

PITCH DETECTION OF SPEECH SIGNALS

Technical Field

5 [001] The present invention relates to the pitch detection of speech signals for various applications, and in particular, to a method and system providing pitch detection of speech signals for use in various audio effects, karaoke, scoring, voice recognition, etc.

10

Background Art

[002] Pitch detection of speech signals finds applications in various audio effects, karaoke, scoring, voice recognition etc. The pitch of a signal is the
15 fundamental frequency of vibration of the source of the tone.

[003] Speech signals can be segregated into two segments: voiced; and unvoiced speech. Voiced speech is
20 produced using the vocal cords and is generally modelled as a filtered train of impulses within a frequency range. Unvoiced speech is generated by forcing air through a constriction in the vocal tract. Pitch detection involves the determination of the continuous pitch period during the
25 voiced segments of speech.

[004] The terms "speech" and "speech signal" are a broad reference to all forms of generated audio or sound. For example, "speech" and its associated "speech signal" can
30 refer to talking, singing, attempted singing, whistling, humming, a recital, etc.. The "speech" and "speech signal" can originate from an individual or a group, being human, animal or otherwise. The "speech" could also be artificially

generated, for example by a computer or other electronic device.

[005] There exist presently known techniques for pitch
5 detection (see W. Hess, "Pitch Determination of Speech
Signals: Algorithms and Devices", Springer-Verlag, 1983). A
time based pitch detector, estimates the pitch period by
determining the glottal closure instant (GCI) and measuring
the time period between each "event". Frequency domain pitch
10 detection can then be used to determine the pitch. Thus, the
speech signal is processed period-by-period.

Autocorrelation Techniques

[006] Correlation is the measure of similarity of two
15 input functions, and in the case of the autocorrelation
function $\Gamma(d)$, the input functions are the same signal $x(n)$,
as shown in Equation 1,

$$\Gamma(d) = \lim_{N \rightarrow \infty} \frac{1}{2N+1} \sum_{n=-N}^{+N} x(n) \cdot x(n+d) \quad (1)$$

20 where, d represents the lag or delay between the signal and
a delayed segment, and N represents the number of samples of
the input under consideration. If the signal is periodic or
quasi-periodic, the similarities between $x(n)$ and $x(n+d)$ are
25 higher. The correlation coefficients are also high if the
lag is equal to a period or a multiple of a period.

[007] As the autocorrelation function (ACF) is the
Inverse Fourier Transform of the power spectrum of the input
30 signal, the pitch is chosen as the frequency (f_s/d) at which
the maximum of the ACF occurs; i.e. where f_s is the sampling
frequency of the speech signal. Complications due to unknown

phase relations and formant structures do not arise, as the technique is independent of these parameters.

Average Magnitude Difference Function

5 [008] Signals that are similar do not exhibit a lot of differences. Thus, periodicity can be detected by investigation of the global deviation between the signals. The Average Magnitude Difference Function (AMDF) is defined as follows:

10

$$AMDF(d) = \frac{1}{K} \sum_{n=q}^{q+K-1} |x(n) - x(n+d)| \quad (2)$$

where, K is the number of samples in a frame and q is the initial sample of the frame. AMDF has a strong minimum when
15 the lag d is equal to the period of the input $x(n)$. This minimum is exactly zero if the input is exactly periodic and the frequency (f_s/d) denotes the pitch of the signal. The algorithm is phase insensitive as the harmonics are removed without regard to their phase.

20

Component Frequency Ratios

[009] An advantage of operating in the frequency domain in contrast to other domains is that the accuracy of the pitch estimate can be improved by interpolation techniques.
25 Due to the Short Time Fourier Transformation principles used, the frequency resolution at the higher end of the spectrum is greater than at the lower end of the spectrum. Also, the fundamental might have a weak amplitude and hence it is usually computed as ratios of harmonic frequencies or
30 the difference between adjacent spectral peaks caused by higher harmonics.

- [010] In cases where the fundamental is absent, it is sufficient to measure the distance between the adjacent or even non-adjacent peaks of the spectrum, representing the higher harmonics of the periodic or quasi-periodic signal.
- 5 The ratios of the higher frequency harmonics are more accurate as the frequency resolution improves at higher frequencies. The greatest common factor is the pitch of the speech signal.

10 Time Domain Techniques

- [011] Autocorrelation techniques are susceptible to frequency overlap problems, also referred to as pitch halving or pitch doubling. Also, an autocorrelation has to be computed over a wide range of lags to determine the
- 15 optimum pitch. Though a rough idea of the pitch can be obtained from the number of zero-crossings, the number of operations required for accurate pitch detection can be computationally intensive.

- 20 [012] The AMDF algorithm is susceptible to intensity variations, noise and low frequency spurious signals, which directly affect the magnitude of the principal minimum at T_0 .

25 Frequency Domain Techniques

- [013] Since it is impractical to handle large segments of the input signal, the discrete version of the Short Time Fourier Transform (STFT), as proposed by Portnoff (M. R. Portnoff, "Implementation of the Digital Phase Vocoder Using
- 30 the Fast Fourier Transform", IEEE Transactions on Acoustics, Speech and Signal Processing, vol. ASSP-24, pp.243-248, June 1969), can be used in the signal analysis. Short time segments of the signal are "windowed" according to Fourier's theorem, which states that, any periodic waveform can be

modeled as a sum of sinusoids with varying amplitudes and frequencies.

[014] A fundamental problem, which arises due to the
5 STFT, is "smearing" of the frequency response, which is
illustrated in Figs. 1a-d (prior art). If the signal
frequency coincides with one of the "bin" frequencies of the
STFT, the original amplitude is retained after the STFT.
However, if the signal frequency lies in between two
10 adjacent bin frequencies of the STFT, the energy is spread
over the entire spectrum, as is comparatively illustrated in
Figs. 1(a) and 1(b), where the y-axis presents signal
amplitude in a logarithmic scale. Also, in the later case,
as the peak frequency lies between two adjacent frequency
15 bins, the amplitude detected is less. This is comparatively
illustrated in Figs. 1(c) and 1(d), which plot the amplitude
spectrum in a linear scale. If the amplitude of the pitch
frequency is too small, it might not be quantified as a
potential candidate. Hence, it is critical to determine the
20 true frequency of the signal.

[015] This identifies a need for pitch detection of
speech signals which overcomes or at least ameliorates the
problems inherent in the prior art.
25

Disclosure Of Invention

[016] By taking into account harmonic relationships
within the signal spectrum while calculating the pitch, the
present invention is able to eliminate the pitch halving and
30 pitch doubling problems faced by standard time domain
algorithms.

[017] To resolve the issue of estimating peak frequencies inaccurately due to frequency "smearing", the exact frequency of a peak is determined by using phase interpolation techniques. The harmonic relationship of the signal and a pitch-tracking algorithm are used to improve the reliability of the pitch estimate.

[018] According to a broad form of the present invention there is provided a system for determining the pitch of speech from a speech signal, the system including:

(1) an input device to receive the speech and generate the speech signal; and,

(2) a processor, the processor adapted to:

(a) distinguish the speech signal into voiced, unvoiced or silence sections using speech signal energy levels;

(b) apply a Fourier Transform to the voiced speech signal and obtain speech signal parameters;

(c) determine peaks of the Fourier transformed speech signal;

(d) track the speech signal parameters of the determined peaks to select partials; and,

(e) determine the pitch from the selected partials using a two-way mismatch error calculation.

[019] According to particular features of an embodiment of the invention, the speech signal is a coded, compressed or real-time audio or data signal, and the system is adapted to perform real-time processing of live speech signals.

[020] According to another broad form of the present invention there is provided a method of determining the pitch of speech from a speech signal, the method including the steps of:

- (1) producing or obtaining the speech signal;
- (2) distinguishing the speech signal into voiced, unvoiced or silence sections using speech signal energy levels;
- 5 (3) applying a Fourier Transform to the voiced speech signal and obtaining speech signal parameters;
- (4) determining peaks of the Fourier transformed speech signal;
- (5) tracking the speech signal parameters of the
- 10 determined peaks to select partials; and,
- (6) determining the pitch from the selected partials using a two-way mismatch error calculation.

[021] Preferably, but not necessarily, prior to applying
15 the Fourier Transform a windowing procedure is applied to the speech signal. Also preferably, the windowing procedure utilises a Blackman window, a Kaiser window, a Raised Cosine window or other sinusoidal model.

20 [022] In a particular embodiment, the Fourier Transform incorporates a frame size. Preferably, the frames are overlapping. In a further particular embodiment, the signal parameters form trajectories that are tracked over a selected number of frames. Preferably, trajectories
25 persisting over more than one frame are utilised.

[023] Also preferably, the signal parameters are frequency, phase and amplitude. In a further particular embodiment, a zero padding procedure is used in determining
30 the peaks of the Fourier transformed speech signal. In still a further particular embodiment, a determined peak falling within a specified frequency range of a harmonic of the pitch is set to the frequency of the harmonic.

[024] Preferably, the two-way mismatch error calculation compares each measured partial to the nearest predicted harmonic and each predicted harmonic to the nearest measured partial to provide a total error.

5

[025] According to yet another broad form of the present invention there is provided a system for estimating the pitch of speech from a speech signal, the system including:

- (1) an input device to receive the speech and produce
10 the speech signal;
- (2) a memory unit or storage unit adapted to communicate required data to a processing unit;
- (3) the processing unit operating on the speech signal and adapted to:
 - 15 (a) section the speech signal into voiced, unvoiced or silence sections using speech signal energy levels;
 - (b) apply a Fast Fourier Transform to the voiced speech signal and generate speech signal parameters;
 - 20 (c) calculate peaks of the Fourier transformed speech signal;
 - (d) track the speech signal parameters of the determined peaks to select partials; and,
 - (e) calculate the pitch from the selected
25 partials using a two-way mismatch error calculation.

[026] According to the invention, frequency domain approaches for pitch detection of speech signals are preferred, as they have been found to provide better
30 results. According to other possible aspects of the invention, an energy estimator can be utilised to help detect the voiced and silence sections of the speech signal. The frequency domain parameters can be obtained from a sinusoidal model by windowing overlapping segments of the

signal and taking a Fast Fourier Transform (FFT). However,
other waveform or function models can be utilised in the
windowing procedure. The accurate determination of the peaks
in the frequency spectrum is important. The harmonic
5 relationship of the signal is considered in the pitch
estimate by considering peaks falling within a specified
range of a harmonic.

[027] A further possible aspect of the invention, which
10 can improve performance, is a pitch-tracking block, which
can assist to obtain accurate estimates of the pitch of the
signal based on previous frames. A pitch-tracking
method/algorithm can be used to estimate the pitch of
successive frames.

15

Brief Description Of Figures

[028] The present invention should become apparent from
the following description, which is given by way of example
only, of a preferred but non-limiting embodiment thereof,
20 described in connection with the accompanying figures.

[029] Figs. 1a and 1b (prior art) illustrate example
logarithmic amplitude responses of: (a) a sinusoid at the
bin frequency, and, (b) a sinusoid between adjacent bins
25 showing spreading;

[030] Figs. 1c and 1d (prior art) illustrate example
linear amplitude responses of: (a) a sinusoid at the bin
frequency, and, (b) a sinusoid between adjacent bins showing
spreading;

30 [031] Fig. 2 illustrates a method for the pitch
detection of speech signals using frequency domain
techniques;

[032] Fig. 3 illustrates a 50% overlap added in a raised cosine window;

[033] Figs. 4a, 4b, and 4c illustrate trajectory continuations for: (a) death of tracks; (b) matching of
5 tracks; and (c) birth of tracks;

[034] Fig. 5 illustrates the spectrum of the raised cosine window;

[035] Fig. 6 illustrates the effect of windowing the signal of the spectrum;

10 [036] Fig. 7 illustrates the effect of zero padding the spectrum;

[037] Fig. 8 illustrates the mismatch error for different fundamental frequencies;

[038] Figs. 9a and 9b illustrate (a) amplitude modulated
15 input with multiple sinusoids (in the time domain), and (b) input with multiple sinusoids (in the frequency domain);

[039] Fig. 10 illustrates pitch estimates of multiple sinusoids;

[040] Fig. 11 illustrates pitch estimates of a frequency
20 chirp;

[041] Fig. 12 illustrates pitch estimates of speech signals for three different speakers;

[042] Fig. 13 illustrates a functional block diagram of a processing system embodiment of the present invention.

25

Modes For Carrying Out The Invention

[043] The following modes are described as applied to the description and claims in order to provide a more precise understanding of the subject matter of the present
30 invention.

Preferred embodiment

[044] In the figures, incorporated to illustrate the features of the present invention, like reference numerals are used to identify like parts throughout the figures.

5 [045] A sinusoidal model (see T. F. Quatieri and R. J. McAulay, "Speech transformations based on a sinusoidal representation", IEEE Transactions on Acoustics, Speech and Signal Processing, December 1986, vol. 34, no. 6, pg. 1449) is utilised, in which the speech signal $x(n)$, can be
10 represented as the sum of sinusoids of varying amplitudes (A_k^l) and frequency peaks (m). (L_k = Signal Bandwidth / Pitch) is the maximum number of frequencies in the frame. That is,

$$x(n) = \sum_{m=1}^{L_k} A_k^l(n) \cdot \cos(\theta_k^l(n)) \quad (3)$$

15

If ϕ_k^l is the starting phase of the of the k^{th} sinusoid in the l^{th} frame, $\theta_k^l(n)$ is defined in Equation 4,

$$\theta_k^l(n) = \frac{2\pi k n}{N} + \phi_k^l \quad (4)$$

20

[046] This allows calculation of the frequency domain parameters of the signal and use of the phase information to determine the true frequency components present in the signal. The flowchart of a preferred method 200 (that can
25 equally be interpreted as a block diagram of system components) according to the present invention is illustrated in Fig. 2.

Parameter Estimation

30 [047] As speech signals 210, consist of silenced and voiced sections, to avoid erroneous pitch detection, these segments of the input 210 are differentiated 220 at the

start of the parameter estimation phase of the algorithm using varying energy levels in the signal 210.

[048] The frequency domain parameters 230 are obtained by windowing 240 a short time segment of the signal 225 and taking its Fourier Transform 250, as described in Equation 5.

$$X(t_a^l, \Omega_k) = \sum_{n=-\infty}^{\infty} h(n) \cdot x(t_a^l + n) e^{-j \Omega_k n} \quad (5)$$

10

[049] At uniform analysis time instants $t_a^l = l \cdot R_a$ where R_a is the analysis hop factor and l is the frame number, the Fourier Transform 250 of the windowed signal 260 is calculated. If N is the size of the Fast Fourier Transform (FFT) 250, $\Omega_k = 2 \cdot \pi \cdot k / N$ is the centre frequency of the k^{th} bin.

[050] The analysis window $h(n)$ is critical for reducing frequency smearing and the window size 270 controls the frequency resolution. "Zero padding" of the frequency spectrum (see J. O. Smith, "Mathematics of the Discrete Fourier Transform (DFT)", Center for Computer Research in Music and Acoustics (CCRMA), Stanford University) is used to obtain an ideally interpolated spectrum, which is used for a better estimate of the peaks in the frequency spectrum at step 280.

Pitch Estimation

[051] Weighted lists of active frequencies within each analysis window are generated, and using basic pattern-matching procedures contiguous frequency tracks are obtained. The track frequency with the maximum number of harmonics is computed using a two-way mismatch procedure 290

and determined to be the pitch 295 of the signal 210. Reliability of the pitch frequency estimate 295 is ensured by using pitch tracking algorithms 285, which minimize the error of prediction based on estimates in the previous
5 frames.

A. Standard Block Level Implementation

Step A1. Input Format (210)

[052] The aforementioned process can be readily
10 implemented as system architecture and can handle Pulse Code Modulated (PCM) signals as input, which is a standard format of coded audio signals. The input is of CD quality, i.e. it is sampled at a rate of 44,100 samples/second. For real-time processing, the signal is processed 2048 samples in a frame,
15 which is approximately 46 milliseconds at the given sampling rate. However to maintain a 50% overlap, only 1024 samples are read in during each frame and the remaining 1024 samples are used from the previous frame.

20 Step A2. Silence / Voice Detection (220)

[053] Speech signals are usually considered as voiced or unvoiced, but in some cases they are something between these two. Voiced sounds consist of fundamental frequency (f_0) and harmonic components produced by the human vocal cords.
25 Purely unvoiced sounds have no fundamental frequency in the excitation signal and therefore harmonic structures are absent in the signal.

[054] The short-term energy is higher for voiced than
30 unvoiced speech, and should also be zero for silent regions in speech. Short-term energy allows one to calculate the amount of energy in a signal at a specific instant in time, and is defined in Equation 6.

$$E_a^l = \sum_{n=(l-1)N+1}^{lN} x(n)^2 \quad (6)$$

[055] The energy in the l^{th} analysis frame of size N is E_a^l . Depending upon the classification of the speech sample into voiced/unvoiced or silenced sections, the following pitch detection algorithm is activated. The pitch detection algorithm is preferably activated only if there is a voiced section in the signal. During noise or silence - neither has any pitch - the pitch detection algorithm is preferably not activated.

Step A3. Window Parameters (270)

[056] The choice of the analysis window is a trade-off of time and frequency resolution, which affects the smoothness of the spectrum and the detection of frequency peaks. Perfect reconstruction is not a criteria for the window shape as the algorithm is used only for pitch estimation and not for signal reconstruction. Hence, the algorithm implements windowing schemes, which provide better frequency resolution. The Blackman window (see http://www-ccrma.stanford.edu/~jos/Windows/Blackman_Harris_Window_Family.html) has a worst-case side-lobe rejection of 58 dB down, which is good for audio applications. However, the Kaiser window (see J. O. Smith, "The window method for digital filter design", Winter 1992, Mathematica notebook for Music 420 (EE367A), <ftp://ccrma-ftp.stanford.edu/pub/DSP/Tutorials/Kaiser.ma.Z>) allows control of the main-lobe width and the highest side-lobe level. If one desires less main-lobe width then a higher side-lobe level is produced, and vice versa.

[057] The windows also serve a dual purpose of reducing spectral leakage or "smearing" by tapering the data record gradually to zero at both end-points of the window. As a

result of the smooth tapering, the main lobe of the frequency response widens and the side-lobe levels decrease.

[058] Using no window is akin to using a rectangular window, unless the signal is exactly periodic in samples. It should be noted that increasing the number of samples in a frame does not reduce spectral leakage. The Raised Cosine window is given by $h(n)$:

$$h(n) = \frac{1}{2} - \frac{1}{2} * \cos\left(\frac{2 * \pi * n}{N}\right) \quad (7)$$

where, N is the same as the frame size in this case and n varies from zero to $(N-1)$. A series of overlap added raised cosine windows are shown in Fig. 3. A detailed discussion on the effect of windows in peak detection follows hereinafter. Overlapping frames ensure that the pitch estimate is updated on a regular basis.

Step A4. Fast Fourier Transform (250)

[059] The N point FFT of the windowed signal returns the amplitude, starting phases and the frequencies of the signal within the frame. For computational efficiency, N is selected as a power of two, though this is not necessarily required. The frame size, as well as the window size are given by N . The FFT can also be interpreted as a Linear Time Invariant filterbank followed by an exponential modulator, which allows one to extract the parameters 230 of the signal 210. The frequency and its corresponding amplitude and phase parameters form trajectories.

30

Step A5. Peak Detection (280)

[060] To determine the pitch of the input signal 210, peaks are detected in the amplitude spectrum. Preferably,

though not necessarily, the peaks are chosen based on their relative magnitude difference between neighbouring frequency bins. An 80dB cut-off criterion is applied to limit the number of peaks. Logarithmic plots can be used for the peak
5 frequency determination, as they are smoother than the amplitude spectrum plots. The transform of the amplitude spectrum is zero padded and the Inverse Fourier transform is computed to increase the frequency resolution and smoothen the spectrum. This step can be discarded if computational
10 efficiency is desired.

Step A6. Harmonic Detection (290)

[061] Pitch is the fundamental frequency of vibration of the source of the tone. In simple mathematical terms, it is
15 the least common divisor of the peak frequencies of the signal if it is harmonic in nature. Speech signals are harmonic in nature and hence, it is easier to determine the signal harmonics using the pitch information.

20 [062] As discussed in S. S. Abeysekera, K. P. Padhi, J. Absar and S. George, "Investigation of different frequency estimation techniques using the phase vocoder", International Symposium on Circuits and Systems, May 2001, the true frequency associated with the k^{th} bin is calculated
25 from the Fourier Transform $X(l, k)$ as defined in Equation 4, over two consecutive frames that are separated by H samples, i.e.

$$\hat{f} = \frac{k}{N} + \frac{\text{Arg}\{X(l, k)\} - \text{Arg}\{X(0, k)\}}{2\pi H} \quad (8)$$

30

[063] Accurate peak determination is essential to determine the exact pitch of the input signal 210. Besides detecting the pitch, this block is also responsible for

detecting the harmonics present in the signal. Once the peak frequencies and the pitch are detected in the signal, any peak falling within a specified range of a harmonic is forced to the frequency of the harmonic. In other words, if

5

$$|f - m.f_0| \leq \delta \quad (9)$$

where, f is the peak frequency, f_0 is the fundamental pitch frequency, m is any integer and δ is an arbitrary constant which determines how close a frequency should be before it is forced to the nearest harmonic frequency. The constant δ is constrained by the accuracy of the parameter estimation system. The higher the accuracy, the smaller the value of δ ; the coarser the parameter estimation algorithm, the larger the value of δ .

Step A7. Pitch Tracking (285)

[064] The frequency, amplitude and phase parameters 230 of the peak frequencies form trajectories, which are tracked across the frames. To avoid detecting spurious peak frequencies, only those trajectories lasting over a number of frames are chosen for harmonic matching.

[065] The tracking procedure consists of piecing together the parameters that fall within certain minimum frequency deviations and choosing trajectories that minimize frequency distance between the parameters. Assuming, all the previous peak frequencies up to bin k in frame l have been matched, and ω^l_k , A^l_k represent the frequency and amplitude parameters of bin k in frame l . The concept of death, continuation and birth of tracks is illustrated in Figs. 4(a), (b) and (c), respectively.

- if $|\omega_k^1 - \omega_q^{1+1}| \geq \Delta \Rightarrow$ the track dies $\Rightarrow A_k^{1+1} = 0$.
 - if $|\omega_k^1 - \omega_q^{1+1}| < \Delta \Rightarrow \omega_k^{1+1}$ is a "tentative" match, i.e. there might be other matching frequencies in the vicinity and hence one should check the entire frequency range.
 - if $|\omega_k^1 - \omega_q^{1+1}| < |\omega_k^1 - \omega_{i+1}^{1+1}| \Rightarrow$ if frequency ω_q^{1+1} is not matched to any other frequency and is the closest to ω_k^1 , ω_q^{1+1} is a "perfect" match.
 - All unmatched peak frequencies in frame $l+1$, are designated as new tracks born
- $\Rightarrow A_k^{1-1} = 0$.

[066] A minimum sleeping time concept ensures that long duration tracks are "killed" only if they do not recur within a specified time.

Step A8. Pitch Determination (290)

[067] The peaks in the amplitude spectrum are herein referred to as "partials" for clarity.

[068] The most likely fundamental frequencies can be chosen from the peaks in the spectrum based on the greatest common divisor of maximum number of partials in the signal spectrum. The initial pitch search could be localised to a frequency range of 110-130 Hz and 200-230 Hz, for male and female speech signals respectively, although other ranges could be selected.

[069] The two-way mismatch error calculation is a two step process in which each measured partial is compared to the nearest predicted harmonic giving the measured-to-predicted error $Err_{p \rightarrow m}$, and each predicted harmonic is compared to the nearest measured partial giving the

predicted-to-measured error $Err_{m \rightarrow p}$. The total error Err_{total} is a weighted combination of these two errors.

[070] The error is normalised by the fundamental
5 frequency and also incorporates factors, which take into account the effect of amplitudes of the partials, i.e. the Signal to Noise Ratio (SNR) on the pitch of the signal.

$$Err_{total} = \frac{Err_{p \rightarrow m}}{N} + \rho \cdot \frac{Err_{m \rightarrow p}}{K}$$

$$= \frac{1}{N} \sum_{n=1}^N \left[\frac{\Delta f_n}{f_n^p} + \frac{a_n}{A_{max}} \cdot \left\{ q \cdot \frac{\Delta f_n}{f_n^p} - r \right\} \right] + \rho \cdot \frac{1}{K} \sum_{k=1}^K \left[\frac{\Delta f_k}{f_k^p} + \frac{a_k}{A_{max}} \cdot \left\{ q \cdot \frac{\Delta f_k}{f_k^p} - r \right\} \right] \quad (10)$$

10

where, N is the number of harmonics of the trial fundamental frequency (f_{fund}) given by $N = \lfloor f_{max} / f_{fund} \rfloor$. The $\lfloor x \rfloor$ operation returns the smallest integer greater than x , f_{max} is the highest frequency and A_{max} is the maximum amplitude of the
15 measured partials. K is the total number of partials, i.e. critical frequencies in each frame.

[071] As the error is a function of the frequency difference ($\Delta f_n = \Delta f_k = |f_n - f_k|$) between the nearest
20 harmonic frequency f_n and the measured peak in the spectrum f_k , maximum error occurs when there are missing harmonics or when the ratio of the amplitudes is small. Similarly, minimum error will occur when most of the harmonics of the trial frequency are present and the ratio of the amplitudes
25 is large. Maher et. al. (see R. C. Maher and J. W. Beauchamp, "Fundamental frequency estimation of musical signals using a two-way mismatch procedure", Journal of the Acoustical Society of America, April 1994, vol. 95(4), pg. 2254) have determined that $p = 0.5$, $q = 1.4$ and $r = 0.5$
30 satisfy the above weighting properties. The frequency which

produces a minimum mismatch error is the pitch of the signal.

B. Improvements in the Pitch Detection Algorithm

5 [072] The human hearing system (the ears and the related perception system in the brain) is more sensitive to frequencies in the range of 1000 Hz - 3000 Hz. However, speech signals have a bandwidth of 20 Hz - 8 kHz. The pitch search can be localised within a range of 50 Hz - 500 Hz, as
10 beyond these frequencies mostly harmonics will be present. However, the peak detection algorithm is used over the entire speech spectrum to capture as many harmonic frequencies as possible. Larger numbers of frequencies chosen lead to an accurate determination of the pitch. In
15 this section, enhancements in the developed pitch detection method/system are discussed.

B1. Effect of Windowing

[073] By considering a sinusoidal model, the spectrum of
20 the window is shifted by the frequency of the sinusoids. The amplitude of the bins adjacent to the peak frequencies is determined by the side-lobe levels of the raised cosine spectrum of the window, as obtained in Equation 11.

$$\begin{aligned}
 W(k) &= \sum_{n=0}^{N-1} w(n) \cdot e^{\frac{-j2\pi kn}{N}} \\
 &= \sum_{n=0}^{N-1} \left\{ \frac{1}{2} + \frac{1}{2} \cos\left(\frac{2\pi n}{N}\right) \right\} \cdot e^{\frac{-j2\pi kn}{N}} \\
 25 \quad &= \sum_{n=0}^{N-1} \left\{ \frac{1}{2} + \frac{1}{2} \left(e^{\frac{j2\pi n}{N}} + e^{\frac{-j2\pi n}{N}} \right) \right\} \cdot e^{\frac{-j2\pi kn}{N}} \\
 &= \frac{N}{4} \cdot e^{\frac{-j\pi k(N-1)}{N}} \left[2 \cdot \frac{\text{Sinc}(k)}{\text{Sinc}(k/N)} + \frac{\text{Sinc}(k-1)}{\text{Sinc}((k-1)/N)} \cdot e^{j\pi(1-\frac{1}{N})} + \frac{\text{Sinc}(k+1)}{\text{Sinc}((k+1)/N)} \cdot e^{-j\pi(1+\frac{1}{N})} \right] \quad (11)
 \end{aligned}$$

[074] As can be seen from the Fig. 5, $W(0) = 2$, $W(\pm 1) = 1$, $W(k) = 0 : \forall$ all other values of k . The worst case spreading of the sinusoid spectrum occurs when the true frequency lies exactly between two frequency bins. Though
 5 the side-lobes enhance undesirable frequency components, they enhance the peak frequency components in the spectrum as shown in Fig. 6.

[075] A complex sinusoid of the form $x(n) = Ae^{j.k_s nT}$, when
 10 windowed, transforms to,

$$X_r(k) = \sum_{n=-\infty}^{\infty} x(n).h(n).e^{-j.knT} = A \sum_{n=-(M-1)/2}^{(M-1)/2} h(n).e^{-j.(k-k_s)nT} = A.W(k-k_s) \quad (12)$$

where $W(k)$ is defined in Equation 11. Thus, the transform of
 15 a windowed sinusoid, isolated or part of a complex tone, is the transform of the window scaled by the amplitude of the sinusoid and centred at the sinusoid's frequency.

B2. Effect of Frequency Padding

[076] The dual of the Zero Padding theorem (J. O. Smith, "Mathematics of the Discrete Fourier Transform (DFT)", Center for Computer Research in Music and Acoustics (CCRMA), Stanford University) states that zero padding in the frequency domain corresponds to ideal bandlimited
 25 interpolation in the time domain. As can be seen in Fig. 7, the interpolated spectrum obtained after computing the inverse transform of the zero padded Fourier spectrum is much smoother than the original spectrum. This further enhances the true peaks in the spectrum.

[077] This further signal processing coupled with an accurate determination of the true frequency of the speech ensures a superior pitch detection algorithm.

5 B3. Enhanced Pitch Detection

[078] The two-way mismatch algorithm for pitch detection solves the pitch halving and pitch doubling problems faced by traditional time domain algorithms. For each trial fundamental frequency, the two-way mismatch error is
10 computed and the frequency with the minimum error is set to be the pitch of the signal.

[079] In the present method/system, Δf_n is defined as follows,

15
$$\Delta f_n = |f_n - f_k| ; \text{ if } f_k \text{ within } \pm f_{fund}/2 \text{ Hz of } f_n$$
$$= |f_n| ; \quad \text{ if } f_k \text{ is not within } \pm f_{fund}/2 \text{ Hz of } f_n$$

[080] The same criteria is also used for calculating Δf_k . This ensures that the error is higher for missing
20 harmonics beyond the search range while putting a limit on the search criteria. This enhances the pitch detection algorithm for speech signals, which are very harmonic in nature.

25 [081] The Applicants considered a test signal containing the series of partials {100, 200, 300, 500, 600, 700, 800} Hz. For a trial fundamental frequency $f_{fund} = 50$ Hz, all the partials are harmonics, however, the harmonics at {50, 150, 250, 350, 400, 450, 550} Hz are missing. Similarly, for f_{fund}
30 $= 100$ Hz, only the harmonic at {400} Hz is missing.

[082] Fig. 8 plots the mismatch error based on Equation
10. As the mismatch error is minimum for a trial

fundamental frequency of 100 Hz, it is the fundamental frequency of the given set of partials.

[083] The different blocks in the architecture ensure
5 that the method algorithm detects the pitch accurately across successive frames.

C. Simulation Results

[084] This section demonstrates the use of frequency
10 domain techniques to determine the pitch of speech audio signals. Both artificially synthesised and natural speech signals are tested. It is essential to use synthesised signals to test the algorithm as there is no standard benchmark to compare the pitch of the signal. Since the
15 signal is synthesised, the pitch of the signal is known and hence a direct comparison is possible.

C1. Sinusoids

[085] As speech signals are represented by a sinusoidal
20 model, the algorithm is first tested on a purely sinusoidal input. The input consists of constant equal amplitude sinusoids at harmonically related frequencies of 440 Hz, and 880 Hz. The input sampling frequency is 8 kHz, the frame size is 2048 samples with a 50% overlap of 1024 samples. The
25 signal is generated over multiple frames and the amplitude is modulated and mixed with noise as presented in Figs. 9(a) and 9(b).

[086] The time-pitch frequency plot of the signal is
30 presented in Fig. 10. The x-axis denotes the time in terms of the number of frames. The y-axis shows the pitch frequency in the STFT, which satisfies the peak detection criteria and the minimum mismatch error criteria as previously discussed. As can be seen from Fig. 10, the

developed method is successfully able to determine the pitch of the input signal depending on whether the input is silence or noise or sinusoidal in nature.

5 C2. Frequency Modulated Sinusoid

[087] To test the pitch tracking algorithm, the frequency of the input is varied from 0 Hz to 4kHz over time. Fig. 11 shows the time-pitch frequency plot of the algorithm as compared to standard autocorrelation
10 techniques. As can be seen, the time domain techniques suffer from pitch halving problems, whereas the present successfully tracks pitch.

C3. Speech Signals

15 [088] Fig. 12 shows the pitch characteristics of three different male speakers speaking "A tiger and a mouse were walking in a field". Both John and Andrew are British English speakers while Dg is an African speaker of English. It can be seen that Dg's voice has a much lower pitch than
20 that of the British speakers. Fig. 12 also shows the change in the pitch of the signal according to the speaker's pronunciation as he speaks.

Various embodiments

25 [089] Other embodiments of the present invention are possible. According to another embodiment of the present invention a processing system, an example of which is shown in Fig. 13 is utilised. In particular, the processing system
1300 generally includes at least a processor or processing
30 unit 1302, a memory 1304, an input device 1306 and an output device 1308, coupled together via a bus or collection of buses 1310. An interface 1312 can also be provided for coupling the processing system 1300 to a storage device 1314 which may house a database 1316. The memory 1304 can be any

form of memory device, for example, volatile or non-volatile memory, solid state storage devices, magnetic devices, etc. The input device 1306 receives speech input 1318 and can include, for example, a microphone, a stored audio device
5 (eg. CD), a voice control device, data acquisition card, etc. The output device 1308 produces a pitch estimate output 1320 and could be, for example, a display device, internal component or electronic device, etc. The storage device 1314 can be any form of storage means, for example, volatile or
10 non-volatile memory, solid state storage devices, magnetic devices, etc.

[090] In use, the processing system 1300 is adapted to allow data or information to be stored in and/or retrieved
15 from the storage device 1314 or database 1316 if required. Alternatively, required data or information could be retrieved from memory 1304. The processor 1302 acts upon speech input 1318 in accordance with the method of the present invention. It should be appreciated that the
20 processing system 1300 may be a specialised electronic device or chip, processing system, computer terminal, server, specialised hardware or firmware, or the like.

[091] The method of the present invention could readily
25 be embodied as software, hardware, firmware or the like, or a combination thereof. Various programming languages could be utilised to realise the method.

[092] The invention may also be said to broadly consist
30 in the parts, elements and features referred to or indicated herein, individually or collectively, in any or all combinations of two or more of the parts, elements or features, and where specific integers are mentioned herein which have known equivalents in the art to which the

invention relates, such known equivalents are deemed to be incorporated herein as if individually set forth.

[093] Although the preferred embodiment has been
5 described in detail, it should be understood that various
changes, substitutions, and alterations can be made herein
by one of ordinary skill in the art without departing from
the scope of the present invention.

The claims:

1. A system for determining the pitch of speech from a speech signal, the system including:
 - 5 (1) an input device to receive the speech and generate the speech signal; and,
 - (2) a processor, the processor adapted to:
 - (a) distinguish the speech signal into voiced, unvoiced or silence sections using speech signal energy
10 levels;
 - (b) apply a Fourier Transform to the voiced speech signal and obtain speech signal parameters;
 - (c) determine peaks of the Fourier transformed speech signal;
 - 15 (d) track the speech signal parameters of the determined peaks to select partials; and,
 - (e) determine the pitch from the selected partials using a two-way mismatch error calculation.
- 20 2. The system according to claim 1, wherein the speech signal is a coded, compressed or real-time audio or data signal.
3. The system according to claim 1 or 2, adapted to
25 perform real-time processing of live speech signals.
4. The system according to any one of the claims 1 to 3, wherein the speech signal is a Pulse Code Modulated signal.
- 30 5. The system according to any one of the claims 1 to 4, wherein the system is incorporated into a karaoke system, computer system or voice recognition system.

6. The system according to any one of the claims 1 to 5, wherein the input device is a microphone or audio receiver.

7. A method of determining the pitch of speech from a speech signal, the method including the steps of:

- (1) producing or obtaining the speech signal;
- (2) distinguishing the speech signal into voiced, unvoiced or silence sections using speech signal energy levels;
- 10 (3) applying a Fourier Transform to the voiced speech signal and obtaining speech signal parameters;
- (4) determining peaks of the Fourier transformed speech signal;
- (5) tracking the speech signal parameters of the
15 determined peaks to select partials; and,
- (6) determining the pitch from the selected partials using a two-way mismatch error calculation.

8. The method according to claim 7, wherein prior to
20 applying the Fourier Transform a windowing procedure is applied to the speech signal.

9. The method according to claim 8, wherein the windowing procedure utilises a Blackman window, a Kaiser window, a
25 Raised Cosine window or other sinusoidal model.

10. The method according to any one of the claims 7 to 9, wherein the Fourier Transform incorporates a frame size.

30 11. The method according to claim 10, wherein the frames are overlapping.

12. The method according to claim 10 or 11, wherein the signal parameters form trajectories that are tracked over a selected number of frames.

5 13. The method according to claim 12, wherein trajectories persisting over more than one frame are utilised.

14. The method according to any one of the claims 7 to 13, wherein the Fourier Transform is a Fast Fourier Transform.

10

15. The method according to any one of the claims 7 to 14, wherein the signal parameters are frequency, phase and amplitude.

15 16. The method according to any one of the claims 7 to 15, wherein a zero padding procedure is used in determining the peaks of the Fourier transformed speech signal.

17. The method according to any one of the claims 7 to 16,
20 wherein a determined peak falling within a specified frequency range of a harmonic of the pitch is set to the frequency of the harmonic.

18. The method according to any one of the claims 7 to 17,
25 wherein peaks are determined in the amplitude spectrum.

19. The method according to any one of the claims 7 to 18, wherein peaks are determined in the amplitude spectrum using a logarithmic scale.

30

20. The method according to any one of the claims 7 to 19, wherein the partials are selected from the peaks based on the greatest common divisor of the maximum number of partials in the speech signal spectrum.

21. The method according to any one of the claims 7 to 20,
wherein the two-way mismatch error calculation compares each
measured partial to the nearest predicted harmonic and each
5 predicted harmonic to the nearest measured partial to
provide a total error.

22. The method according to claim 21, wherein the error is
normalised, and adjusted using a signal-to-noise ratio.

10

23. The method according to any one of the claims 7 to 22,
wherein the pitch is determined as corresponding to the
minimum two-way mismatch error.

15 24. The method according to any one of the claims 7 to 23,
wherein the signal energy levels are short-term signal
energy levels.

25. The method according to any one of the claims 7 to 24,
20 wherein distinguishing the speech signal utilises an energy
estimation calculation.

26. The method according to any one of the claims 7 to 25,
wherein the pitch is determined using a localised frequency
25 range.

27. The method according to claim 26, wherein the localised
frequency range is about 50 - 500 Hz.

30 28. The system as claimed in claim 1, the processor being
adapted to perform the method of any one of the claims 8 to
27.

29. A system for estimating the pitch of speech from a speech signal, the system including:

(1) an input device to receive the speech and produce the speech signal;

5 (2) a memory unit or storage unit adapted to communicate required data to a processing unit;

(3) the processing unit operating on the speech signal and adapted to:

10 (a) section the speech signal into voiced, unvoiced or silence sections using speech signal energy levels;

(b) apply a Fast Fourier Transform to the voiced speech signal and generate speech signal parameters;

15 (c) calculate peaks of the Fourier transformed speech signal;

(d) track the speech signal parameters of the determined peaks to select partials; and,

(e) calculate the pitch from the selected partials using a two-way mismatch error calculation.

20

30. The system as claimed in claim 29, wherein the Fast Fourier Transform operates on a frame of a windowed portion of the speech signal, and the speech signal parameters are tracked over more than one frame.

25