

REDUCING SPARSENESS IN CODED SPEECH SIGNALS

This application claims the priority under 35 USC 119(e)(1) of copending U.S. Provisional Application No. 06/057,752, filed on September 2, 1997, and is a continuation -in-part of copending U.S. Serial No. 09/034,590 (docket 34645-405),
5 filed on March 4, 1998.

FIELD OF THE INVENTION

The invention relates generally to speech coding and, more particularly, to the problem of sparseness in coded speech signals.

BACKGROUND OF THE INVENTION

10 Speech coding is an important part of modern digital communications systems, for example, wireless radio communications systems such as digital cellular telecommunications systems. To achieve the high capacity required by such systems both today and in the future, it is imperative to provide efficient compression of speech signals while also providing high quality speech signals. In this connection, when the
15 bit rate of a speech coder is decreased, for example to provide additional communication channel capacity for other communications signals, it is desirable to obtain a graceful degradation of speech quality without introducing annoying artifacts.

Conventional examples of lower rate speech coders for cellular telecommunications are illustrated in IS-641 (D-AMPS EFR) and by the G.729 ITU
20 standard. The coders specified in the foregoing standards are similar in structure, both including an algebraic codebook that typically provides a relatively sparse output. Sparseness refers in general to the situation wherein only a few of the samples of a given codebook entry have a non-zero sample value. This sparseness condition is particularly prevalent when the bit rate of the algebraic codebook is reduced in an
25 attempt to provide speech compression. With very few non-zero samples in the codebook to begin with, and with the lower bit rate requiring that even fewer codebook samples be used, the resulting sparseness is an easily perceived degradation in the coded speech signals of the aforementioned conventional speech coders.

-2-

It is therefore desirable to avoid the aforementioned degradation in coded speech signals when the bit rate of a speech coder is reduced to provide speech compression.

5 In an attempt to avoid the aforementioned degradation in coded speech signals, the present invention provides an anti-sparseness operator for reducing the sparseness in a coded speech signal, or any digital signal, wherein sparseness is disadvantageous.

BRIEF DESCRIPTION OF THE DRAWINGS

FIGURE 1 is a block diagram which illustrates one example of an anti-sparseness operator of the present invention.

10 FIGURE 2 illustrates various positions in a Code Excited Linear Predictive encoder/decoder where the anti-sparseness operator of FIGURE 1 can be applied.

FIGURE 2A illustrates a communications transceiver that can use the encoder/decoder structure of FIGURES 2 and 2B.

15 FIGURE 2B illustrates another exemplary Code Excited Linear Predictive decoder including the anti-sparseness operator of FIGURE 1.

FIGURE 3 illustrates one example of the anti-sparseness operator of FIGURE 1.

FIGURE 4 illustrates one example of how the additive signal of FIGURE 3 can be produced.

20 FIGURE 5 illustrates in block diagram form how the anti-sparseness operator of FIGURE 1 can be embodied as an anti-sparseness filter.

FIGURE 6 illustrates one example of the anti-sparseness filter of FIGURE 5.

FIGURES 7-11 illustrate graphically the operation of an anti-sparseness filter of the type illustrated in FIGURE 6.

25 FIGURES 12-16 illustrate graphically the operation of an anti-sparseness filter of the type illustrated in FIGURE 6 and at a relatively lower level of anti-sparseness operation than the anti-sparseness filter of FIGURES 7-11.

FIGURE 17 illustrates another example of the anti-sparseness operator of FIGURE 1.

30 FIGURE 18 illustrates an exemplary method of providing anti-sparseness modification according to the invention.

-3-

DETAILED DESCRIPTION

FIGURE 1 illustrates an example of an anti-sparseness operator according to the present invention. The anti-sparseness operator ASO of FIGURE 1 receives at input A thereof a sparse, digital signal received from a source 11. The anti-sparseness operator ASO operates on the sparse signal A and provides at an output thereof a digital signal B which is less sparse than the input signal A.

FIGURE 2 illustrates various example locations where the anti-sparseness operator ASO of FIGURE 1 can be applied in a Code Excited Linear Predictive (CELP) speech encoder provided in a transmitter for use in a wireless communication system, or in a CELP speech decoder provided in a receiver of a wireless communication system. As shown in FIGURE 2, the anti-sparseness operator ASO can be provided at the output of the fixed (e.g, algebraic) codebook 21, and/or at any of the locations designated by reference numerals 201-206. At each of the locations designated in FIGURE 2, the anti-sparseness operator ASO of FIGURE 1 would receive at its input A the sparse signal and provide at its output B a less sparse signal. Thus, the CELP speech encoder/decoder structure shown in FIGURE 2 includes several examples of the sparse signal source of FIGURE 1.

The broken line in FIGURE 2 illustrates the conventional feedback path to the adaptive codebook as conventionally provided in CELP speech encoders/decoders. If the anti-sparseness operator ASO is provided where shown in FIGURE 2 and/or at any of locations 201-204, then the anti-sparseness operator(s) will affect the coded excitation signal reconstructed by the decoder at the output of summing circuit 210. If applied at locations 205 and/or 206, the anti-sparseness operator(s) will have no effect on the coded excitation signal output from summing circuit 210.

FIGURE 2B illustrates an example CELP decoder including a further summing circuit 25 which receives the outputs of codebooks 21 and 23, and provides the feedback signal to the adaptive codebook 23. If the anti-sparseness operator ASO is provided where shown in FIGURE 2B, and/or at locations 220 and 240, then such anti-sparseness operator(s) will not affect the feedback signal to the adaptive codebook 23.

FIGURE 2A illustrates a transceiver whose receiver (RCVR) includes the CELP decoder structure of FIGURE 2 (or FIGURE 2B) and whose transmitter

-4-

(XMTR) includes the CELP encoder structure of FIGURE 2. FIGURE 2A illustrates that the transmitter receives as input an acoustical signal and provides as output to the communications channel reconstruction information from which a receiver can reconstruct the acoustical signal. The receiver receives as input from the communications channel reconstruction information, and provides a reconstructed acoustical signal as an output. The illustrated transceiver and communications channel could be, for example, a transceiver in a cellular telephone and the air interface of a cellular telephone network, respectively.

FIGURE 3 illustrates one example implementation of the anti-sparseness operator ASO of FIGURE 1. In FIGURE 3, a noise-like signal $m(n)$ is added to the sparse signal as received at A. FIGURE 4 illustrates one example of how the signal $m(n)$ can be produced. A noise signal with a Gaussian distribution $N(0,1)$ is filtered by a suitable high pass and spectral coloring filter to produce the noise-like signal $m(n)$.

As illustrated in FIGURE 3, the signal $m(n)$ can be applied to the summing circuit 31 with a suitable gain factor via multiplier 33. The gain factor of FIGURE 3 can be a fixed gain factor. The gain factor of FIGURE 3 can also be a function of the gain conventionally applied to the output of adaptive codebook 23 (or a similar parameter describing the amount of periodicity). In one example, the FIGURE 3 gain would be 0 if the adaptive codebook gain exceeds a predetermined threshold, and linearly increasing as the adaptive codebook gain decreases from the threshold. The FIGURE 3 gain can also be analogously implemented as a function of the gain conventionally applied to the output of the fixed codebook 21 of FIGURE 2. The FIGURE 3 gain can also be based on power-spectrum matching of the signal $m(n)$ to the target signal used in the conventional search method, in which case the gain needs to be encoded and transmitted to the receiver.

In another example, the addition of a noise-like signal can be performed in the frequency domain in order to obtain the benefit of advanced frequency domain analysis.

FIGURE 5 illustrates another example implementation of the ASO of FIGURE 2. The arrangement of FIGURE 5 can be characterized as an anti-sparseness filter

-5-

designed to reduce sparseness in the digital signal received from the source 11 of FIGURE 1.

One example of the anti sparseness filter of FIGURE 5 is illustrated in more detail in FIGURE 6. The anti-sparseness filter of FIGURE 6 includes a convolver
5 section 63 that performs a convolution of the coded signal received from the fixed (e.g. algebraic) codebook 21 with an impulse response (at 65) associated with an all-pass filter. The operation of one example of the FIGURE 6 anti-sparseness filter is illustrated in FIGURES 7-11.

FIGURE 10 illustrates an example of an entry from the codebook 21 of
10 FIGURE 2 having only two non-zero samples out of a total of forty samples. This sparseness characteristic will be reduced if the number (density) of non-zero samples can be increased. One way to increase the number of non-zero samples is to apply the codebook entry of FIGURE 10 to a filter having a suitable characteristic to disperse the energy throughout the block of forty samples. FIGURES 7 and 8 respectively
15 illustrate the magnitude and phase (in radians) characteristics of an all-pass filter which is operable to appropriately disperse the energy throughout the forty samples of the FIGURE 10 codebook entry. The filter of FIGURES 7 and 8 alters the phase spectrum in the high frequency area between 2 and 4 kHz, while altering the low frequency areas below 2 kHz only very marginally. The magnitude spectrum remains
20 essentially unaltered by the filter of FIGURES 7 and 8.

Example FIGURE 9 illustrates graphically the impulse response of the all-pass filter defined by FIGURES 7 and 8. The anti-sparseness filter of FIGURE 6 produces a convolution of the FIGURE 9 impulse response on the FIGURE 10 block of samples. Because the codebook entries are provided from the codebook as blocks of
25 forty samples, the convolution operation is performed in blockwise fashion. Each sample in FIGURE 10 will produce 40 intermediate multiplication results in the convolution operation. Taking the sample at position 7 in FIGURE 10 as an example, the first 34 multiplication results are assigned to positions 7-40 of the FIGURE 11 result block, and the remaining 6 multiplication results are "wrapped around"
30 according to a circular convolution operation such that they are assigned to positions 1-6 of the result block. The 40 intermediate multiplication results produced by each of the remaining FIGURE 10 samples are assigned to positions in the FIGURE 11

-6-

result block in analogous fashion, and sample 1 of course needs no wrap around. For each position in the result block of FIGURE 11, the 40 intermediate multiplication results assigned thereto (one multiplication result per sample in FIGURE 10) are summed together, and that sum represents the convolution result for that position.

5 It is clear from inspection of FIGURES 10 and 11 that the circular convolution operation alters the Fourier spectrum of the FIGURE 10 block so that the energy is dispersed throughout the block, thereby dramatically increasing the number (or density) of non-zero samples in the block, and correspondingly reducing the amount of sparseness. The effects of performing the circular convolution on a block-by-block basis can be smoothed out by the synthesis filter 211 of FIGURE 2.

10 FIGURES 12-16 illustrate another example of the operation of an anti-sparseness filter of the type shown generally in FIGURE 6. The all-pass filter of FIGURES 12 and 13 alters the phase spectrum between 3 and 4 kHz without substantially altering the phase spectrum below 3 kHz. The impulse response of the filter is shown in FIGURE 14. Referencing the result block of FIGURE 16, and noting that FIGURE 15 illustrates the same block of samples as FIGURE 10, it is clear that the anti-sparseness operation illustrated in FIGURES 12-16 does not disperse the energy as much as shown in FIGURE 11. Thus, FIGURES 12-16 define an anti-sparseness filter which modifies the codebook entry less than the filter defined by

15 FIGURES 7-11. Accordingly, the filters of FIGURES 7-11 and FIGURES 12-16 define respectively different levels of anti-sparseness filtering.

20 A low adaptive codebook gain value indicates that the adaptive codebook component of the reconstructed excitation signal (output from adder circuit 210) will be relatively small, thus giving rise to the possibility of a relatively large contribution from the fixed (e.g. algebraic) codebook 21. Because of the aforementioned sparseness of the fixed codebook entries, it would be advantageous to select the anti-sparseness filter of FIGURES 7-11 rather than that of FIGURES 12-16 because the filter of FIGURES 7-11 provides a greater modification of the sample block than does the filter of FIGURES 12-16. With larger values of adaptive codebook gain, the fixed codebook contribution is relatively less, so the filter of FIGURES 12-16 which

25 provides less anti-sparseness modification could be used.

30

-7-

The present invention thus provides the capability of using the local characteristics of a given speech segment to determine whether and how much to modify the sparseness characteristic associated with that segment.

5 The convolution performed in the FIGURE 6 anti-sparseness filter can also be linear convolution, which provides smoother operation because blockwise processing effects are avoided. Moreover, although blockwise processing is described in the above examples, such blockwise processing is not required to practice the invention, but rather is merely a characteristic of the conventional CELP speech encoder/decoder structure shown in the examples.

10 A closed-loop version of the method can be used. In this case, the encoder takes the anti-sparseness modification into account during search of the codebooks. This will give improved performance at the price of increased complexity. The (circular or linear) convolution operation can be implemented by multiplying the filtering matrix constructed from the conventional impulse response of the search filter
15 by a matrix which defines the anti-sparseness filter (using either linear or circular convolution).

FIGURE 17 illustrates another example of the anti-sparseness operator ASO of FIGURE 1. In the example of FIGURE 17, an anti-sparseness filter of the type illustrated in FIGURE 5 receives input signal A, and the output of the anti-sparseness
20 filter is multiplied at 170 by a gain factor g_2 . The noise-like signal $m(n)$ from FIGURES 3 and 4 is multiplied at 172 by a gain factor g_1 , and the outputs of the g_1 and g_2 multipliers 170 and 172 are added together at 174 to produce output signal B. The gain factors g_1 and g_2 can be determined, for example, as follows. The gain g_1 can first be determined in one of the ways described above with respect to the gain of FIGURE
25 3, and then the gain factor g_2 can be determined as a function of gain factor g_1 . For example, gain factor g_2 can vary inversely with gain factor g_1 . Alternatively, the gain factor g_2 can be determined in the same manner as the gain of FIGURE 3, and then the gain factor g_1 can be determined as a function of gain factor g_2 , for example g_1 can vary inversely with g_2 .

30 In one example of the FIGURE 17 arrangement: the anti-sparseness filter of FIGURES 12-16 is used; gain factor $g_2 = 1$; $m(n)$ is obtained by normalizing the Gaussian noise distribution $N(0,1)$ of FIGURE 4 to have an energy level equal to the

-8-

fixed codebook entries, and setting the cutoff frequency of the FIGURE 4 high pass filter at 200 Hz; and gain factor g_1 is 80% of the fixed codebook gain.

FIGURE 18 illustrates an exemplary method of providing anti-sparseness modification according to the invention. At 181, the level of sparseness of the coded speech signal is estimated. This can be done off-line or adaptively during speech processing. For example, in algebraic codebooks and multi-pulse codebooks the samples may be close to each other or far apart, resulting in varying sparseness; whereas in a regular pulse codebook, the distance between samples is fixed, so the sparseness is constant. At 183, a suitable level of anti-sparseness modification is determined. This step can also be performed off-line or adaptively during speech processing as described above. As another example of adaptively determining the anti-sparseness level, the impulse response (see FIGURES 6, 9 and 14) can be changed from block to block. At 185, the selected level of anti-sparseness modification is applied to the signal.

It will be evident to workers in the art that the embodiments described above with respect to FIGURES 1-18 can be readily implemented using, for example, a suitably programmed digital signal processor or other data processor, and can alternatively be implemented using, for example, such suitably programmed digital signal processor or other data processor in combination with additional external circuitry connected thereto.

Although exemplary embodiments of the present invention have been described above in detail, this does not limit the scope of the invention, which can be practiced in a variety of embodiments.

WHAT IS CLAIMED IS:

1. An apparatus for reducing sparseness in an input digital signal which includes a first sequence of sample values, comprising:
 - an input to receive the input digital signal;
 - 5 an anti-sparseness operator coupled to said input and responsive to the input digital signal for producing an output digital signal which includes a further sequence of sample values, said further sequence of sample values having a greater density of non-zero sample values than the first sequence of sample values; and
 - an output coupled to said anti-sparseness operator to receive therefrom
 - 10 said output digital signal.
2. The apparatus of Claim 1, wherein said anti-sparseness operator includes a circuit for adding to the input digital signal a noise-like signal.
3. The apparatus of Claim 1, wherein said anti-sparseness operator includes a filter coupled to said input to filter the input digital signal.
- 15 4. The apparatus of Claim 3, wherein said filter is an all-pass filter.
5. The apparatus of Claim 3, wherein said filter uses one of circular convolution and linear convolution to filter respective blocks of sample values in said first sequence of sample values.
- 20 6. The apparatus of Claim 3, wherein said filter modifies a phase spectrum of said input digital signal but leaves a magnitude spectrum thereof substantially unaltered.
7. The apparatus of Claim 1, wherein said anti-sparseness operator includes a signal path extending from said input to said output, said signal path including a filter, and said anti-sparseness operator also including a circuit for adding
- 25 a noise-like signal to a signal carried by said signal path.



-10-

8. The apparatus of Claim 7, wherein said filter is an all-pass filter.

9. The apparatus of Claim 7, wherein said filter uses one of circular convolution and linear convolution to filter respective blocks of sample values in the first sequence of sample values.

5 10. The apparatus of Claim 7, wherein said filter modifies a phase spectrum of the input digital signal but leaves a magnitude spectrum thereof substantially unaltered.

11. An apparatus for processing acoustical signal information, comprising:
an input for receiving the acoustical signal information;
10 a coding apparatus coupled to said input and responsive to said information for providing a digital signal, said digital signal including a first sequence of sample values; and
an anti-sparseness operator having an input coupled to said coding apparatus and responsive to said digital signal for producing an output digital signal
15 which includes a second sequence of sample values, said second sequence of sample values having a greater density of non-zero sample values than the first sequence of sample values.

12. The apparatus of Claim 11, wherein said coding apparatus includes a plurality of codebooks, a summing circuit and a synthesis filter, said codebooks having
20 respective outputs coupled to respective inputs of said summing circuit, and said summing circuit having an output coupled to an input of said synthesis filter.

13. The apparatus of Claim 12, wherein said anti-sparseness operator input is coupled to one of said codebook outputs.

14. The apparatus of Claim 12, wherein said anti-sparseness operator input
25 is coupled to said output of said summing circuit.



-11-

15. The apparatus of Claim 12, wherein said anti-sparseness operator input is coupled to an output of said synthesis filter.

16. The apparatus of Claim 12, wherein said coding apparatus is an encoding apparatus and the acoustical signal information includes an acoustical signal.

5 17. The apparatus of Claim 12, wherein said coding apparatus is a decoding apparatus and the acoustical signal information includes information from which an acoustical signal is to be constructed.

18. A method of reducing sparseness in an input digital signal which includes a first sequence of sample values, comprising:
10 receiving the input digital signal;
producing in response to the input digital signal an output digital signal which includes a second sequence of sample values, said second sequence of sample values having a greater density of non-zero sample values than the first sequence of sample values; and
15 outputting the output digital signal.

19. The method of Claim 18, wherein said producing step includes filtering the input digital signal.

20. The method of Claim 19, wherein said filtering step includes using an all-pass filter.

20 21. The method of Claim 19, wherein said filtering step includes using one of circular convolution and linear convolution to filter respective blocks of sample values in the first sequence of sample values.

22. The method of Claim 19, wherein said filtering step includes modifying a phase spectrum of the input digital signal but leaving the magnitude spectrum thereof
25 substantially unaltered.



-12-

23. The method of Claim 18, wherein said producing step includes filtering a first signal to obtain a filtered signal, and adding a noise-like signal to one of said first signal and said filtered signal.

5 24. The method of Claim 23, wherein said filtering step includes using an all-pass filter.

25. The method of Claim 23, wherein said filtering step includes using one of circular convolution and linear convolution to filter respective blocks of sample values in the first sequence of sample values.

10 26. The method of Claim 23, wherein said filtering step includes modifying a phase spectrum of the input digital signal but leaving a magnitude spectrum thereof substantially unaltered.

27. The method of Claim 18, wherein said producing step includes adding a noise-like signal to the input digital signal.

15 28. A method of processing acoustical signal information, comprising:
receiving the acoustical signal information;
providing in response to the information a digital signal including a first sequence of sample values; and
producing in response to the digital signal an output digital signal which includes a further sequence of sample values, the further sequence of sample
20 values having a greater density of non-zero sample values than the first sequence of sample values.

