



159159



G00002

- 1 -

VOICE ACTIVITY DETECTOR

Field of the Invention

5

The present invention relates to a voice activity detector, and a process for detecting a voice signal.

Background of the Invention

10

In a number of speech processing applications it is important to determine the presence or absence of a voice component in a given signal, and in particular, to determine the beginning and ending of voice segments. Detection of simple energy thresholds has been used for this purpose, however, satisfactory results only tend to be obtained where
15 relatively high signal to noise ratios are apparent in the signal.

Voice activity detection generally finds applications in speech compression algorithms, karaoke systems and speech enhancement systems. Voice activity detection processes typically dynamically adjust the noise level detected in the signals to facilitate detection of
20 the voice components of the signal.

The International Telecommunication Union (ITU) prescribes the following standards for a voice activity detector (VAD):

- 25
1. ITU-T G.723.1 Annex A, Series G: Transmission Systems and Media, "Silence compression scheme", 1996.
 2. ITU-T G.729 Annex B, Series G: Transmission Systems and Media, "A
30 silence compression scheme for G.729 optimized for terminals conforming to recommendation V.70", 1996.

- 2 -

The European Telecommunication Standards Institute (ETSI) prescribes the following standard for a VAD:

1. ETSI EN 301 708 V7.1.1, Digital cellular telecommunications system
5 (Phase 2+); "Voice Activity Detector (VAD) for adaptive Multi-Rate
(AMR) speech traffic channels: general description", 1999.

The basic function of the ETSI VAD is to indicate whether each 20 ms frame of an input
signal sampled at 16kHz contains data that should be transmitted, i.e. speech, music or
10 information tones. The ETSI VAD sets a flag to indicate that the frame contains data that
should be transmitted. A flow diagram of the processing steps of the ETSI VAD is shown
in Figure 1. The ETSI VAD uses parameters of the speech encoder to compute the flag.

The input signal is initially pre-emphasized and windowed into frames of 320 samples.
15 Each windowed frame is then transformed into the frequency domain using a Discrete
Time Fourier Transform (DTFT).

The channel energy estimate for the current sub-frame is then calculated based on the
following:

20

1. the minimum allowable channel energy;
2. a channel energy smoothing factor;
3. the number of combined channels; and
4. elements of the respective low and high channel combining tables.

25

The channel Signal to Noise Ratio (SNR) vector is used to compute the voice metrics of
the input signal. The instantaneous frame SNR and the long-term peak SNR are used to
calibrate the responsiveness of the ETSI VAD decision.

- 3 -

The quantized SNR is used to determine the respective voice metric threshold, hangover count and burst count threshold parameters. The ETSI VAD decision can then be made according to the following process:

```

5      if ( v(m)>v th +  $\mu(m)$ )
      {          /* if the voice metric > voice metric threshold*/

          VAD(m)=ON
          b(m)=b(m-1) + 1      /* increment burst counter */
10
          if ( b(m)>b th )
          {          /* compare counter with threshold */

              h(m)=h cnt      /* set hangover */
15      }
      }
      else
      {
20
          b(m) = 0      /* clear burst counter */
          h(m)=h(m-1) -1      /* decrement hangover /

          if ( h(m) <= 0 )
25      {          /* check for expired hangover */

              VAD(m)=OFF
              h(m)=0
30      }
          else
          {          /* hangover not yet expired */

              VAD(m) = ON
35
          }
      }

```

To avoid being over-sensitive to fluctuating, non-stationary, background noise conditions, a bias factor may be used to increase the threshold on which the ETSI VAD decision is based. This bias factor is typically derived from an estimate of the variability of the background noise estimate. The variability estimate is further based on negative values of the instantaneous SNR. It is presumed that a negative SNR can only occur as a result of fluctuating background noise, and not from the presence of voice. Therefore, the bias factor is derived by first calculating the variability factor. The spectral deviation estimator

is used as a safeguard against erroneous updates of the background noise estimate. If the spectral deviation of the input signal is too high, then the background noise estimate update may not be permitted.

- 5 The ETSI VAD needs at least 4 frames to give a reliable average speech energy with which the speech energy of the current data frame can be compared.

A typical problem faced by a VAD is misclassification of the input signal into voice / silence regions. Some standard algorithms vary the noise threshold dynamically across a
 10 number of frames and produce more accurate VAD estimates with time. However, the complexity of these VADs is relatively high. The complexity of the ETSI VAD may be given as follows:

$$\text{ETSI VAD} = \{ 2.O(L) + O(M.\log_2(M)) + 4.O(N_c) \} \text{ operations}$$

15

where

N_c is the number of combined channels;

L is the subframe length; and

20

M is the DFT length.

Windowing and pre-emphasis both have an order of $O(L)$. The Discrete Time Fourier Transform has an order of $O(M.\log_2(M))$. The channel energy estimator, Channel SNR estimator, voice metric calculator and Long-term Peak SNT calculator each have
 25 complexity of the order of $O(N_c)$.

These VADs are typically not efficient for applications that require low-delay signal dependant estimation of voice / silence regions of speech. Such applications include pitch detection of speech signals for karaoke. If a noisy signal is determined to be a speech
 30 track, the pitch detection algorithm may return an erroneous estimate of the pitch of the signal. As a result, most of the pitch estimates will be lower than expected, as shown in

Figure 2. The ETSI VAD supports a low-delay VAD estimate based on pre-fixed noise thresholds, however, these thresholds are not signal dependent.

An object of the present invention is to overcome or ameliorate one or more of the above
5 mentioned difficulties, or at least provide a useful alternative.

Summary of the Invention

In accordance with the present invention, there is provided a method for determining
10 whether a data frame of a coded speech signal corresponds to voice or to noise, including the steps of:

- determining the cross-correlation of the data of said data frame;
- determining the periodicity of the cross-correlation;
- determining the variance of the periodicity;
- 15 determining said data frame corresponds to noise if the cross-correlation is lower than a predetermined cross-correlation value; and
- determining the data corresponds to voice if the variance is less than a predetermined variance value.

20 The present invention also provides a method for determining whether a data frame of a coded speech signal corresponds to voice or to noise, including the steps of:

- determining an energy of said frame;
- determining an average speech energy of the coded speech signal;
- if the data frame is one of a predetermined number of initial data frames of the
- 25 coded speech signal, performing the method referred to above; and
- else, comparing the energy of the frame with the average speech energy, and the data frame corresponds to speech if the average speech energy is less than or equal to that of the energy of the frame.

30 The present invention also provides a voice activity detector for determining whether a data frame of a coded speech signal corresponds to voice or to noise, including:

- 6 -

- means for determining the cross-correlation of the data of said data frame;
- means for determining the periodicity of the cross-correlation;
- means for determining the variance of the periodicity;
- means for determining said data frame corresponds to noise if the cross-correlation
- 5 is lower than a predetermined cross-correlation value; and
- means for determining the data corresponds to voice if the variance is less than a predetermined variance value.

Brief Description of the Drawings

10

Preferred embodiments are hereafter described, by way of non-limiting example only, with reference to the accompanying drawings in which:

- Figure 1 is a block diagram showing an ESTI Voice Activity Detector;
- 15 Figure 2 is a graphical illustration of pitch estimation of speech determined using a known and voice activity detector;
- Figure 3 is a diagrammatic illustration of a voice activity detector in accordance with a preferred embodiment of the invention;
- Figure 4 is a flow diagram showing a process preferred by the voice activity detector;
- 20 Figure 5 shows the frequency spectrum and cross-correlation of speech and noise signals;
- Figure 6 is a graphical illustration showing the distance between adjacent peaks in the cross-correlation of speech signals;
- Figure 7 is a graphical illustration showing the distance between adjacent peaks in the cross-correlation of brown noise signals;
- 25 Figure 8 is a graphical illustration of pitch estimation of speech determined using a voice activity detector in accordance with a preferred embodiment of the invention.
- Figure 9 is a flow diagram showing a process preferred by the voice activity detector

30

Detailed Description of Preferred Embodiments of the Invention

A voice activity detector (VAD) 10, as shown in Figure 3, receives coded speech input signals, partitions the input signals into data frames and determines, for each frame,
 5 whether the data relates to voice or noise. The VAD 10 operates in the time domain and takes into account the inherent characteristics of speech and coloured noise to provide improved distinction between speech and silenced sections of speech. The VAD 10 preferably executes a VAD process 12, as shown in Figure 4.

10 Coloured noise has the following fundamental properties:

1. White noise: the power of the noise is randomly distributed over the entire frequency spectrum and the correlation is very low.
2. Brown noise: the frequency spectrum, $(1 / f^2)$, is mostly dominant in the
 15 very low frequency regions. Brown noise has a high cross correlation like speech signals.
3. Pink noise: the frequency spectrum, $(1 / f)$, is mostly present in the low frequencies. The cross-correlation values of Pink noise are not comparable to those of speech signals.

20

Figure 5 shows the frequency spectrum and cross-correlation of speech and coloured noise signals, where the cross-correlation is computed by varying the lag from 0 to 2048 samples. As can be observed from Figure 5(a), speech is highly correlated due to the higher number of harmonics in the spectrum. The correlation is also highly periodic.

25

The VAD 10 takes into account the above-described statistical parameters to improve the estimate of the initial frames. The cross-correlation of the signal is determined to obtain a VAD estimate in the initial frames of the input. Speech samples are highly correlated and the correlation is periodic in nature due to harmonics in the signal. Figure 6 shows the
 30 distance between adjacent peaks in speech cross-correlation. Figure 7 shows the distance between adjacent peaks in brown noise cross-correlation. As can be observed, the

- 8 -

estimates of the periodicity of the peaks in the speech samples are more stable than those of pink and brown noise. A variance estimation method is described below that successfully differentiates between speech and noise.

- 5 After a certain number of frames, the energy threshold estimator also helps to improve the distinction between the voiced and silenced sections of the speech signal. The short-term energy signal is determined to adaptively improve the voiced/silence detection across a large number of frames.
- 10 The VAD 10 receives, at step 20 of the process shown in Figure 4, Pulse Code Modulated (PCM) signals as input. The input signal is sampled at 12,000 samples per second. The sampled PCM signals are divided into data frames, each frame containing 2048 samples. Each input frame is further partitioned into two sub-frames of 1024 samples each. Each pair of sub-frames is used to determine cross-correlation.
- 15 The VAD 10 then determines, at step 22, the amount of short-term energy in the input signal. The short-term energy is higher for voiced than un-voiced speech and should be zero for silent regions in speech. Short-term energy is calculated using the following formula:

20

$$E^l = \sum_{n=(l-1)N+1}^{lN} x(n)^2 \quad (1)$$

The energy in the l^{th} analysis frame of size N is E^l . The average energy thresholds are determined, at step 22, as follows:

25

$$E_s^a = \frac{1}{k} \sum_{l=1}^k E^l \quad \text{and} \quad E_n^a = \frac{1}{k} \sum_{l=1}^k E^l \quad (2)$$

where

E_s^a is the average speech energy over k frames classified as speech and
 E_n^a is the average noise energy over k frames classified as speech.

- 5 Where the current data frame being processed is the fifth or greater in a series of data frames, the VAD 10 compares, at step 23, the energy of the current frame with the average speech energy E_s^a to determine whether it contains speech or noise.

Otherwise, the VAD 10 determines, at step 24, the cross-correlation, $Y(\tau)$, of the first and
 10 second sub frames of the data frame under consideration as follows:

$$Y(\tau) = \sum_{n=0}^{N/2-1} x_1(n)x_2(n+\tau) \quad (3)$$

where,

- τ is the lag between the sequences,
 15 $x_1(n)$ is the first half of the input frame under consideration
 $x_2(n)$ is the second half of the input frame under consideration and
 N is the size of the frame.

Input signals with cross-correlation lower than 0.4 are considered as noise. This test
 20 therefore detects the presence of either white or pink noise in the data frame under consideration. Further tests are conducted to determine whether the current data frame is speech or brown noise.

As discussed above, the cross-correlation of speech samples is highly periodic. The
 25 periodicity of the cross-correlation of the current data frame is determined, at step 26, to segregate speech and noisy signals. The periodicity of the cross-correlation can be measured, with reference to Figure 6, by determining the:

1. Distance between positive peaks: Diff_{pp}
- 30 2. Distance between negative peaks: Diff_{nn}

- 10 -

3. Distance between consecutive positive and negative peaks: Diff_{pn}
4. Distance between consecutive negative and positive peaks: Diff_{np}

The peaks can be identified by using:

5

$$\begin{aligned} Y(\tau-1) < Y(\tau) > Y(\tau+1) & \quad \text{for maxima and} \\ Y(\tau-1) > Y(\tau) < Y(\tau+1) & \quad \text{for minima.} \end{aligned}$$

To ensure spurious peaks are not chosen, the process is extended to cover five lags on
10 either side of a trial peak lag. In doing so, makes the peak detection criteria is stringent and
entails the risk of leaving out genuine peaks in the cross correlation.

The variance of periodicity is determined at step 28. The variance σ^2 is a measure of how
spread out a distribution is and is defined as the average squared deviation of each number
15 in the sequence from its mean, i.e.

$$\sigma^2 = \frac{\sum (x - \mu)^2}{L} \quad (4)$$

where

- x is the sequence whose variance is being measured and can be any of the
20 Diff_{xx} sequences mentioned in the previous section;
- μ is the mean of sequence x ; and
- L is the number of samples in the sequence i.e. the number of peaks in the
different cases.

25 The estimate is normalised by L as the number of peaks in the correlation of speech and
noisy samples will be different. To obtain an accurate estimate of the variance of the
periodicity, a linear combination of the variances of the Diff_{xx} is taken.

From Figure 6, it can be seen that the mean of the Diff_{xx} sequences of speech signals is
30 higher as compared to that of noisy signals. To take into account the percentage variation

- 11 -

of the $Diff_{xx}$ sequences from their respective means rather than the absolute variation, σ^2 is further normalised by μ^2 .

$$\varepsilon = \frac{\sigma^2}{\mu^2} = \frac{\sum (x - \mu)^2}{L \cdot \mu^2} = \frac{1}{L} \sum \left\{ \left(\frac{x}{\mu} \right)^2 - 1 \right\} \quad (5)$$

5

Equation 5 varies according to $0 < \varepsilon < 1$. The variance of the periodicity of the cross-correlation of speech signals is therefore lower than that of noise. The content of the relevant data frame may be considered to be voice if $\varepsilon < 0.2$, for example.

- 10 The VAD 10 experiences a delay of one data frame, ie the time taken for the first 2048 bits of sampled input signal to fill the first data frame. With a sampling frequency of 12 kHz., the VAD 10 will experience a lag of 0.17 seconds. The computation of the cross-correlation values for different lags takes minimal time. The VAD 10 may reduce the lag by reducing the frame size to 1024 samples. However, the reduced lag comes at the
- 15 expense of increasing the error margin in the computation of the variance of the periodicity of the cross-correlation. This error can be reduced by overlapping the sub-frames used for the correlation.

Figure 8 shows the effect of the VAD 10 when used for pitch detection in a karaoke

20 application. The average pitch estimate has improved in comparison with the pitch estimation shown in Figure 2 obtained using a known VAD that gradually adapts the energy thresholds over a number of frames.

The number of computations required for the computation of the correlation values

25 initially, reduce with higher number of frames, which dynamically adapt to the SNR of the input signal. The initial order of computational complexity is:

$$O(N) + O(N^2/2) + 5.O(K) \quad (7)$$

where

- 30 N is the number of samples in a frame; and

K is the number of peaks detected in the auto-correlation function.

In the steady state, when the energy thresholds have been determined, the order of complexity of the process VAD 10 reduces to $2.O(N)$.

5

The VAD 10 may alternatively execute a VAD process 50, as shown in Figure 9. The VAD 10 receives, at step 52, Pulse Code Modulated (PCM) signals as input. The input signal is sampled at 12,000 samples per second. The sampled PCM signals are divided into data frames, each frame containing 2048 samples. Each input frame is further
10 partitioned into two sub-frames of 1024 samples each. Each pair of sub-frames is used to determine cross-correlation.

The VAD 10 determines, at step 54, the cross-correlation, $Y(\tau)$, of the first and second sub frames of the data frame under consideration using Equation (3) Input signals with cross-
15 correlation lower than 0.4 are considered as noise. This test therefore detects the presence of either white or pink noise in the data frame under consideration. Further tests are conducted to determine whether the current data frame is speech or brown noise.

As discussed above, the cross-correlation of speech samples is highly periodic. The
20 periodicity of the cross-correlation of the current data frame is determined, at step 56, to segregate speech and noisy signals. The periodicity of the cross-correlation can be measured in the above-described manner with reference to Figure 6.

The variance of periodicity is determined at step 58 in the above-described manner. The
25 estimate is normalised by L as the number of peaks in the correlation of speech and noisy samples will be different. To obtain an accurate estimate of the variance of the periodicity, a linear combination of the variances of the $Diff_{xx}$ is taken.

From Figure 6, it can be seen that the mean of the $Diff_{xx}$ sequences of speech signals is
30 higher as compared to that of noisy signals. To take into account the percentage variation of the $Diff_{xx}$ sequences from their respective means rather than the absolute variation, σ^2 is

- 13 -

further normalised by μ^2 as given by Equation 5. The variance of the periodicity of the cross-correlation of speech signals is therefore lower than that of noise. The content of the relevant data frame may be considered to be voice if $\epsilon < 0.2$, for example.

- 5 The VAD 10 sets a flag indicating whether the contents of the relevant data frame is voice.

THE CLAIMS OF DEFINING THE INVENTION ARE AS FOLLOWS:-

1. A method for determining whether a data frame of a coded speech signal corresponds to voice or to noise, including the steps of:
 - 5 determining the cross-correlation of the data of said data frame;
 - determining the periodicity of the cross-correlation;
 - determining the variance of the periodicity;
 - determining said data frame corresponds to noise if the cross-correlation is lower than a predetermined cross-correlation value; and
 - 10 determining the data corresponds to voice if the variance is less than a predetermined variance value.
2. The method claimed in claim 1, wherein the cross-correlation, $Y(\tau)$, is calculated in accordance with the following:
 - 15
$$Y(\tau) = \sum_{n=0}^{N/2-1} x_1(n)x_2(n+\tau)$$
 - where,
 - 20 τ is the lag between the sequences $x_1(n)$ and $x_2(n)$;
 - $x_1(n)$ is the first half of data frame;
 - $x_2(n)$ is the second half of the data frame; and
 - N is the size of the frame.
- 25 3. The method claimed in claim 1 or claim 2, wherein the predetermined cross-correlation value corresponds to that of white or pink noise.
4. The method claimed in any one of claims 1 to 3, wherein the predetermined correlation value is 0.4.

- 15 -

5. The method claimed in any one of claims 2 to 4, wherein the periodicity is determined by measuring:

- (a) distance between positive peaks: Diff_{pp} ;
- (b) distance between negative peaks: Diff_{nn} ;
- 5 (c) distance between consecutive positive and negative peaks: Diff_{pn} ; and
- (d) distance between consecutive negative and positive peaks: Diff_{np}

where the peaks are identified by using:

$$\begin{array}{ll} 10 & Y(\tau-1) < Y(\tau) > Y(\tau+1) \quad \text{for maxima and} \\ & Y(\tau-1) > Y(\tau) < Y(\tau+1) \quad \text{for minima.} \end{array}$$

6. The method claimed in claim 5, wherein the variance, σ^2 , is calculated as follows:

$$15 \quad \sigma^2 = \frac{\sum (x - \mu)^2}{L}$$

where

- x is the sequence whose variance is being measured;
- 20 μ is the mean of sequence x ; and
- L is the number of samples in the sequence.

7. The method claimed in claim 6, wherein the variance is normalised by μ^2 substantially as follows:

$$25 \quad \varepsilon = \frac{\sigma^2}{\mu^2} = \frac{\sum (x - \mu)^2}{L \cdot \mu^2} = \frac{1}{L} \sum \left\{ \left(\frac{x}{\mu} \right)^2 - 1 \right\}$$

8. The method claimed in claim 7, wherein the predetermined variance value is 0.2

9. A method for determining whether a data frame of a coded speech signal corresponds to voice or to noise, including the steps of:

determining an energy of said frame;

determining an average speech energy of the coded speech signal;

- 5 if the data frame is one of a predetermined number of initial data frames of the coded speech signal, performing the method claimed in any one of claims 1 to 8; and

10 else, comparing the energy of the frame with the average speech energy, and the data frame corresponds to speech if the average speech energy is less than or equal to that of the energy of the frame.

10. The method claimed in claim 9, wherein determining the energy of the frame by determining:

15
$$E^l = \sum_{n=(l-1)N+1}^{lN} x(n)^2$$

where the energy in the l^{th} analysis frame of size N is E^l .

11. The method claimed in claim 10, wherein the average speech energy determined over k data frames is as follows:
- 20

$$E_s^a = \frac{1}{k} \sum_{l=1}^k E^l$$

12. A voice activity detector for determining whether a data frame of a coded speech signal corresponds to voice or to noise, including:
- 25

means for determining the cross-correlation of the data of said data frame;

means for determining the periodicity of the cross-correlation;

means for determining the variance of the periodicity;

- 30 means for determining said data frame corresponds to noise if the cross-correlation is lower than a predetermined cross-correlation value; and

means for determining the data corresponds to voice if the variance is less than a predetermined variance value.

13. The voice activity detector claimed in claims 12, wherein the cross-correlation,
5 $Y(\tau)$, is calculated in accordance with the following:

$$Y(\tau) = \sum_{n=0}^{N/2-1} x_1(n)x_2(n+\tau)$$

where,

10

τ is the lag between the sequences $x_1(n)$ and $x_2(n)$;
 $x_1(n)$ is the first half of data frame;
 $x_2(n)$ is the second half of the data frame; and
 N is the size of the frame.

15

14. The voice activity detector claimed in claim 12 or claim 13, wherein the predetermined cross-correlation value corresponds to that of white or pink noise.

15. The voice activity detector claimed in any one of claims 12 to 14, wherein the
20 predetermined correlation value is 0.4.

16. The voice activity detector claimed in any one of claims 14 to 15, wherein the periodicity is determined by measuring:

- (a) distance between positive peaks: Diff_{pp} ;
 25 (b) distance between negative peaks: Diff_{nn} ;
 (c) distance between consecutive positive and negative peaks: Diff_{pn} ; and
 (d) distance between consecutive negative and positive peaks: Diff_{np}

wherein the peaks are identified by using:

30

$Y(\tau-1) < Y(\tau) > Y(\tau+1)$ for maxima and
 $Y(\tau-1) > Y(\tau) < Y(\tau+1)$ for minima.

17. The voice activity detector claimed in claim 16, wherein the variance, σ^2 , is calculated as follows:

$$\sigma^2 = \frac{\sum (x - \mu)^2}{L}$$

where

- x is the sequence whose variance is being measured;
 μ is the mean of sequence x ; and
 L is the number of samples in the sequence.

18. The voice activity detector claimed in claim 17, wherein the variance is normalised by μ^2 substantially as follows:

$$\varepsilon = \frac{\sigma^2}{\mu^2} = \frac{\sum (x - \mu)^2}{L \cdot \mu^2} = \frac{1}{L} \sum \left\{ \left(\frac{x}{\mu} \right)^2 - 1 \right\}$$

19. The voice activity detector claimed in claim 18, wherein the predetermined variance value is 0.2

20. A method for determining whether a data frame of a coded speech signal corresponds to voice or to noise substantially as hereinbefore described with reference to drawings 3 to 8.

21. A voice activity detector substantially as hereinbefore described with reference to Figures 3 to 8.