

NOISE REDUCTION

Broadband noise when added to a speech signal can impair the quality of the signal, reduce intelligibility, and increase listener fatigue. Since in practice much speech is recorded and transmitted in the presence of
5 noise, the problem of noise reduction is vital to the world of telecommunications, and has gained much attention in recent years.

Various classes of noise reduction algorithm have been developed, including noise suppression filtering, comb
10 filtering, and model based approaches. Known noise suppression techniques include spectral and cepstral subtraction, and Wiener filtering.

Spectral subtraction is a very successful technique for reducing noise in speech signals. This operates (see
15 for example, Boll "Suppression of Acoustic Noise in Speech using Spectral Subtraction", IEEE Trans. on Acoustics, Speech and Signal Processing, Vol. ASSP-27, No. 2, April 1979, p.113) by converting a time domain (waveform) representation of the speech signal into the frequency
20 domain, for example by taking the Fourier transform of segments of speech to obtain a set of signals representing the short term power spectrum of the speech. An estimate is generated (during speech-free periods) of the noise power spectrum and these values are subtracted from the
25 speech power spectrum signals; the inverse Fourier transform is then used to reconstruct the time-domain signal from the noise-reduced power spectrum and the unmodified phase spectrum.

A related technique is that of spectral scaling,
30 described by Eger "A Nonlinear Processing Technique for Speech Enhancement" Proc. ICASSP 1983 (IEEE) pp 18A.1.1-18.A.1.4; again the signals are transformed into frequency domain signals which are then multiplied by a nonlinear transfer characteristic so as preferentially to attenuate

low-magnitude frequency components, prior to inverse transformation. Developments of this technique, are described in our International patent application No. PCT/GB89/00049 (published as WO89/06877) or US patent
5 5,133,013.

Due to non-stationarity in the noise, the estimated noise spectrum used for spectral subtraction will be different from the actual noise spectrum during speech activity. This error in noise estimation tends to affect
10 small spectral regions of the output, and is perceived as short duration random tones, or musical noise. Whilst much lower in overall energy than the original noise, this musical noise tends to be very irritating to listen to. A similar effect occurs in the case of spectral scaling.

15 Several methods have been employed in an attempt to minimise the musical noise. Magnitude averaging can be used to reduce these artifacts, although this can result in temporal smearing, due to the non-stationarity of the speech. Another method consists of subtracting an
20 overestimate of the noise spectrum, and preventing the output spectrum from going below a pre-set minimum level. This technique can be very effective, but can lead to greater distortion to the speech.

According to the present invention there is provided
25 a noise reduction apparatus comprising:

- conversion means for converting a time-varying input signal into signals representing the magnitudes of spectral components of the input signals;
- processing means operable to effect a reduction in
30 the magnitude of low-magnitude ones of the said spectral component signals relative to that of higher magnitude ones of the said spectral component signals; and
- reconversion means to convert the said spectral component signals into a time-varying signal;
- 35 characterised by means to identify formant regions of the speech spectrum; and

This Sheet Has Been Replaced



THIS SHEET HAS BEEN
REPLACED

means to attenuate those frequency components lying outside the formant regions.

Some embodiments of the invention will now be described, by way of example, with reference to the accompanying drawings.

The known method of spectral subtraction involves, as illustrated in Figure 1, subtracting an estimate of the short term noise power spectrum from the short term power spectrum of the speech plus noise. Noisy speech signals, in the form of digital samples at a sampling rate of, for example, 10 kHz are received at an input 1. The speech is segmented (2) into 50% overlapping Hanning windows of 51ms duration and a unit 3 generates for each segment a set of Fourier coefficients using a discrete short-time Fourier transform.

If a segment of speech $\{s(t)\}$ is corrupted by additive noise $\{n(t)\}$, Then the corrupted signal $\{y(t)\}$ can be written as

$$y(t) = s(t) + n(t).$$

It can be shown that the short term power spectrum of the corrupted signal, $P_y(\omega)$, can likewise be written as the sum of the noise and speech power spectra, viz.

$$P_y(\omega) = P_s(\omega) + P_n(\omega)$$

If an estimate of the noise power spectrum, $\hat{P}_n(\omega)$, can be obtained, then an approximation $\hat{P}_s(\omega)$ to the speech power spectrum can be obtained from

$$\hat{P}_s(\omega) = P_y(\omega) - \hat{P}_n(\omega).$$

The short term power spectrum $P_y(\omega)$ is obtained by squaring (4) the Fourier coefficients from the unit 3.

The noise spectrum cannot be calculated precisely, but can be estimated during periods when no speech is present in the input signal. This condition is recognised by a voice activity detector 5 to produce a control signal C which permits the updating of a store 6 with $P_y(\omega)$ when speech is absent from the current segment. This spectrum is smoothed, for example by firstly making each frequency

This Sheet Has Been Replaced



THIS SHEET HAS BEEN
REPLACED

sample of $P_y(\omega)$ the average of several surrounding frequency samples, giving $\bar{P}_y(\omega)$, the smoothed short term power spectrum of the current frame. With a frame length of 512 samples, the smoothing may for example be performed by
 5 averaging nine adjacent samples.

This smoothed power spectrum may then be used to update a spectral estimate of the noise, which consists of a proportion of the previous noise estimate and a proportion of the smoothed short term power spectrum of the
 10 current segment. Thus the noise power spectrum gradually adapts to changes in the actual spectrum of the noise. This may be written as $\hat{P}_n(\omega) = \lambda \cdot \hat{P}_{old}(\omega) + (1-\lambda) \cdot \bar{P}_y(\omega)$ (3) where $\hat{P}_n(\omega)$ is the updated noise spectral estimate, $\hat{P}_{old}(\omega)$ is the old noise spectral estimate, $\bar{P}_y(\omega)$ is the smoothed
 15 noise spectrum from the present frame, and λ is a decay factor (e.g. a value of $\lambda=0.85$). The contents of the store 6 thus represent the current estimate $\hat{P}_n(\omega)$ of the short term noise power spectrum.

This estimate is subtracted from the noisy speech
 20 power spectrum in a subtractor 7. The harshness of the subtraction can be varied by applying a scaling factor α (in a multiplier 8) so that

$$\hat{P}_s(\omega) = P_y(\omega) - \alpha \cdot \hat{P}_n(\omega).$$

The scaling factor α would have a value of about 2.3
 25 for standard spectral subtraction, with a signal to noise ratio of 10 dB. A higher value would be used for lower signal to noise ratios. Any resulting negative terms are set to zero, since a frequency component cannot have a negative power; alternatively a non zero minimum power
 30 level may be defined, for example defining $\hat{P}_s(\omega)$ as the maximum of $P_y(\omega) - \alpha \cdot \hat{P}_n(\omega)$ and $\beta \cdot \hat{P}_n(\omega)$ where β determines the minimum power level or 'spectral floor'. A non zero value of β may reduce the effect of musical noise by retaining a small amount of the original noise signal.

35 After subtraction, the square root of the power terms is taken by a unit 9 to provide corresponding Fourier

amplitude components, and the time domain signal segments reconstructed by an inverse Fourier transform unit 10 from these along with phase components $\phi_p(\omega)$ directly from the FFT unit 3 (via a line 11). The windowed speech segments
5 are overlapped in a unit 12 to provide the reconstructed output signal at an output 13.

As already discussed in the introduction, the spectral subtraction technique employed in the apparatus of Figure 1 has the disadvantage that the output, though less noisy
10 than the input signal, contains musical noise. The majority of information in a segment of noise-free speech is contained within one or more high energy frequency bands, known as formants. In the case of speech corrupted by white additive noise, the musical noise remaining after
15 spectral subtraction is equally likely at all frequencies. It follows that the formant regions of the frequency spectrum will have a local signal-to-noise ratio (s.n.r.) which is higher than the mean s.n.r. for the signal as a whole.

20 Within the formant regions themselves, the musical noise is largely masked out by the speech itself. Figure 2 illustrates a first embodiment of the present invention which aims to reduce the audible musical noise by attenuating the signal in the regions of the frequency
25 spectrum lying between the formant regions. Attenuation of the regions between the formants has little effect on the perceived quality of the speech itself, so that this approach is able to effect a substantial reduction in the musical noise without significantly distorting the speech.

30 This attenuation is performed by a unit 20, which multiplies the Fourier coefficients by respective terms of a frequency response $H(\omega)$ (those parts of the apparatus of Figure 2 having the same reference numerals as in Figure 1 being as already described).

35 The response $H(\omega)$ is derived from the L.P.C. (Linear Predictive Coding) spectrum $L(\omega)$ which is obtained by means

of a Linear Prediction analysis unit 21. L.P.C. analysis is a well known technique in the field of speech coding and processing and will not, therefore, be described further here. The attenuation operation is such that any
5 coefficient of the spectrally subtracted speech $\hat{P}_s(\omega)$ is attenuated only if the corresponding frequency term of the L.P.C. spectrum is below a threshold value τ . Thus the response $H(\omega)$ is a nonlinear function of $L(\omega)$ and is obtained by a nonlinear processing unit 22 according to the
10 rule:

$$\text{- if } L(\omega) \geq \tau \text{ then } H(\omega) = 1$$

$$\text{- if } L(\omega) < \tau \text{ then } H(\omega) = \left[\frac{L(\omega)}{\tau} \right]^\sigma$$

Preferably the threshold value τ is a constant for all
15 frequencies and for all speech segments; therefore in a strongly voiced segment of speech, only small portions of the spectrum will be attenuated, whereas in quiet segments most or all of the spectrum may be attenuated. A typical value of about 0.1% of the peak amplitude of the speech is
20 found to work well. A lower value of τ will produce a more harsh filtering operation. Thus the value could be increased for higher signal to noise ratios, and lowered for lower signal to noise ratios. The power term σ is used to vary the harshness of the attenuation; a larger value of
25 σ will make the attenuation more harsh. Values of σ from 2 to 4 have been found to work well in practice. Figure 3 is a graph showing the values of $H(\omega)$ for a typical L.P.C. spectrum $L(\omega)$.

As is well known, the L.P.C. analysis is very
30 sensitive to the presence of noise in the speech signal being analysed. However, the estimation of L.P.C. parameters in the presence of noise is improved by using

spectral subtraction prior to the L.P.C. analysis, and for this reason the estimator 21 in Figure 2 takes as its input the output of the subtractor 7.

When the spectral subtraction is followed by the
5 weighting function $H(\omega)$ a lower value of the scaling factor can be used (α_1 in Figures 4 and 5). A value of 1.5 for a signal to noise ratio of 10dB has been found to work well.

It has been found that a higher value of α gives better results for the auxiliary spectral subtraction
10 (α_2 in Figures 4 and 5). (A value of 2.5 has been found to work well at a signal noise ratio of 10 dB); thus in Figure 4 a separate multiplier 8¹ and subtractor stage 7¹, are used to feed the LPC spectrum estimation 21.

As the response $H(\omega)$ is applied to the amplitude
15 terms, and does not affect the phase spectrum $\phi_s(\omega)$, this attenuation is not strictly a filtering operation; though it would in principle be possible to apply filtering by $H(\omega)$ after the inverse Fourier transformation in 10. Alternatively it is also possible to apply the attenuation
20 before the square root (9).

It is noted in passing that the estimation of L.P.C. parameters is not as critical in this context as in coding or recognition applications, since a small error in the bandwidth or frequency of a pole of the filter will affect
25 the filtering only slightly; consequently L.P.C. algorithms generally considered unsuitable for noisy situations may nevertheless be of use here.

However, there are a number of further steps that can be taken to improve the accuracy of the L.P.C. estimation,
30 as will now be described with reference to Figure 4. When a segment of speech containing uncorrelated noise is analysed, the contribution of the speech component (as opposed to the noise component) to the results is enhanced by a factor dependent on the segment length. Theory
35 predicts that when the speech is entirely stationary (i.e. $P_s(\omega)$ is not changing with time) the degree of enhancement

is proportional to the square root of the segment length. Consequently it is preferable to use, for the spectral subtraction preceding the L.P.C. analysis, a longer segment length when the speech is stationary. Thus the apparatus
5 of Figure 5 includes an auxiliary spectral subtraction arrangement comprising units 2' to 8' which are identical to units 2 to 8 in all respects except for the segment length. The L.P.C. estimator 21 now takes its input from the auxiliary subtractor 7'.

10 The speech is divided into stationary sections and the segment length adjusted to match. A further unit 23 monitors the stationarity of the input speech signal and provides to the windowing unit 2' (and units 3' to 8', via connections not illustrated) a control signal CSL
15 indicating the segment length that is to be used. Tests have indicated that a typical range of segment length variation is from 38 to 205 ms.

The mode of operation of the detector 23 might be as follows:

20 (i) The LP spectrum of the central 25 ms of the present frame of noisy speech is calculated.

(ii) LP spectra of neighbouring 25 ms portions are also calculated, and spectral distances between the central LP spectrum and the neighbouring LP spectra are calculated.

25 (iii) Any neighbouring 25 ms portions judged sufficiently similar to the present portion are included in the 'stationary section'. A maximum of four 25 ms segments forward and back from the present portion are used. Thus stationary sections might range in length from 25 ms to 225
30 mS, and will not necessarily be centred around the present windowed frame.

(iv) Spectral subtraction is then performed on the stationary section as a whole, and the LP spectral estimate is calculated.

35 Additionally, it is found that L.P.C parameters derived from spectrally subtracted speech tend to move the

poles of the response - compared with the true positions that would be obtained by analysing a noise-free version of the speech - towards the unit circle (i.e. the opposite of what occurs when L.P.C. parameters are calculated directly from noisy speech). This effect can be mitigated by damping the parameters prior to calculation of the L.P.C. spectrum $L(\omega)$. Thus the L.P.C. estimation unit 21 in Figure 5 proceeds by:

(i) deriving the coefficients a_i ($1 \leq i \leq p$) of an L.P.C. filter of order p .

(ii) Damping the coefficients using the transformation

$$a_i' = a_i \cdot \sigma_i$$

where σ is a constant less than unity (e.g. 0.97).

(iii) Computing the filter response $L(\omega)$ from the damped coefficients a_i' .

Figure 6 shows graphically a comparison of the results obtained.

The first plot shows a short term spectrum of the corrupted vowel sound 'o' from the word 'hogs' after enhancement by spectral subtraction. The second plot shows the same frame of corrupted speech after spectral subtraction followed by the post processing algorithm. The peaks marked # in the first plot have been removed by the spectral weighting function in the second plot. It can be shown that these peaks are uncorrelated with the speech, and are the cause of the musical noise. Secondly, the attenuation of the lower amplitude formants is greater in the first plot, due to higher value of α , leading to more distorted speech.

A further embodiment of the invention employs spectral scaling rather than spectral subtraction. Figure 7 shows the basic principle of this, where the transformed coefficients are subjected to processing (in unit 30) by a nonlinear transfer characteristic which progressively attenuates lower intensity spectral components (assumed to

consist mainly of noise) but passes higher intensity spectral components relatively unattenuated. As described by Munday (U.S. patent No. 5,133,013) different transfer characteristics may be used for different frequency components, and/or level automatic gain control or other arrangements may be provided for scaling the nonlinear characteristic according to signal amplitude.

Spectral attenuation as envisaged by the present invention may be employed in this case also, as shown in Figure 8 where the unit 20 is inserted between the nonlinear processing 30 and the inverse FFT unit 10. As in the case of Figure 4, the response $H(\omega)$ is provided by an L.P.C. estimation unit 21 and nonlinear unit 22, which function as described above, save that the input to the spectrum estimation is now obtained from the nonlinear processing stage 30. Analogously to the case of the apparatus of Figure 4 or 5, this input may be obtained from an auxiliary spectral scaling arrangement having a different value of α and/or a different, or adaptively variable segment length.

It should be noted that the preprocessing for the L.P.C. spectrum estimation and the main spectral subtraction or scaling do not necessarily have to be of the same type; thus, if desired, the apparatus of Figure 5 could utilise spectral scaling to feed the L.P.C. analysis unit 21, or the apparatus of Figure 8 could employ spectral subtraction.



CLAIMS

1. A noise reduction apparatus comprising:
 - conversion means for converting a time-varying input signal into signals representing the magnitudes of spectral components of the input signals;
 - processing means operable to effect a reduction in the magnitude of low-magnitude ones of the said spectral component signals relative to that of higher magnitude ones of the said spectral component signals; and
 - reconversion means to convert the said spectral component signals into a time-varying signal; characterised by means to identify formant regions of the speech spectrum; and means to attenuate those frequency components lying outside the formant regions.
2. A noise reduction apparatus according to Claim 1 in which the conversion means is operable to perform a discrete Fourier transform on segments of the input signal.
3. A noise reduction apparatus according to Claim 1 or 2 including means for recognising periods during which speech is absent from the input signal and to store signals representing the power spectrum of the input signal during such periods to represent an estimated noise spectrum of the input signal and the processing means is operable to subtract, from signals representing the power spectrum of the input signal, the signals representing an estimated noise spectrum of the input signal.
4. A noise reduction apparatus according to Claim 1 or 2 in which the processing means is operable to apply to the said magnitude signals a nonlinear transfer characteristic such as to attenuate low magnitude spectral component signals relative to high magnitude ones.

This Sheet Has Been Replaced



THIS SHEET HAS BEEN
REPLACED



5. A noise reduction apparatus according to any one of claims 1 to 4 in which the means to identify formant regions is responsive to the input signal or a derivative of it to produce frequency response signals and the
5 attenuation means is operable to multiply the power spectrum of the signal by the frequency response signals.

6. A noise reduction apparatus according to Claim 5 in which the means to identify formant regions includes Linear Predictive Analysis means to produce an LP spectrum.

10 7. A noise reduction apparatus according to Claim 6 in which the means to identify formant regions includes thresholding means such that the frequency response signals are unity wherever the LP spectrum is above a threshold value and otherwise are a function of the LP spectrum.

15 8. A noise reduction apparatus according to Claim 5, 6 or 7 in which the means to identify formant regions is responsive to the output of the processing means.

9. A noise reduction apparatus according to Claim 5, 6 or 7 in which the means to identify the formant regions is
20 responsive to the spectral signals following processing by auxiliary processing means operable to effect a reduction in the magnitude of low-magnitude ones of the said spectral component signals relative to that of higher magnitude ones of the said spectral component signals.

25 10. A noise reduction apparatus according to Claim 5, 6 or 7 including auxiliary conversion means for converting the time-varying input signal into signals representing the magnitudes of spectral components of the input signals and auxiliary processing means operable to effect a reduction
30 in the magnitude of low-magnitude ones of the said spectral component signals relative to that of higher magnitude ones

This Sheet Has Been Replaced



THIS SHEET HAS BEEN
REPLACED



of the said spectral component signals; and in which the means to identify the formant regions is responsive to the output of the auxiliary processing means.

11. A noise reduction apparatus according to Claim 10 in
5 which the conversion means is operable to produce spectral
component signals for each of successive fixed time periods
of the input signal and the auxiliary conversion means is
operable to produce spectral component signals for each
10 successive time period of speech, those periods having
durations differing from the said fixed time periods.

12. A noise reduction apparatus according to Claim 11
including means for monitoring the stationarity of the
input speech signal and to control the duration of the time
periods employed by the auxiliary conversion means.

13. Noise reduction apparatus substantially as herein
15 described with reference to figures 2 to 6 and 8 of the
accompanying drawings.

This Sheet Has Been Replaced



THIS SHEET HAS BEEN
REPLACED