. However if the watermark signal tracks the audio signal too closely, certain sounds can result in the watermark becoming audible. To avoid this, any decrease in the watermark signal is made to precede by an advance 360, the corresponding decrease in the audio signal level at t1.




Fig. 11




Fig. 1




Fig. 3




Fig. 2

3/14




Fig. 4




Fig. 2b




Fig. 5

5/14




Fig. 7




Fig. 6

6/14




Fig. 8




Fig. 9




Fig. 10




Fig. 11




Fig. 12

8/14




Fig. 13

9/14




Fig. 14




Fig. 15




Fig. 16

10/14




Fig. 17

11/14




Fig. 18




Fig. 19

13/14



Fig. 20




Fig. 21




Fig. 22

14/14




Fig. 23




Fig. 24




Fig. 25

AUDIO PROCESSING

This invention relates to audio processing.

In applications such as digital fingerprinting or watermarking (which may 5 collectively be referred to by the term forensic marking), a payload signal may be inserted into a primary audio signal in the form of a noise pattern such as a pseudo-random noise signal. The aim is generally that the noise signal is near to imperceptible and, if it can be heard, is not subjectively disturbing. This type of technique allows various types of payload to be added in a way which need not alter the overall bandwidth, bitrate and format of the 10 primary audio signal. The payload data can be recovered later by a correlation technique, which often still works even if the watermarked audio signal has been manipulated or damaged in various ways between watermark application and watermark recovery.

Examples of the type of payload data which can be added include security data (e.g. 15 for identifying pirate or illegal copies), broadcast monitoring data and metadata describing the audio signal represented by the primary audio signal.

The noise signal can be modulated before being added to the primary audio signal. This means in general terms that the level of the noise signal is increased when the level of the primary audio signal increases, and is decreased when the level of the primary audio signal decreases. In this way, more of the payload data's noise signal (giving a potentially 20 better recovery of the payload data) can be included when it can be masked by louder passages in the primary audio signal.

However, if the noise signal tracks the primary audio signal *too* closely it can become audible and potentially subjectively disturbing, especially with sounds such as drum beats and the like.

25 In envelope-controlled audio processing systems a time constant can be applied to the rise time and fall time of the controlled signal (in this example, the noise signal). These are known as the attack and decay (or release) time constants. If such measures are applied to the present example, the result is that a rapid rise in the primary audio signal level causes a slower rise in the noise signal. This is quite acceptable – even desirable in some 30 circumstances. But it is more of a problem that a sudden decrease in the primary audio signal level would lead to a slower decrease in the noise signal level. In an extreme case this could lead to undesirable situation of the noise signal being instantaneously larger than the primary audio signal.

This invention provides audio processing apparatus in which a payload signal is inserted into a primary audio signal, the apparatus comprising:

a noise generator operable to generate a noise signal in dependence on the payload signal;

5 a level detector for detecting a signal level of the primary signal;

a modulator for respectively increasing or decreasing the level of the noise signal in response to an increase or a decrease of the detected signal level of the primary audio signal, to generate a modulated noise signal;

a combiner for combining the primary signal and the modulated noise signal; and

10 a signal delay arrangement;

the modulator operating with respect to the signal delay arrangement so that a decrease in the level of the noise signal is time-advanced with respect to the corresponding decrease in the signal level of the primary audio signal.

The invention addresses the problem described above by providing a time-advanced release function, so that a decrease in the level of the noise signal is time-advanced with respect to the corresponding decrease in the signal level of the primary audio signal. In other words, with respect to the primary audio signal, the noise signal starts to fall before the primary audio signal starts to fall. The amount of this time advance can be set, relative to any release time constant in the system and the audio bandwidth of the primary audio signal,

20 so that either the noise signal is never larger than the primary audio signal, or so that any difference between them is within limits considered to be acceptable.

Further respective aspects and features of the invention are defined in the appended claims.

Embodiments of the invention will now be described, by way of example only, with reference to the accompanying drawings in which:

Figure 1 schematically illustrates a digital cinema arrangement including a fingerprint encoder;

Figure 2 schematically illustrates a fingerprint detector;

Figure 3 is a schematic overview of the operation of a fingerprint encoder;

Figure 4 schematically illustrates a payload generator;

Figure 5 schematically illustrates a fingerprint stream generator;

Figure 6 schematically illustrates a spectrum analyser;

Figure 7 schematically illustrates a spectrum follower;

Figures 8 to 11 schematically illustrate the operation of an envelope follower;

Figure 12 is a schematic overview of the operation of a fingerprint detector;

Figure 13 is a schematic flowchart showing a part of the operation of a temporal alignment unit;

5 Figure 14 schematically illustrates suspect material and proxy material divided into blocks;

Figure 15 schematically illustrates a low pass filter arrangement;

Figure 16 schematically illustrates a thresholded signal;

Figure 17 schematically illustrates a correlation operation;

Figure 18 schematically illustrates a power curve;

10 Figure 19 schematically illustrates a deconvolver training operation;

Figure 20 schematically illustrates a magnitude curve;

Figure 21 schematically illustrates a thresholded and interpolated magnitude curve;

Figure 22 schematically illustrates an intermediate result of the process shown in Figure 19;

15 Figure 23 schematically illustrates an impulse response;

Figure 24 schematically illustrates a smoothing curve;

Figure 25 schematically illustrates a smoothed impulse response; and

Figure 26 schematically illustrates a data processing apparatus.

20 Introduction

Fingerprinting or watermarking techniques – more generically referred to as forensic marking techniques - have been proposed which are suitable for video signals. See for example EP-A-1 324 262. While the general mathematical framework may appear in principle to be applicable to audio signals, several significant technical differences are present. In the present description, both “fingerprint” and “watermark” will be used to denote a forensic marking of material.

One of the main factors to be considered is how the fingerprint data should be encoded into the audio signal. The human ear is very different from the human eye in terms of sensitivity and dynamic range, and this has made many previous commercial fingerprinting schemes fail in subjective listening (“A/B”) tests.

30 The human ear is capable of hearing phase differences of less than one sample at a 48kHz sampling rate, and it has a working dynamic range of 9 orders of magnitude at any one time. With this in mind, an appropriate encoding method is considered to be encoding the fingerprint data as a low-level noise signal that is simply added to the media.

Noise has many psycho-acoustic properties that make it favourable to this task, not least of which is that the ear tends to ignore it when it is at low levels, and it is a sound that is generally calming (in imitation of the natural sounds of wind, rushing streams or ocean waves), rather than generally irritating. The random nature of noise streams also implies 5 there is little possibility of interfering with brain function in the way that, for example, strobe effects or malicious use of subliminal information can do to visual perception.

An implementation of this type of technique will now be described.

Mathematical Foundation

10 Consider a fingerprint payload “vector” $\mathbf{P} = p[1]..p[n]$.

For the *embedding* process, this payload is added to an audio signal vector $\mathbf{V} = v[1]..v[n]$ to yield a watermarked payload vector $\mathbf{W} = \mathbf{V} + \mathbf{P}$.

15 The elements of the payload vector \mathbf{P} are statistically independent random variables of mean value 0, and standard deviation α^2 , where α is referred to as the strength of the watermark, written as $N(0, \alpha^2)$. Simply stated, this notation is used to indicate that the payload is a Gaussian random noise stream. The noise stream is scaled so that the standard deviation is in the range $+/- 1.0$ as an audio signal. This scaling is important because if this is not done correctly, the similarity indicator (“SimVal”) calculated below will not be correct. Note that the convention here is that $+/- 1.0$ is considered to be “full scale” in the audio 20 domain, and so in the present case many samples of the Gaussian noise stream will actually be greater than full scale.

For the *extraction* process, the original proxy vector \mathbf{V} is subtracted from a watermarked suspect vector (e.g. a pirate copy of the audio material in question) \mathbf{W}_s to yield the suspect payload vector $\mathbf{P}_s = \mathbf{W}_s - \mathbf{V}$. In other words, $\mathbf{P}_s = \text{Suspect-audio-stream} - \text{Proxy-audio-stream}$.

25 To test whether the content was watermarked with a candidate payload vector \mathbf{P} , an inner-loop correlation (written as “ \bullet ”) is performed between the candidate payload vector \mathbf{P} and the *normalised* suspect payload vector \mathbf{P}_s to yield a *similarity value*, hereafter termed a SimVal:

30

$$\text{SimVal} = (\mathbf{P}_s / |\mathbf{P}_s|) \bullet \mathbf{P}$$

where $|\mathbf{Ps}|$ is the vector magnitude of \mathbf{Ps} , meaning $|\mathbf{Ps}| = \sqrt{\mathbf{Ps} \cdot \mathbf{Ps}}$. Here, $\sqrt{\cdot}$ indicates a square root function. Note that to *normalise* a vector means to scale the values within the vector so they add up to a magnitude of exactly 1.

5 This formula indicates the degree of statistical correlation between \mathbf{Ps} and \mathbf{P} , with a maximum value that is close to the square root of the length of the vector. We say that if the SimVal is greater than a particular threshold value T , then the payload \mathbf{P} is present in \mathbf{Ps} , and if the SimVal $\leq T$, then it is *not* present.

In order to give the values of SimVal some statistical meaning, the value of T is related to the probability of a false positive by the following formula:

10

$$T = \sqrt{2 \ln(M^2 / p \sqrt{2\pi})}$$

15 where p is the false positive probability, \ln is the natural logarithm, and M is the population size (i.e. the number of unique payload vectors issued for the given audio content). For example, if the false probability is required to be better than 1 in 100,000,000, and the population size is 1000, the value SimVal will need to be greater than 8.

20 Generally speaking, a SimVal of 10 is a useful aim in forensic analysis of pirate audio material using the present techniques. For particularly large populations M , a value of 12 might be more appropriate. In empirical trials, it has been found that if a value of 8 is reached within analysis of a few seconds of the suspect audio material, a value of 12 will generally be reached within another few seconds.

25 Figure 1 schematically illustrates a digital cinema arrangement in which a secure playout apparatus 10 receives encrypted audio/video material along with a decryption key. A decrypter 20 decrypts the audio and video material. The decrypted video material is supplied to a projector 30 for projection onto a screen 40. The decrypted audio material is provided to a fingerprint encoder 50 which applies a fingerprint as described above.

Generally, the fingerprint might be unique to that material, that cinema and that instance of replay. This would allow piracy to be retraced to a particular showing of a film.

30 The fingerprinted audio signal is passed to an amplifier 60 which drives multiple loudspeakers 70 and sub-woofer(s) 80 in a known cinema sound configuration.

Fingerprinting may also be applied to the video information. Known video fingerprinting means (not shown) may be used.

Preferably, the playout apparatus is secure, in that it is a sealed unit with no external connections by which non-fingerprinted audio (or indeed, video) can be obtained. Of course, the amplifier 60 and projector 30 need not necessarily form part of the secure system.

If an illegal copy is made of the material from that cinema performance, for example 5 by the use of a camcorder within the cinema, the audio content associated with the film will have the fingerprint information encoded by the fingerprint encoder 50 included within it. In order to establish this, for investigative or legal reasons, a suspect copy of the material can be supplied to a fingerprint detector 80 of Figure 2 along with the original (or “proxy”) material and a key used to generate the original fingerprint. In its simplest terms, the fingerprint 10 detector 80 generates a probability that the particular fingerprint is present in the suspect material. The detection process will be described in more detail below.

Embedding Process

In video fingerprinting the techniques are generally frame based (a frame being a 15 natural processing block size in the video domain), and the whole of the fingerprint payload vector is buried (at low level) in each frame. In some systems the strength of the fingerprint is set to be greater in “busier” image areas of the frame, and also at lower spatial frequencies which are difficult or impossible to remove without seriously changing the nature of the video content. The idea is that over many frames the correlations on each frame can be 20 accumulated, as if the correlation were being done on a single vector; if there is a real statistical correlation between the suspect payload \mathbf{Ps} and the candidate payload \mathbf{P} , the correlation will continue to rise from frame to frame.

For audio, there is generally no such natural processing block.

In the present embodiments, for reasons of efficiency of fast Fourier transform (FFT) 25 operations, a processing block size of the audio version is set to a power of 2 audio samples, for example 64k samples (65536 samples). Note also that the vector lengths will be the same size as the processing block.

Successive correlations for these audio frames can be accumulated in the same way as for the video system.

30 There is one sample of payload vector for each sample of content. Also, the payload is concentrated in the “mid-frequencies” because both the high frequency content (say $> 5\text{KHz}$) and the low frequency content (say $< 150\text{Hz}$) can be completely lost without intolerable loss of audio quality. The loss of these frequencies could be an artefact of poor recording equipment or techniques on the part of a pirate, or they could be deliberately

removed by a pirate to try to inhibit a fingerprint recovery process. It is therefore more appropriate to concentrate the payload into the more subjectively important mid frequencies, i.e. frequencies that cannot be easily removed without seriously degrading the quality.

In general terms:

- 5 1. The payload seeds an AES Rijndael-based pseudo-random number stream to generate a noise stream.
2. The noise stream is “shaped” according to a perceptual analysis of the audio stream.
3. The shaped noise stream is added at low level to the audio stream.

10 The generated noise stream contains multiple layers within it, each generated from a different subset of the payload data. It will be appreciated that other data could be included within the payload, such as a frame number and/or the date/time.

15 The random number streams are generated by repeated application of 256-bit Rijndael encryption to a moving counter. The numbers are then scaled to be within +/-1.0, to produce full scale *white* noise. The white noise stream is turned into Gaussian noise by applying the Box-Muller transform to pairs of points.

In the present embodiment there are 16 layers to the noise stream. A first layer of the pseudo-random noise generator is seeded by the first 16bits of the payload, the second layer seeded by the first 32 bits of the payload, and so on until the 16th layer which is seeded by the entire 256 bit payload.

20 Perceptual analysis involves a simple spectral analysis in order to establish a gain value to scale the Fingerprint noise stream for each sample in the audio stream. The idea is that louder sections in the audio stream will hide louder intensity of fingerprint noise.

25 Extending this concept further, the mid-frequency content of the audio stream (where the fingerprint is to be hidden) is split into several bands (say 8 or 12) which are preferably spread evenly on a logarithmic frequency scale (though of course any band-division could be used). This means, for example, that the frequency spectrum is roughly divided into the octaves. Each band is then processed separately to generate a respective gain envelope that is used to modulate the amplitude of the corresponding frequency band in the fingerprint noise stream. When the envelope modulation is used in all bands, the result is that the noise stream sounds very much like a “ghostly” rendition of the original audio signal. More importantly, this ghostly rendition, because of its similarity to the content, when added to the original material, becomes inaudible to the ear, despite being added at relatively high signal levels. For example, even if the modulated noise is added at a level as high as -30dB (decibels) relative to the audio, it can subjectively be almost inaudible.

The present embodiment uses 2049 sample impulse response kernels to implement “brick wall” (steep-sided response) convolution band filters to separate the information in each frequency band. The convolutions are done in the FFT domain for speed. One important reason for using convolution filters for the band pass filter rather than recursive filters is that the convolution filters can be made to have a fixed delay that is independent of frequency. The reason this is important is that the modulations of the noise-stream for any given frequency band must be made to line up with the actual envelope of the original content when the noise stream is added. If the filters were to have a delay that depends on frequency, the resultant misalignment would be difficult to correct, which could lead to increased perceptibility of the noise and possible variation of correlation values with frequency.

Figure 3 is a schematic overview of the operation of a fingerprint encoder such as the encoder 50 of Figure 1. A payload generator 100 produces payload data to be encoded as a fingerprint. As mentioned above, this could include various content and other identifiers and may well be unique to that instance of the replay of the content. The payload generator will be described further below with reference to Figure 4.

The payload is supplied to a fingerprint stream generator 110. As described above, this is fundamentally a random number generator using AES-Rijndael encryption based on an encryption key to produce an output sequence which depends on the payload supplied from the payload generator 100. The fingerprint stream generator will be described further below with reference to Figure 5.

The source material (to which the fingerprint is to be applied) is supplied to a spectrum analyser 120. This analyses the amplitude or envelope of the source material in one or more frequency bands. The spectrum analyser supplies envelope information to a spectrum follower 130. The spectrum follower modulates the noise signal output by the fingerprint stream generator 110 in accordance with the envelope information from the spectrum analyser 120. The spectrum analyser will be described further below with reference to Figure 6 and the spectrum follower with reference to Figure 7.

The output of the spectrum follower 130 is a noise signal at a significantly lower level than the source material but which generally follows the envelope of the source material. The noise signal is added to the source material by an adder 140. The output of the adder 140 is therefore a fingerprinted audio signal.

A delay element 150 is shown schematically in the source material path. This is to indicate that the spectrum analysis and envelope determination may take place on a time-

advanced version of the source material compared to that version which is passed to the adder 140. This time-advance feature will be described further below.

Figure 4 schematically illustrates a payload generator. As mentioned above, this takes various identification data such as a serial number, a location identifier and a location 5 private key and generates payload data 160 which is supplied as a seed to the fingerprint stream generator 110. The location private key may be used to encrypt the location identifier by an encryption device 170. The various components of the payload data are bit-aligned for output as the seed by logic 180.

Figure 5 schematically illustrates a fingerprint stream generator 110. This receives 10 the seed data 160 from the payload generator 100 and key data 190 which is expanded by expansion logic 200 into sixteen different keys K-1...K-16.

A frame number may optionally be added to the seed data 160 by an adder 210.

The stream generator has sixteen AES-Rijndael number generators 220..236. Each of 15 these receives a respective key from the key expansion logic 200. Each is also seeded by a respective set of bits from the seed data 160. The number generator 220 is seeded by the first 16 bits of the seed data 160. The number generator 221 is seeded by the first 32 bits of the seed data 160 and so on. This arrangement allows a hierarchy of payloads to be established which can make it easier to search for a particular fingerprint at the decoding stage by first searching for all possible values of the first 16 bits, then searching for possible 20 values of the 17th to 32nd bits (knowing the first 16 bits) and so on.

The output of each number generator 220..236 is provided to a Gaussian mapping arrangement 240..256. This takes the output of the number generator, which is effectively white noise, and applies a known mapping process to produce noise with a Gaussian profile.

The Gaussian noise signals from each instance of the mapping logic 240..256 are 25 added by an adder 260 to generate a noise signal 270 as an output.

Figure 6 schematically illustrates a spectrum analyser 120. This receives the source material (to be fingerprinted) as an input and generates envelope information 280 as outputs.

The spectrum analyser comprises a set of eight (in this example) band filters 290..297, each of which filters a respective band of frequencies from the source material.

30 The filters may be overlapping or non-overlapping in frequency, and the extent of the entire available frequency range which is covered by the eight filters may be one hundred percent or, more usually, much less than this. The respective bands relating to the eight filters may be contiguous (i.e. adjacent to one another) or not. The number of filters (bands) used could

be less than or more than eight. It will accordingly be realised that the present description is merely one example of the way in which these filters could operate.

5 In the present case, a mid-frequency range is handled by the filters, from about 150Hz to about 5kHz. This is divided into eight logarithmically equal bands, each of which therefore extends over about one octave. The filtering technique used for the band filters 290..297 is in accordance with that described above.

At the output of each band filter, is an envelope detector 300..307. This generates an envelope signal relating to the envelope of the filtered source material at the output of the respective band filter.

10 Figure 7 schematically illustrates a spectrum follower. The spectrum follower receives the envelope information 280 from the spectrum analyser 120 and the Gaussian noise signal 270 from the fingerprint stream generator 110.

15 The Gaussian noise signal 270 is supplied to a set of band filters 310..317. These are set up to have the same (or as near as practical) responses as the corresponding filters 290..297 of the spectrum analyser 120. This generates eight bands within the noise spectrum. Each of the filtered noise bands is supplied to a respective envelope follower 320..327. This takes the envelope signal relating to the envelope of that band in the source material and modulates the filtered noise signal in the same band. The outputs of all of the envelope followers 320..327 are summed by an adder 330 to generate a shaped noise signal 20 340.

The envelope followers can include a scaling arrangement so that the eventual shaped noise signal 340 is at an appropriate level with respect to the source material, for example minus 30dB with respect to the source material.

25 As mentioned above, the shaped noise signal 340 is added to the source material by the adder 140 to generate fingerprinted source material as an output signal.

The fingerprinting process can take place on different audio channels (such as left and right channels) separately or in synchronism. It is however preferred that a different noise signal is used for each channel to avoid a pirate attempting to derive (and then remove or defeat) the fingerprint by comparing multiple channels. In either case, the envelope 30 signals 280 preferably relates to the individual audio channel being fingerprint encoded.

The operation of the envelope detection and envelope following described above will now be explained in more detail with reference to Figures 8 to 11. Note that in the case of the spectrum follower described above, envelope following would take place in respect of each channel or band. Also, the time constants to be described below can be made

dependent on the audio frequency or frequency range applicable to a band, e.g. dependent on the fastest rise time of a signal within that band. This would allow them to be adjusted as a group, by simply changing the relationship between time constant and fastest rise time.

5 In Figures 8 to 11, the horizontal axis represents time on an arbitrary scale, the solid curve represents an example (in schematic form) of an envelope signal relating to the source material and the broken lines represent (in schematic form) the modulation applied by the envelope followers 320..327.

In Figure 8, a time constant is applied by the envelope follower to restrict the rise time of the noise signal in response to a sudden rise of the envelope of the source material. 10 This is represented by a left hand section of the broken line, lagging in time behind the more vertical rise of the solid line. Such a time constant is often referred to as an “attack” time constant. However, it will be noted in all of Figures 8 to 11 that although the *rate* of rise of the noise signal is limited, the time at which the noise signal *starts* to rise is the same as the time at which the envelope signal starts to rise (subject only to trivial time differences caused 15 by detection delays). It would be possible to delay (or even, with the time-advanced arrangements described below, advance) the start of the noise signal’s rise with respect to the envelope signal, but this appears to give little benefit. In particular, delaying the rise of the noise signal restricts the useful payload which can be concealed behind a rising signal, and advancing the noise signal’s start time could give audible artefacts similar to those to be 20 described with reference to the trailing edge of the envelope of Figure 8.

Similarly, at the trailing edge of the source material envelope, the decrease of the noise envelope shown by the trailing dotted line is also restricted by a “decay” time constant. Unfortunately, this means that over a period from t_1 to t_2 the noise signal is larger than the source material signal and so the noise could be subjectively disturbing to the listener.

25 Figure 9 illustrates the situation common in envelope following audio effects processors, whereby a “sustain” period 350 is defined which delays the onset of the decay of the envelope-following signal (in this case, the noise signal). This makes the situation described above even worse, in that the noise signal is now larger than the source material signal between times t_1 and t_3 . Accordingly, a sustain period is not used in the present 30 embodiments.

Measures to address this problem will be described with reference to Figures 10 and 11.

In Figure 10, the time at which the noise signal starts to decrease is advanced with respect to the time at which the source material’s envelope decreases by an advanced time

360. In this example, this means that the noise signal has decayed to insignificant levels by the time t_1 .

In Figure 11, if the advance period 360 is reduced slightly, then the noise signal starts to decrease before the source material's envelope decreases, but it has not finished 5 decreasing by the time t_1 . This means that between the times t_1 and t_4 there is a small amount of noise still present, but the problem is much less than that shown in Figure 8.

Accordingly, by starting the decrease of the noise signal at an earlier time than the decrease of the source material envelope which prompts that noise reduction, the subjectively disturbing excess noise shown in Figures 8 and 9 can be reduced or avoided.

10 In order to achieve this, it is necessary to include a delay somewhere within the system so that envelope information for the source material can be acquired in a time-advanced relationship to the addition of the source material to the noise at the adder 140. The delay shown in Figure 3 is a very schematic example of how this might be achieved. The skilled person will appreciate that many other possibilities are available.

Extraction Process

The major stages of fingerprint extraction are as follows:

1. The suspect material is treated to attempt to reverse any damage or distortion.
2. So-called proxy content (a term used to describe an unwatermarked original version 20 of the content) is subtracted from the suspect content to leave the suspect fingerprint. This relies on being able to align temporally the suspect material and the proxy content. In some circumstances a watermarked proxy may be used. Of course the watermark in the proxy is likely to be detected by correlation, but it does not prevent other watermark(s) being detected, and can be ignored. In this way secured copies 25 may be sent to third parties contracted to operate the extraction process.
3. The suspect fingerprint is “unshaped” according to a spectral analysis of the proxy content.
4. For each candidate payload in the population for this content, compare candidate payload to the suspect payload over a relatively short section of content. If the value 30 SimVal looks promising, add this candidate to the short-list of candidates that will be subjected to a much longer analysis.

Figure 12 is a schematic overview of the operation of a fingerprint detector such as the detector 80 of Figure 2. The detector receives suspect material, such as a suspected

pirate copy of a piece of content, and so-called proxy material which is a plain (non-watermarked) copy of the same material.

The suspect material is first supplied to a temporal alignment unit 400. The operation of this will be described below with reference to Figures 13 to 18. In brief, however, the 5 temporal alignment unit detects any temporal offset between the proxy material and the suspect material and so allows the two sets of material to be aligned temporarily. The alignment which can potentially be achieved by the temporal alignment 400 is to within a certain tolerance such as a tolerance of \pm one sample. Further time corrections to allow a complete alignment between the two signals are carried out by a deconvolver 410 to be 10 described below.

The deconvolver applies an impulse response to the suspect material to attempt to render it more like the proxy material. The aim here is to reverse (at least partially) the effects of signal degradations in the suspect material; examples of such degradations are listed below.

15 In order to do this, the deconvolver 410 is “trained” by a deconvolver training unit 420. The operation of the deconvolver training unit will be described below with reference to Figures 19 to 25, but in brief, the deconvolver training unit compares the time-aligned suspect material and proxy material in order to derive a transform response which represents what might have happened to the proxy material to turn it into the suspect material. This 20 transform response is applied “in reverse” by the deconvolver 410. Preferably, the transform response is updated at different positions within the suspect material so as to represent the degradation present at that particular point. In the embodiment to be described below, the transform response detected by the deconvolver training unit is based upon a rolling average of responses detected over a predetermined member of most-recent portions for blocks of the 25 suspect material and proxy material.

A delay 430 may be provided to compensate for the deconvolver and deconvolver training operation.

A cross normalisation unit 440 then acts to normalise the magnitudes of the deconvolved suspect material and the proxy material. This is shown in Figure 12 as acting 30 on the suspect material but it will be appreciated that the magnitude of the proxy material could be adjusted, or alternatively, the magnitudes of both could be adjusted.

After normalisation, a subtractor 450 establishes the difference between the normalised, deconvolved suspect material and the proxy material. This difference signal is passed to an “unshaper” 460 which is arranged to reverse the effects of the noise shaping

carried out by the spectrum follower 130. In order to do this, the proxy material is subjected to a spectrum analysis stage 470 which operates in an identical way to the spectrum analyser 120 of Figure 3.

So, the spectrum analyser 470 and the unshaper 460 can be considered to operate in an identical manner to the spectrum analyser 120 and the spectrum follower 130, except that a reciprocal of the envelope-controlled gain value is used with the aim of producing a generally uniform noise envelope as the output of the unshaper 460. The noise signal generated by the unshaper 460, P_s is passed to a comparator 480. The other input to the comparator, P , is generated as follows.

A fingerprint generator 490 operates in the same way as the payload generator 100 and fingerprint stream generator 110 of Figure 3. Accordingly, these operations will not be described in detail here. The fingerprint generator 490 operates, in turn, to produce all possible variants of the fingerprint which might be present in the suspect material. Each is tested in turn to derive a respective likelihood value SimVal.

Of course it would be possible to employ multiple fingerprint generators 490 and to use multiply comparators 480 acting in parallel so that the noise stream P_s is compared with more than one fingerprint at a time.

Delays 500, 510 are provided to compensate for the processing delays applied to the suspect material, in order that the fingerprint generated by the fingerprint generator 490 is properly time-aligned with the fingerprint which may be contained within the suspect material.

Temporal Alignment

The first thing to do with the suspect pirated signal is to find the true synchronisation with the proxy signal.

A sub-sample delay may be included to allow, if necessary, to compensate for any sub-sample delay/advance imposed by re-sampling or MP3 encoding effects.

Figure 13 is a schematic flowchart showing a part of the operation of the temporal alignment unit 400. Each step of the flowchart is implemented by a respective part or function of the temporal alignment unit 400.

While it would be possible, in theory, to align the suspect and proxy material by a (single) direct correlation process, in the case of substantial material such as a film soundtrack, the correlation processing required would be enormous, as the processing operations increase generally with the square of the number of audio samples involved.

Accordingly, the present process aimed to provide at least an approximate alignment without the need for a full correlation of the two signals.

Referring to Figure 13, at a step 600, the two audio signals are divided into portions or blocks. These blocks are of equal size for each of the two signals, but need not be a predetermined size. So, one option would be to have a fixed size of (say) 64k samples, whereas another option is to have a fixed number of blocks so that the total length of the longer of the two pieces of material (generally the proxy material) is divided by a predetermined number of blocks to arrive at a required block size for this particular instance of the time alignment processing. In any event, the block size should be at least two samples.

A low pass pre-filtering stage (not shown) can be included before the step 600 of Figure 13. This can reduce any artefacts caused by the arbitrary misalignment between the two signals with respect to the block size.

At a step 605, the absolute value of each signal is established and the maximum power detected (with reference to the absolute value) for each block. Of course, different power characteristics could be established instead, such as mean power. The aim is to end up with a power characteristic signal from each of the proxy and suspect signals, having a small number (e.g. 1 or 2) of values per block. The present example has one value per block.

At a step 610, the two power characteristic signals are low-pass filtered or smoothed.

Figure 14 schematically illustrates the division of the two signals into blocks, whereby in this example the proxy material represents the full length of a movie film and the suspect material represents a section taken from that movie film.

Figure 15 schematically illustrates a low pass filter applied to the two power characteristic signals separately. Each sample is multiplied (at a multiplier 611 by a coefficient, and added at a adder 612 to the product of the adder's output and a second coefficient. This takes place at a multiplier 613. This process produces a low-pass filtered version of each signal.

At this stage, the two power characteristic signals have a magnitude generally between zero and one. The filtering process may have introduced some minor excursions above one, but there are no excursions below zero because of the absolute value detection in the step 605.

At a step 630, a threshold is applied. This is schematically illustrated in Figure 16. An example of such a threshold might be 0.3, although of course various other values can be used.

The threshold is applied as follows.

The aim is to map the power characteristic signal value corresponding to the threshold to a revised value of one. Any signal values falling below the threshold will be mapped to signal values between zero and one. Any signal values falling above the 5 threshold will be mapped to signal values greater than one. So, one straightforward way of achieving this is to multiply the entire power characteristic signal by a value of 1/threshold, which in this case would be 3.33....

The reason why this is relevant is that the next step 640 is to apply a power law to the signals. An example here is that each signal is squared, which is to say that each sample 10 value is multiplied by itself. However, other powers greater than 1, integral or non-integral, could be used. The overall effect of the step 630 and 640 is to emphasise higher signal values and diminish the effect of lower signal values. This arises because any number between zero and one which is raised to a power greater than one (e.g. squared) gets smaller, whereas any signal value greater than one which is raised to a power greater than one 15 becomes larger.

After application of the power law, the resulting signals are subjected to an optional high-pass filtering process at a step 650. At a step 660, the mean value of each signal is subtracted so as to generate signals having a mean of zero. (This step is useful for better operation of the following correlation step 670).

20 Finally, at a step 670, the power characteristic signals are subjected to a correlation process. This is illustrated schematically in Figure 17, where the power values from the suspect material are padded with zeros to provide a data set of the same length as the proxy material. The correlation process will (hopefully) generate a peak correlation, whose offset 701 from a centre position 702 indicates a temporal offset between the two files. This offset 25 can be corrected by applying a relative delay to either the proxy or the suspect signals.

The process described with reference to Figure 13 to 17 can be repeated with a smaller block size and a restricted range about which correlation is performed (taking the offset 701 from the first stage as a starting position and an approximate answer). Indeed, the process can be repeated more than twice at appropriately decreasing block sizes. To gain a 30 benefit, the block size should remain at least two samples.

Figure 18 schematically illustrates a power characteristic signal as generated by the step 605, and a filtered power characteristic signal as generated by the step 660. Here, the threshold is 0.3, the power factor in step 640 is 1.5 and a 1/10 scaling has been applied.

Damage Reversal

5 The purpose of damage reversal is to transform the pirated content in such a way that it becomes as close as possible to the original proxy version. This way the suspect payload Ps that results from subtracting the proxy from the pirated version will be as small as possible, which should normally result in larger values of SimVal.

For audio, there is a long list of possible distortions that can be accidentally or purposefully imposed by the pirate, each potentially resulting in a reduction in the SimVal value:

- 10 ▪ High, Low, Notch, Band or Parametric Filtering
- Compression, Expansion, Limiting, Gating
- Overdrive, clipping.
- Inflation, valve-sound, and other sound enhancement effects
- Re-sampling, ADC and DAC re-conversion
- Freq drift, wow-and-flutter, Phase reversal, vari-speed.
- 15 ▪ MP3-family lossy encoding/decoding techniques.
- Echo, Reverb, Spatialisation.
- So-called de-essing, de-hissing, de-crackling.

20 To counter as many of these damages as possible, the fingerprint recovery arrangement includes a general purpose deconvolver, which with reference to the Proxy signal can be trained to significantly reduce/remove any effect that could be produced by the action of a convolution filter. Other previous uses of deconvolvers can be found in telecommunications (to remove the unwanted echoes imposed by a signal taking a number of different paths through a system) and in archived material restoration projects (to remove age damage, or to remove the artefacts of imperfect recording equipment).

25 Briefly, the deconvolver is trained by transforming the suspect pirated audio material and the proxy version into the FFT domain. The Real/Imaginary values of the desired signal (the proxy) are divided (using complex division) by the Real/Imaginary values of the actual signal (the pirated version), to gain the FFT of an impulse response kernel that will transform the actual response to the desired response. The resulting FFT is smoothed and then 30 averaged with previous instances to derive an FFT that represents a general transform for that audio signal in the recent past. The FFT is then turned into a time domain impulse response kernel ready for application as a convolution filter (a process that involves rotating

the time domain signal and applying a window-sync function to it such as a “Hamming” window to reduce aliasing effects).

5 A well trained deconvolver can in principle reduce by a factor of ten the effect of non-linear gain effects applied to a pirated version, for example by microphone compression circuitry. In an empirical test, it was found that the deconvolver was capable of increasing a per-block value of SimVal from 15 to 40.

Figure 19 schematically illustrates a deconvolver training operation, as applied by the deconvolver training unit 420.

10 The process starts with a block-by-block fast Fourier transform (FFT) of both the suspect material (700) and the proxy material (710), where the block size might be, for example, 64k consecutive samples. A divider 720 divides one of the FFTs by the other. In the present case, because it is desired to generate a transform response which will be applied to the suspect material, the divider operates to divide the proxy FFT by the suspect FFT.

15 An averager 730 averages a current division from the divider 720 and n most recent division results stored in a buffer 740. Of course, the most recent result is also added to the buffer and a least-recently stored result discarded. An example of n is 5. It would of course be possible to store the raw FFTs, form two averages (one for the proxy and one for the suspect material) and divide the averages, but this would increase the storage requirement.

20 A converter then converts the averaged division result, which is a complex result, into a magnitude and phase representation.

Logic 750 removes any small magnitude values. Here, while the magnitude value is deleted, the corresponding phase value is left untouched. The logic 750 operates only on magnitude values. The deleted small magnitude values are replaced by values interpolated from the nearest surrounding non-deleted magnitude values, by a linear interpolation.

25 This process is illustrated schematically in Figures 20 and 21, where Figure 20 schematically illustrates the output of the magnitude/phase converter 740 as a set of magnitude values (the phase values are not shown). Any magnitude values falling below a threshold T_{mag} are deleted and replacement values 751, 752, 753 generated by linear interpolation between the nearest non-deleted values.

30 The resulting magnitude values are smoothed by a low-pass filter 760 before being converted back to a complex representation at a converter 770. An inverse FFT 780 is then applied. This generates an impulse response rather like that shown in Figure 22. In order to arrive at a suitable form for a deconvolution with the suspect material, the impulse response

is rotated by half of the window size so as to adjoin the two half-lobes into a central peak such as that shown in Figure 23. This is carried out by logic 790.

However, the output from the logic 790, shown in Figure 23, is still not entirely suitable for the deconvolution. This is because the side lobes 791 of this response extend 5 across the entire window. This can cause aliasing problems if such a response was used in the deconvolver 410. Therefore, a modulator 800 multiplies the response of Figure 23 by a sync window function such as that shown in Figure 24, to produce a required impulse response such as that shown in Figure 25. It is this impulse response which is supplied to the deconvolver 410.

10 Level Matching

After the deconvolving operation, the pirated signal is made to match the level of the proxy signal as closely as possible. In practice, empirical tests showed that a useful way to do this is to match the mean magnitudes of the two signals, rather than matching the peak values.

15 Once these three steps (Time alignment, Deconvolution and Level Matching) has been achieved, the proxy signal is subtracted from the pirated material to leave the suspect payload **Ps**.

Suspect Payload Extraction

Note that the payload signal that comes out of the Noise Shaper in the embedding 20 process is very different from the Gaussian noise stream that went into it. In order to recover a suspect payload signal that more closely matches the candidate payload Gaussian noise stream (in the statistical sense) for purposes of finding the value SimVal, it is appropriate to reverse the effect of noise-shaping – i.e. to “unshape” the payload signal.

The “unshaping” is achieved by using the same noise-shaping component, except that 25 instead of multiplying the gain values with the noise stream, a division is applied.

Another possible method, that of noise-shaping the candidate payload stream prior to comparison, is possible from a technical point of view but is not favoured for legal reasons. This is because it would be in violation of the mathematical principle adopted in digital rights management systems that the candidate stream be composed of statistically 30 independent samples. The application of filters to a noise stream automatically relates the samples.

Another reason is that the technique of convolution tends to operate more successfully if the signal being sought is buried in noise. Looking for a noise stream amongst noise is generally more effective and reliable (since it yields a much more *stable* cross-correlation) than looking for a shaped signal amongst similarly shaped residual audio signals.

Finally, Figure 26 illustrates a data processing apparatus. This is provided merely as one example of how the encoder 50 of Figure 1 or the detector 80 of Figure 2 may be implemented. However, it should be noted that at least in Figure 1, the entire digital cinema arrangement 10 is preferably a secure unit with no external connections, so it may be that the 10 fingerprint encoder, at least, is better implemented as a hard-wired device such as one or more field programmable gate arrays (FPGA) or application specific integrated circuits (ASIC).

Referring to Figure 26, the data processing apparatus comprises a central processing unit 900 memory 910 (such as random access memory, read only memory, non-volatile memory or the like), a user interface controller 920 providing an interface to, for example, a display 930 and a user input device 945 such as a keyboard, a mouse or both, storage 930 such as hard disk storage, optical disk storage or both, a network interface 940 for connecting to a local area network or the internet 950 and a signal interface 960. In Figure 26, the signal interface is shown in a manner appropriate to the fingerprint encoder 50, in that 20 it receives unfingerprinted material and output fingerprinted material. However, the apparatus could of course be used to embody the fingerprint detector.

The elements 900, 910, 940, 920, 930, 960 are interconnected by a bus 970. In operation, a computer program is provided by a storage medium (e.g. an optical disk) or 25 over the network or Internet connection 950 and is stored in memory 910. Successive instructions are executed by the CPU 900 to carry out the function described in relation to fingerprint encoding or detecting as described above.

CLAIMS

1. Audio processing apparatus in which a payload signal is inserted into a primary audio signal, the apparatus comprising:

5 a noise generator operable to generate a noise signal in dependence on the payload signal;

a level detector for detecting a signal level of the primary signal;

a modulator for respectively increasing or decreasing the level of the noise signal in response to an increase or a decrease of the detected signal level of the primary audio signal,
10 to generate a modulated noise signal;

a combiner for combining the primary signal and the modulated noise signal; and

a signal delay arrangement;

the modulator operating with respect to the signal delay arrangement so that a decrease in the level of the noise signal is time-advanced with respect to the corresponding
15 decrease in the signal level of the primary audio signal.

2. Apparatus according to claim 1, in which the modulator operates with respect to the signal delay arrangement so that an increase in the level of the noise signal is not time-advanced with respect to the corresponding increase in the signal level of the primary audio
20 signal.

3. Apparatus according to claim 1 or claim 2, in which the payload signal is a forensic marking signal.

25 4. Apparatus according to any one of the preceding claims, in which the noise generator is a pseudo-random noise generator seeded by the at least some of the payload signal.

5. Apparatus according to claim 4, in which the noise generator is operable to encrypt at least some of the payload data in accordance with a security key.

30

6. Apparatus according to any one of the preceding claims, in which the combiner is operable to add the noise signal to the primary audio signal.

7. Apparatus according to any one of the preceding claims, comprising two or more sets of the level detector, the modulator and the combiner, arranged to operate in respect of two or more respective audio frequency bands, the apparatus comprising means for combining the outputs of the two or more combiners.

5

8. Audio processing apparatus substantially as hereinbefore described with reference to the accompanying drawings.

9. An audio processing method in which a payload signal is inserted into a primary audio signal, the method comprising the steps of:

generating a noise signal in dependence on the payload signal;
detecting a signal level of the primary signal;
respectively increasing or decreasing the level of the noise signal in response to an increase or a decrease of the detected signal level of the primary audio signal, to generate a modulated noise signal; and

combining the primary signal and the modulated noise signal;
the arrangement being such that a decrease in the level of the noise signal is time-advanced with respect to the corresponding decrease in the signal level of the primary audio signal.

20

10. An audio processing method substantially as hereinbefore described with reference to the accompanying drawings.

11. Computer software having program code which, when run on a computer, causes the computer to carry out a method according to claim 9 or claim 10.

12. A medium by which software according to claim 11 is provided.

13. A medium according to claim 12, the medium being a storage medium.

30

14. A medium according to claim 12, the medium being a transmission medium.



For Innovation

23

Application No: GB0522051.2

Examiner: Mr John Betts

Claims searched: All

Date of search: 15 May 2006

Patents Act 1977: Search Report under Section 17

Documents considered to be relevant:

Category	Relevant to claims	Identity of document and passage or figure of particular relevance
A	-	US5768426 A (Rhoads) see Fig. 6 and col. 23 lines 35-41
A	-	US5940429 A (Solana) see col. 5 lines 38-53

Categories:

X	Document indicating lack of novelty or inventive step	A	Document indicating technological background and/or state of the art.
Y	Document indicating lack of inventive step if combined with one or more other documents of same category.	P	Document published on or after the declared priority date but before the filing date of this invention.
&	Member of the same patent family	E	Patent document published on or after, but with priority date earlier than, the filing date of this application.

Field of Search:

Search of GB, EP, WO & US patent documents classified in the following areas of the UKC^X :

G5R; H4F; H4R

Worldwide search of patent documents classified in the following areas of the IPC

G06T; G10L; G11B

The following online and other databases have been used in the preparation of this search report

On-line: WPI, EPDOC, INSPEC