Description

This invention is concerned with speech coding, and more particularly to systems in which a speech signal can be generated by feeding the output of an excitation source through a synthesis filter. The coding problem then becomes one of generating, from input speech, the necessary excitation and filter parameters. LPC (linear predictive coding) parameters for the filter can be derived using well-established techniques, and the present invention is concerned with excitation source.

Systems in which a voiced/unvoiced decision on the input speech is made to switch between a noise source and a repetitive pulse source tend to give the speech output an unnatural quality, and it has been proposed to employ a single "multipulse" excitation source in which a sequence of pulses is generated, no prior assumptions being made as to the nature of the sequence. It is found that, with this method, only a few pulses (say 8 in 10ms frame) are sufficient for obtaining reasonable results. See B S Atal and J R Remde: "A New Model of LPC Excitation for producing Natural-sounding Speech at Low Bit Rates", Proc. IEEE ICASSP, Paris, pp.614, 1982.

More particularly, the present invention concerns a speech coder comprising means for deriving, from an input speech signal, parameters of a synthesis filter; means for generating a coded representation of an excitation consisting of a plurality of pulses within a time frame corresponding to a larger plurality of speech samples, being arranged in operation to select the amplitudes and timing of pulses so as to reduce the difference between the input speech signal and the response of the filter to the excitation by:

deriving the amplitude and timing of a first pulse, which alone represents an excitation tending to reduce the said difference, and successively deriving one or more further pulses which in combination with the first and any intervening pulses represent an excitation tending to reduce the said difference.

Such a coder is described in "Efficient Computation and Encoding of the Multipulse Excitation for LPC", M. Berouti et al, ICASSP' 84, IEEE International Conference on Acoustics, Speech and Signal Processing, 19th-21st March 1984, San Diego, CA, vol. 1, pages 10. 1. 1 - 10. 1. 4, IEEE New York, US.

In accordance with the present invention the coder also includes means for multiplying the pulse amplitudes by factors which depend only on their position in the derivation sequence, the factors for each pulse after the first being greater than the factor used for the first pulse and greater than or equal than the factor(s) used for any intervening pulses, and a backward adaptive quantiser for quantising the products.

Some embodiments of the invention will now be described with reference to the accompanying drawings, in which:

Figure 1 is a block diagram of one embodiment of speech coder;

Figure 2 is a block diagram of a decoder for use with the coder of Figure 1; and

Figure 3 is a block diagram of a second embodiment of coder.

In the coder of Figure 1, input speech signals, in sampled (preferably digital) form at an input 1 are processed by a predictor 2 to produce an output (e.g. in the form of a set of filter coefficients) defining a synthesis filter having a spectral response akin to that the of the speech signals. The predictor analysis can be any of those conventionally used in so-called LPC (linear predictive coding) speech coders. As is common in such systems, the analysis is performed on frames of speech into which the input samples are divided. Typically the frame length may be 20ms; hence a set of coefficients is produced every 20ms and supplied via lines 3 to an output multiplexer 4.

As well as the filter representation, the coder also produces a representation of an excitation which is to be generated at the decoder to drive the synthesis filter in order to produce an approximation to the original speech. The coder of Figure 1 has a multipulse derivation unit 5 which derives from the input speech samples and the LPC coefficients the amplitudes (on output 6) and positions (on output 7) of the pulses in a "multipulse" excitation frame as mentioned above. Whilst the typical sub-block (i.e. portion of LPC frame) size of 10ms with eight pulses may be employed, the embodiment of Figure 1 employs a sub-block duration of 4ms, with three pulses. This is preferred as introducing less delay into the coding process. The object of the multipulse derivation is to find the pulse positions and amplitudes which minimise the error between the decoded synthetic speech and the original speech.

If it is assumed that a sub-block consists of n speech samples, this represents n input speech samples $s_0..s_{n-1}$ and n synthesised samples $s'_{0...}s'_{n-1}$, which can be regarded as vectors s, s' . The excitation consists of pulses of amplitude a_m which are, it is assumed, permitted to occur at any of the n possible time instants within the frame, but there are only a limited number of them (say k). Thus the excitation can be expressed as an n-dimensional vector a with components $a_0....a_{n-1}$, but only k of them are non-zero. The objective is to find the 2k unknowns (k amplitudes, k pulse positions) which minimise the error:

$$e^2 = (s - s')^2 \quad (1)$$

The amount of computation required to do this is considerable and the procedure proposed by Atal and Remde was as follows:

(1) Find the amplitude and position of one pulse, alone, to give a minimum error.

(2) Find the amplitude and position of a second pulse which, in combination with this first pulse, give a minimum error; the positions and amplitudes of the pulse(s) previously found are fixed during this stage.

(3) Repeat for further pulses.

This method is employed in derivation unit 5 of Figure 1; that the earlier derived pulses are taken into account in the later derivations within a sub-block is indicated in Figure 1 by feedback paths 8, 9. Note that the sequence in which the pulses are derived is not related to their actual position within the sub-block.

The pulse amplitudes a_i are passed via a backward-adaptive quantiser 10, described below. First however they are multiplied (in a multiplier 11) by a statistical factor f_i. In practice it is found that the first pulse to be derived is generally the largest, and successively derived pulses tend to be progressively smaller, at least for the first few pulses. Although the pulse sizes vary, a statistical analysis on training sequences shows that on average this is so, and the multiplier is supplied with factors such that on average the pulse amplitudes at the multiplier output tend to be the same irrespective of which pulse in the derivation sequence it is. For the case considered here of three pulses, the factors employed are:

first pulse to be derived $f_0 = 1$

second pulse to be derived $f_1 = 8/5$

third pulse to be derived $f_2 = 8/3$

(the fourth to sixth pulses, if present, may be given the factors 8/3, 8/3 and 4). The object of this step is to make the adaptive quantisation more efficient and enable either the quantisation noise or the number of bits used to encode the amplitude (or both) to be reduced.

Where larger numbers of pulses are used, suitable factors can be derived by analysis of sample sequences of speech to find the average magnitudes of the pulses compared with that of the first derived pulse. The multiplication factor is then the reciprocal of this. A simple (albeit non-optimum) approach for such a situation is to use a factor of unity for the first derived pulse, and 2 for the remainder.

The adaptive quantiser 9 is a 3-bit Jayant quantiser and has an optimum non-linear Max quantiser 12 having the following characteristic:

TABLE 1

INPUT RANGE	OUTPUT	OUTPUT CODE
below-1. 748	-2. 152	1/4
-1. 748 to -1. 5	-1. 344	1/3
-1. 5 to 0. 50006	-0. 7560	1/2
-0. 50006 to 0	-0. 2451	1/1
0 to 0. 50006	0. 2451	0/1
0. 50006 to 1. 5	0. 7560	0/2
1. 5 to 1. 748	1. 344	0/3
above 1. 748	2. 152	0/4

The output code simply represents the values of the three output bits - the number before the "/" is the sign bit and the number 1.....4 following signifies the binary number 0....11.

A scaling unit 13 provides a scale factor to a divider 14 at the quantiser input. The scale factor S (initially unity) is varied in that, depending on the quantiser codeword output for a given pulse amplitude value, the scale factor S is increased or decreased from its current value to a new value to be used for the next pulse amplitude.

$$S_k = S_{k-1} \cdot m_{k-1}$$

Where m is given by:

Table 2

output code	m
1	0. 875
2	0. 875
3	1. 000
4	1. 500

Note that these factors are different from those proposed by Jayant; also that the scale factor is not reset at the end of a sub-block or frame.

An additional feature that may be employed for speeding up adaptation is that, if two consecutive output codes have the value 4, then the second occurrence results in an increase of scale factor by a factor of 2.25 (i.e. two increases of 1.5). This is illustrated in frame 1 by a delay 15 and 4,4 detector 16.

The output multiplexer receives the quantised amplitudes from the quantiser 10 and the position information from the derivation unit 5, as well as the LPC coefficients and combines these into a single output 17.

A decoder is shown in Figure 2, where a demultiplexer 24 separates the coefficients, amplitudes and position information and feeds the coefficients to update a synthesis filter 30. The pulse amplitude codewords are passed via an "inverse quantiser" 22 which removes the non-linearity introduced by the quantiser 12 - i.e. it converts the received codewords into the values given in the middle column of table 1. The scaling factor S is obtained from the amplitude codewords by units 23, 25, 26 in all respects identical to units 13, 15, 16 of Figure 1 and the inverse quantiser output is multiplied by S in a multiplier 31. The factors f_i are then applied to a divider 32 whose output represents the original amplitudes (but with quantisation error) and is supplied along with the pulse position information to an excitation generator 33.

The output of the excitation generator 33 is filtered by the filter 30 to produce decoded speech at an output 34.

It has already been mentioned that the multipulse derivation unit takes account, in the later pulse derivations, of the effect of the earlier derived pulses, via the feedback paths 8, 9. It is preferable to take account of the actual effect of these pulses at the decoder and therefore the quantisation is preferably included within this loop. Thus, in the modified coder shown in Figure 3, the pulse amplitudes are fed back from the output via a local decoder 40 which has an inverse quantiser 22', multiplier 31' and divider 32'. The scale factor can be obtained from the quantiser 10, of course. The decoder is Figure 2 may again be used with this coder.

Some multipulse coding schemes involving sequential pulse derivation involve re-optimisation steps. This is because the earlier derived pulses are derived without reference to the nature of those derived later, and the results can be improved by applying a correction to the amplitudes and/or positions of the pulses. See, for example our UK patent applications nos. 8608031 (Patent no. 2173679B) and 8720604 (Patent no. 2195220B).

In the case of Figure 1, any of these techniques may be applied as in the past. In the case of Figure 2, position re-optimisation may be used, if desired. However, in Figure 3, where in-loop quantisation of pulse i is carried out before pulse i+1 is derived, and further adjustment of pulse i may not then be possible without seriously affecting the quantisation process.

Claims

1. A speech coder comprising:

means (2) for deriving, from an input speech signal, parameters of a synthesis filter;

means (5) for generating a coded representation of an excitation consisting of a plurality of pulses within a time frame corresponding to a larger plurality of speech samples, being arranged in operation to select the amplitudes and timing of pulses so as to reduce the difference between the input speech signal and the response of the filter to the excitation by:

deriving the amplitude and timing of a first pulse, which alone represents an excitation tending to reduce the said difference, and successively deriving one or more further pulses which in combination with the first and any intervening pulses represent an excitation tending to reduce the said difference;

characterised by means (11) for multiplying the pulse amplitudes by factors (f_i)which depend only on their position in the derivation sequence, the factors for each pulse after the first being greater than the

* C L M S *

Table 2

output code	m
1	0. 875
2	0. 875
3	1. 000
4	1. 500

Note that these factors are different from those proposed by Jayant; also that the scale factor is not reset at the end of a sub-block or frame.

An additional feature that may be employed for speeding up adaptation is that, if two consecutive output codes have the value 4, then the second occurrence results in an increase of scale factor by a factor of 2.25 (i.e. two increases of 1.5). This is illustrated in frame 1 by a delay 15 and 4,4 detector 16.

The output multiplexer receives the quantised amplitudes from the quantiser 10 and the position information from the derivation unit 5, as well as the LPC coefficients and combines these into a single output 17.

A decoder is shown in Figure 2, where a demultiplexer 24 separates the coefficients, amplitudes and position information and feeds the coefficients to update a synthesis filter 30. The pulse amplitude codewords are passed via an "inverse quantiser" 22 which removes the non-linearity introduced by the quantiser 12 - i.e. it converts the received codewords into the values given in the middle column of table 1. The scaling factor S is obtained from the amplitude codewords by units 23, 25, 26 in all respects identical to units 13, 15, 16 of Figure 1 and the inverse quantiser output is multiplied by S in a multiplier 31. The factors f_i are then applied to a divider 32 whose output represents the original amplitudes (but with quantisation error) and is supplied along with the pulse position information to an excitation generator 33.

The output of the excitation generator 33 is filtered by the filter 30 to produce decoded speech at an output 34.

It has already been mentioned that the multipulse derivation unit takes account, in the later pulse derivations, of the effect of the earlier derived pulses, via the feedback paths 8, 9. It is preferable to take account of the actual effect of these pulses at the decoder and therefore the quantisation is preferably included within this loop. Thus, in the modified coder shown in Figure 3, the pulse amplitudes are fed back from the output via a local decoder 40 which has an inverse quantiser 22', multiplier 31' and divider 32'. The scale factor can be obtained from the quantiser 10, of course. The decoder is Figure 2 may again be used with this coder.

Some multipulse coding schemes involving sequential pulse derivation involve re-optimisation steps. This is because the earlier derived pulses are derived without reference to the nature of those derived later, and the results can be improved by applying a correction to the amplitudes and/or positions of the pulses. See, for example our UK patent applications nos. 8608031 (Patent no. 2173679B) and 8720604 (Patent no. 2195220B).

In the case of Figure 1, any of these techniques may be applied as in the past. In the case of Figure 2, position re-optimisation may be used, if desired. However, in Figure 3, where in-loop quantisation of pulse i is carried out before pulse i+1 is derived, and further adjustment of pulse i may not then be possible without seriously affecting the quantisation process.

Claims

1. A speech coder comprising:

means (2) for deriving, from an input speech signal, parameters of a synthesis filter;

means (5) for generating a coded representation of an excitation consisting of a plurality of pulses within a time frame corresponding to a larger plurality of speech samples, being arranged in operation to select the amplitudes and timing of pulses so as to reduce the difference between the input speech signal and the response of the filter to the excitation by:

deriving the amplitude and timing of a first pulse, which alone represents an excitation tending to reduce the said difference, and successively deriving one or more further pulses which in combination with the first and any intervening pulses represent an excitation tending to reduce the said difference;

characterised by means (11) for multiplying the pulse amplitudes by factors (f_i)which depend only on their position in the derivation sequence, the factors for each pulse after the first being greater than the

* C L M S *

factor used for the first pulse and greater than or equal to the factor(s) used for any intervening pulses, and a backward adaptive quantiser (10) for quantising the products.

2. A speech coder according to Claim 1 in which the factor is unity for the first pulse.

3. A speech coder according to Claim 1 or 2 in which at least three pulses are derived.

4. A speech coder according to Claim 3 in which the factors for the first three pulses in order of derivation are substantially 1, 8/5 and 8/3.

5. A speech coder according to any one of he preceding claims in which the deriving means (5) are arranged when deriving the further pulse(s), to employ the values of the amplitudes of the first and any intervening pulses obtained from the quantiser output via a local decoder (40).

Patentansprüche

1. Sprachkodierer, der aufweist:
eine Einrichtung (2) zum Ableiten von Parametern eines Synthesefilters von einem eingegebenen Sprachsignal;
eine Einrichtung (5) zum Erzeugung einer kodierten Darstellung einer Angregung, die aus einer Vielzahl von Impulsen innerhalb eines Zeitrahmens besteht, der einer größeren Vielzahl von Sprachabtastungen entspricht, wobei die Einrichtung im Betrieb eingerichtet ist, um die Amplituden und die Zeitgabe der Impulse auszuwählen, um die Amplituden und die Zeitgabe der Impulse auszuwählen, um die Differenz zwischen dem eingegebenen Sprachsignal und der Antwort des Filters auf die Anregung zu reduzieren, und zwar durch:
Ableiten der Amplitude und der Zeitgabe eines ersten Impulses, der allein eine Anregung darstellt, die dazu neigt, die Differenz zu reduzieren, und aufeinanderfolgendes Ableiten eines oder mehrerer weiterer Impulse, die in Kombination mit dem ersten und irgendwelchen dazwischenliegenden Impulsen eine Anregung darstellen, die dazu neigt, die Differenz zu reduzieren;
gekennzeichnet durch
eine Einrichtung (11) zum Vervielfachen der Impulsamplituden durch Faktoren (f_i), die nur von ihrer Position in der Ableitungsfolge abhängen, wobei die Faktoren für jeden Impuls nach dem ersten größer sind als der Faktor, der für den ersten Impuls benutzt wird, und größer als oder gleich dem Faktor (den Faktoren) ist (sind), der (die) für irgendwelche dazwischenliegenden Impulse benutzt wird (werden), und einen rückwärts adaptiven Quantisierer (10) zum Quantisieren der Produkte.

2. Sprachkodierer nach Anspruch 1, wobei der Faktor eine Einheit für den ersten Impuls ist.

3. Sprachkodierer nach Anspruch 1 oder 2, wobei zumindest drei Impulse abgeleitet werden.

4. Sprachkodierer nach Anspruch 3, wobei die Faktoren für die ersten drei Impulse in einer Reihenfolge der Ableitung im wesentlichen 1, 8/5 und 8/3 sind.

5. Sprachkodierer nach einem der vorangehenden Ansprüche, wobei die Ableitungseinrichtung (5) beim Ableiten des (der) weiteren Impulses (Impulse) eingerichtet sind, die Werte der Amplituden des ersten und irgendwelcher dazwischenliegender Impulse zu verwenden, die von dem Quantisiererausgang über einen lokalen Dekodierer (40) erhalten werden.

Revendications

1. Un codeur de parole comprenant:
un moyen (2) de dérivation de paramètres d'un filtre de synthèse à partir d'un signal de parole d'entrée;
un moyen (5) de génération d'une représentation codée d'une excitation consistant en une série d'impulsions situées à l'intérieur d'une base de temps correspondant à une série plus importante d'échantillons de parole, agencé pour choisir en fonctionnement les amplitudes et la synchronisation d'impulsions de façon à réduire la différence entre le signal de parole d'entrée et la réponse du filtre à l'excitation:

en dérivant l'amplitude et la synchronisation d'une première impulsion, qui représente seule une excitation tendant à réduire ladite différence, et en dérivant successivement une ou plusieurs impulsions additionnelles qui représentent, en combinaison avec la première impulsion et toute impulsion intermédiaire, une excitation tendant à réduire ladite différence;

caractérisé par un moyen (11) de multiplication des amplitudes d'impulsions par des facteurs (f_i) qui ne dépendent que de leur position dans la séquence de dérivation, les facteurs de chaque impulsion suivant la première étant supérieurs aux facteurs utilisés pour la première impulsion et supérieurs ou égaux au(x) facteur(s) utilisé(s) pour toutes impulsions intermédiaires, et un quantificateur adaptatif en rétroaction (10) pour quantifier les produits.

2. Un codeur de parole selon la revendication 1 dans lequel le facteur est l'unité pour la première impulsion.

3. Un codeur de parole selon la revendication 1 ou 2 dans lequel au moins trois impulsions sont dérivées.

4. Un codeur de parole selon la revendication 3 dans lequel les facteurs des trois premières impulsions dans l'ordre de dérivation sont sensiblement 1, 8/5 et 8/3.

5. Un codeur de parole selon l'une quelconque des revendications précédentes dans lequel les moyens de dérivation (5) sont agencés pour employer, lorsqu'ils dérivent l'impulsion ou les impulsion additionnelle(s), les valeurs des amplitudes de la première impulsion et des autres impulsions intermédiaires obtenues a partir de la sortie du quantifieur au moyen d'un décodeur local (40).