

TIME-WARPING FRAMES OF WIDEBAND VOCODER

BACKGROUND

Field

[0001] This invention generally relates to time-warping, i.e., expanding or compressing, frames in a vocoder and, in particular, to methods of time-warping frames in a wideband vocoder.

Background

[0002] Time-warping has a number of applications in packet-switched networks where vocoder packets may arrive asynchronously. While time-warping may be performed either inside or outside the vocoder, performing it in the vocoder offers a number of advantages such as better quality of warped frames and reduced computational load.

SUMMARY

[0003] The invention comprises an apparatus and method of time-warping speech frames by manipulating a speech signal. In one aspect, a method of time-warping Code-Excited Linear Prediction (CELP) and Noise-Excited Linear Prediction (NELP) frames of a Fourth Generation Vocoder (4GV) wideband vocoder is disclosed. More specifically, for CELP frames, the method maintains a speech phase by adding or deleting pitch periods to expand or compress speech, respectively. With this method, the lower band signal may be time-warped in the residual, i.e., before synthesis, while the upper band signal may be time-warped after synthesis in the 8 kHz domain. The method disclosed may be applied to any wideband vocoder that uses CELP and/or NELP for the low band and/or uses a split-band technique to encode the lower and upper bands separately. It should be noted that the standards name for 4GV wideband is EVRC-C.

[0004] In view of the above, the described features of the invention generally relate to one or more improved systems, methods and/or apparatuses for communicating speech. In one embodiment, the invention comprises a method of communicating speech comprising time-warping a residual low band speech signal to an expanded or compressed version of the residual low band speech signal, time-warping a high band

speech signal to an expanded or compressed version of the high band speech signal, and merging the time-warped low band and high band speech signals to give an entire time-warped speech signal. In one aspect of the invention, the residual low band speech signal is synthesized after time-warping of the residual low band signal while in the high band, synthesizing is performed before time-warping of the high band speech signal. The method may further comprise classifying speech segments and encoding the speech segments. The encoding of the speech segments may be one of code-excited linear prediction, noise-excited linear prediction or 1/8 (silence) frame coding. The low band may represent the frequency band up to about 4 kHz and the high band may represent the band from about 3.5 kHz to about 7 kHz.

[0005] In another embodiment, there is disclosed a vocoder having at least one input and at least one output, the vocoder comprising an encoder comprising a filter having at least one input operably connected to the input of the vocoder and at least one output; and a decoder comprising a synthesizer having at least one input operably connected to the at least one output of the encoder and at least one output operably connected to the at least one output of the vocoder. In this embodiment, the decoder comprises a memory, wherein the decoder is adapted to execute software instructions stored in the memory comprising time-warping a residual low band speech signal to an expanded or compressed version of the residual low band speech signal, time-warping a high band speech signal to an expanded or compressed version of the high band speech signal, and merging the time-warped low band and high band speech signals to give an entire time-warped speech signal. The synthesizer may comprise means for synthesizing the time-warped residual low band speech signal, and means for synthesizing the high band speech signal before time-warping it. The encoder comprises a memory and may be adapted to execute software instructions stored in the memory comprising classifying speech segments as 1/8 (silence) frame, code-excited linear prediction or noise-excited linear prediction.

[0006] Further scope of applicability of the present invention will become apparent from the following detailed description, claims, and drawings. However, it should be understood that the detailed description and specific examples, while indicating preferred embodiments of the invention, are given by way of illustration only, since various changes and modifications within the spirit and scope of the invention will become apparent to those skilled in the art.

BRIEF DESCRIPTION OF THE DRAWINGS

[0007] The present invention will become more fully understood from the detailed description given here below, the appended claims, and the accompanying drawings in which:

[0008] FIG. 1 is a block diagram of a Linear Predictive Coding (LPC) vocoder;

Comment [r1]: Is this figure incomplete in terms of labels?

[0009] FIG. 2A is a speech signal containing voiced speech;

[0010] FIG. 2B is a speech signal containing unvoiced speech;

[0011] FIG. 2C is a speech signal containing transient speech;

[0012] FIG. 3 is a block diagram illustrating time-warping of low band and high band;

[0013] FIG. 4A depicts determining pitch delays through interpolation;

[0014] FIG. 4B depicts identifying pitch periods;

[0015] FIG. 5A represents an original speech signal in the form of pitch periods;

[0016] FIG. 5B represents a speech signal expanded using overlap/add; and

[0017] FIG. 5C represents a speech signal compressed using overlap/add.

DETAILED DESCRIPTION

[0018] The word “illustrative” is used herein to mean “serving as an example, instance, or illustration.” Any embodiment described herein as “illustrative” is not necessarily to be construed as preferred or advantageous over other embodiments.

[0019] Time-warping has a number of applications in packet-switched networks where vocoder packets may arrive asynchronously. While time-warping may be performed either inside or outside the vocoder, performing it in the vocoder offers a number of advantages such as better quality of warped frames and reduced computational load. The techniques described herein may be easily applied to other vocoders that use similar techniques such as 4GV-Wideband, the standards name for which is EVRC-C, to vocode voice data.

Description of Vocoder Functionality

[0020] Human voices comprise of two components. One component comprises fundamental waves that are pitch-sensitive and the other is fixed harmonics that are not pitch sensitive. The perceived pitch of a sound is the ear's response to frequency, i.e., for most practical purposes the pitch is the frequency. The harmonics components add

distinctive characteristics to a person's voice. They change along with the vocal cords and with the physical shape of the vocal tract and are called formants.

[0021] Human voice may be represented by a digital signal $s(n)$ 10 (see FIG. 1). Assume $s(n)$ 10 is a digital speech signal obtained during a typical conversation including different vocal sounds and periods of silence. The speech signal $s(n)$ 10 may be portioned into frames 20 as shown in FIGs. 2A – 2C. In one aspect, $s(n)$ 10 is digitally sampled at 8 kHz. In other aspects, $s(n)$ 10 may be digitally sampled at 16kHz or 32kHz or some other sampling frequency.

[0022] Current coding schemes compress a digitized speech signal 10 into a low bit rate signal by removing all of the natural redundancies (i.e., correlated elements) inherent in speech. Speech typically exhibits short term redundancies resulting from the mechanical action of the lips and tongue, and long term redundancies resulting from the vibration of the vocal cords. Linear Predictive Coding (LPC) filters the speech signal 10 by removing the redundancies producing a residual speech signal. It then models the resulting residual signal as white Gaussian noise. A sampled value of a speech waveform may be predicted by weighting a sum of a number of past samples, each of which is multiplied by a linear predictive coefficient. Linear predictive coders, therefore, achieve a reduced bit rate by transmitting filter coefficients and quantized noise rather than a full bandwidth speech signal 10.

[0023] A block diagram of one embodiment of a LPC vocoder 70 is illustrated in FIG. 1. The function of the LPC is to minimize the sum of the squared differences between the original speech signal and the estimated speech signal over a finite duration. This may produce a unique set of predictor coefficients which are normally estimated every frame 20. A frame 20 is typically 20 ms long. The transfer function of a time-varying digital filter 75 may be given by:

$$H(z) = \frac{G}{1 - \sum a_k z^{-k}},$$

where the predictor coefficients may be represented by a_k and the gain by G .

[0024] The summation is computed from $k = 1$ to $k = p$. If an LPC-10 method is used, then $p = 10$. This means that only the first 10 coefficients are transmitted to a LPC synthesizer 80. The two most commonly used methods to compute the coefficients are, but not limited to, the covariance method and the auto-correlation method.

[0025] Typical vocoders produce frames 20 of 20 msec duration, including 160 samples at the preferred 8 kHz rate or 320 samples at 16 kHz rate. A time-warped compressed version of this frame 20 has a duration smaller than 20 msec, while a time-warped expanded version has a duration larger than 20 msec. Time-warping of voice data has significant advantages when sending voice data over packet-switched networks, which introduce delay jitter in the transmission of voice packets. In such networks, time-warping may be used to mitigate the effects of such delay jitter and produce a “synchronous” looking voice stream.

Comment [r2]: 4GV wideband produces 320 samples at 16k sampling. Can you add both “320 samples” and “8khz” and where this is used in the docu

[0026] Embodiments of the invention relate to an apparatus and method for time-warping frames 20 inside the vocoder 70 by manipulating the speech residual. In one embodiment, the present method and apparatus is used in 4GV wideband. The disclosed embodiments comprise methods and apparatuses or systems to expand/compress different types of 4GV wideband speech segments encoded using Code-Excited Linear Prediction (CELP) or (Noise-Excited Linear Prediction (NELP) coding.

[0027] The term “vocoder” 70 typically refers to devices that compress voiced speech by extracting parameters based on a model of human speech generation. Vocoders 70 include an encoder 204 and a decoder 206. The encoder 204 analyzes the incoming speech and extracts the relevant parameters. In one embodiment, the encoder comprises the filter 75. The decoder 206 synthesizes the speech using the parameters that it receives from the encoder 204 via a transmission channel 208. In one embodiment, the decoder comprises the synthesizer 80. The speech signal 10 is often divided into frames 20 of data and block processed by the vocoder 70.

[0028] Those skilled in the art will recognize that human speech may be classified in many different ways. Three conventional classifications of speech are voiced, unvoiced sounds and transient speech.

[0029] FIG. 2A is a voiced speech signal $s(n)$ 402. FIG. 2A shows a measurable, common property of voiced speech known as the pitch period 100.

[0030] FIG. 2B is an unvoiced speech signal $s(n)$ 404. An unvoiced speech signal 404 resembles colored noise.

[0031] FIG. 2C depicts a transient speech signal $s(n)$ 406, i.e., speech which is neither voiced nor unvoiced. The example of transient speech 406 shown in FIG. 2C might represent $s(n)$ transitioning between unvoiced speech and voiced speech. These three classifications are not all inclusive. There are many different classifications of speech

that may be employed according to the methods described herein to achieve comparable results.

4GV Wideband Vocoder

[0032] The fourth generation vocoder (4GV) provides attractive features for use over wireless networks as further described in co-pending patent application Serial Number 11/123,467, filed on May 5, 2005, entitled "Time Warping Frames Inside the Vocoder by Modifying the Residual," which is fully incorporated herein by reference. Some of these features include the ability to trade-off quality vs. bit rate, more resilient vocoding in the face of increased packet error rate (PER), better concealment of erasures, etc. In the present invention, the 4GV wideband vocoder is disclosed that encodes speech using a split-band technique, i.e., the lower and upper bands are separately encoded.

[0033] In one embodiment, an input signal represents wideband speech sampled at 16 kHz. An analysis filterbank is provided generating a narrowband (low band) signal sampled at 8 kHz, and a high band signal sampled at 7 kHz. This high band signal represents the band from about 3.5 kHz to about 7 kHz in the input signal, while the low band signal represents the band up to about 4 kHz, and the final reconstructed wideband signal will be limited in bandwidth to about 7 kHz. It should be noted that there is an approximately 500 Hz overlap between the low and high bands, allowing for a more gradual transition between the bands.

[0034] In one aspect, the narrowband signal is encoded using a modified version of the narrowband EVRC-B speech coder, which is a CELP coder with a frame size of 20 milliseconds. Several signals from the narrowband coder are used by the high band analysis and synthesis; these are: (1) the excitation (i.e., quantized residual) signal from the narrowband coder; (2) the quantized first reflection coefficient (as an indicator of the spectral tilt of the narrowband signal); (3) the quantized adaptive codebook gain; and (4) the quantized pitch lag.

[0035] The modified EVRC-B narrowband encoder used in 4GV wideband encodes each frame voice data in one of three different frame types: Code-Excited Linear Prediction (CELP); Noise-Excited Linear Prediction (NELP); or silence $1/8^{\text{th}}$ rate frame.

[0036] CELP is used to encode most of the speech, which includes speech that is periodic as well as that with poor periodicity. Typically, about 75% of the non-silent frames are encoded by the modified EVRC-B narrowband encoder using CELP.

[0037] NELP is used to encode speech that is noise-like in character. The noise-like character of such speech segments may be reconstructed by generating random signals at the decoder and applying appropriate gains to them.

[0038] $1/8^{\text{th}}$ rate frames are used to encode background noise, i.e., periods where the user is not talking.

Time-Warping 4GV Wideband Frames

[0039] Since the 4GV wideband vocoder encodes lower and upper bands separately, the same philosophy is followed in time-warping the frames. The lower band is time-warped using a similar technique as described in the above-mentioned co-pending patent application entitled "Time Warping Frames Inside the Vocoder by Modifying the Residual."

[0040] Referring to FIG. 3, there is shown a lower-band warping 32 that is applied on a residual signal 30. The main reason for doing time-warping 32 in the residual domain is that this allows the LPC synthesis 34 to be applied to the time-warped residual signal. The LPC coefficients play an important role in how speech sounds and applying synthesis 34 after warping 32 ensures that correct LPC information is maintained in the signal. If time-warping is done after the decoder, on the other hand, the LPC synthesis has already been performed before time-warping. Thus, the warping procedure may change the LPC information of the signal, especially if the pitch period estimation has not been very accurate.

Time-Warping of Residual Signal When Speech Segment is CELP

[0041] In order to warp the residual, the decoder uses pitch delay information contained in the encoded frame. This pitch delay is actually the pitch delay at the end of the frame. It should be noted here that even in a periodic frame, the pitch delay might be slightly changing. The pitch delays at any point in the frame may be estimated by interpolating between the pitch delay of the end of the last frame and that at the end of the current frame. This is shown in FIG. 4. Once pitch delays at all points in the frame are known, the frame may be divided into pitch periods. The boundaries of pitch periods are determined using the pitch delays at various points in the frame.

[0042] FIG. 4A shows an example of how to divide the frame into its pitch periods. For instance, sample number 70 has pitch delay of approximately 70 and sample number 142 has pitch delay of approximately 72. Thus, pitch periods are from [1-70] and from [71-142]. This is illustrated in FIG. 4B.

- [0043] Once the frame has been divided into pitch periods, these pitch periods may then be overlap/added to increase/decrease the size of the residual. The overlap/add technique is a known technique and FIGS. 5A–5C show how it is used to expand/compress the residual.
- [0044] Alternatively, the pitch periods may be repeated if the speech signal needs to be expanded. For instance, in FIG. 5B, pitch period PP1 may be repeated (instead of overlap-added with PP2) to produce an extra pitch period.
- [0045] Moreover, the overlap/adding and/or repeating of pitch periods may be done as many times as is required to produce the amount of expansion/compression required.
- [0046] Referring to FIG. 5A, the original speech signal comprising of 4 pitch periods (PPs) is shown. FIG. 5B shows how this speech signal may be expanded using overlap/add. In FIG. 5B, pitch periods PP2 and PP1 are overlap/added such that PP2's contribution goes on decreasing and that of PP1 is increasing. FIG. 5C illustrates how overlap/add is used to compress the residual.
- [0047] In cases when the pitch period is changing, the overlap-add technique may require the merging of two pitch periods of unequal length. In this case, better merging may be achieved by aligning the peaks of the two pitch periods before overlap/adding them.
- [0048] The expanded/compressed residual is finally sent through the LPC synthesis.
- [0049] Once the lower band is warped, the upper band needs to be warped using the pitch period from the lower band, i.e., for expansion, a pitch period of samples is added, while for compressing, a pitch period is removed.
- [0050] The procedure for warping the upper band is different from the lower band. Referring back to FIG. 3, the upper band is not warped in the residual domain, but rather warping 38 is done after synthesis 36 of the upper band samples. The reason for this is that the upper band is sampled at 7 kHz, while the lower band is sampled at 8 kHz. Thus, the pitch period of the lower band (sampled at 8 kHz) may become a fractional number of samples when the sampling rate is 7 kHz, as in the upper band. As an example, if the pitch period is 25 in the lower band, in the upper band's residual domain, this will require $25 * 7/8 = 21.875$ samples to be added/removed from the upper band's residual. Clearly, since a fractional number of samples cannot be generated, the upper band is warped 38 after it has been resampled to 8 kHz, which is the case after synthesis 36.

[0051] Once the lower band is warped 32, the unwarped lower band excitation (consisting of 160 samples) is passed to the upper band decoder. Using this unwarped lower band excitation, the upper band decoder produces 140 samples of upper band at 7 kHz. These 140 samples are then passed through a synthesis filter 36 and resampled to 8 kHz, giving 160 upper band samples.

[0052] These 160 samples at 8 kHz are then time-warped 38 using the pitch period from the lower band and the overlap/add technique used for warping the lower band CELP speech segment.

[0053] The upper and lower bands are finally added or merged to give the entire warped signal.

Time-Warping of Residual Signal When Speech Segment is NELP

[0054] For NELP speech segments, the encoder encodes only the LPC information as well as the gains of different parts of the speech segment for the lower band. The gains may be encoded in "segments" of 16 PCM samples each. Thus, the lower band may be represented as 10 encoded gain values (one each for 16 samples of speech).

[0055] The decoder generates the lower band residual signal by generating random values and then applying the respective gains on them. In this case, there is no concept of pitch period and as such, the lower band expansion/compression does not have to be of the granularity of a pitch period.

[0056] In order to expand/compress the lower band of a NELP encoded frame, the decoder may generate a larger/smaller number of segments than 10. The lower band expansion/compression in this case is by a multiple of 16 samples, leading to $N = 16 * n$ samples, where n is the number of segments. In case of expansion, the extra added segments can take the gains of some function of the first 10 segments. As an example, the extra segments may take the gain of the 10th segment.

[0057] Alternately, the decoder may expand/compress the lower band of a NELP encoded frame by applying the 10 decoded gains to sets of y (instead of 16) samples to generate an expanded ($y > 16$) or compressed ($y < 16$) lower band residual.

[0058] The expanded/compressed residual is then sent through the LPC synthesis to produce the lower band warped signal.

[0059] Once the lower band is warped, the unwarped lower band excitation (comprising of 160 samples) is passed to the upper band decoder. Using this unwarped lower band excitation, the upper band decoder produces 140 samples of upper band at 7 kHz.

These 140 samples are then passed through a synthesis filter and resampled to 8 kHz, giving 160 upper band samples.

[0060] These 160 samples at 8 kHz are then time-warped in a similar way as the upper band warping of CELP speech segments, i.e., using overlap/add. When using overlap/add for the upper-band of NELP, the amount to compress/expand is the same as the amount used for the lower band. In other words, the “overlap” used for the overlap/add method is assumed to be the amount of expansion/compression in the lower band. As an example, if the lower band produced 192 samples after warping, the overlap period used in the overlap/add method is $192 - 160 = 32$ samples.

[0061] The upper and lower bands are finally added to give the entire warped NELP speech segment.

[0062] Those of skill in the art would understand that information and signals may be represented using any of a variety of different technologies and techniques. For example, data, instructions, commands, information, signals, bits, symbols, and chips that may be referenced throughout the above description may be represented by voltages, currents, electromagnetic waves, magnetic fields or particles, optical fields or particles, or any combination thereof.

[0063] Those of skill would further appreciate that the various illustrative logical blocks, modules, circuits, and algorithm steps described in connection with the embodiments disclosed herein may be implemented as electronic hardware, computer software, or combinations of both. To clearly illustrate this interchangeability of hardware and software, various illustrative components, blocks, modules, circuits, and steps have been described above generally in terms of their functionality. Whether such functionality is implemented as hardware or software depends upon the particular application and design constraints imposed on the overall system. Skilled artisans may implement the described functionality in varying ways for each particular application, but such implementation decisions should not be interpreted as causing a departure from the scope of the present invention.

[0064] The various illustrative logical blocks, modules, and circuits described in connection with the embodiments disclosed herein may be implemented or performed with a general purpose processor, a Digital Signal Processor (DSP), an Application Specific Integrated Circuit (ASIC), a Field Programmable Gate Array (FPGA) or other programmable logic device, discrete gate or transistor logic, discrete hardware

components, or any combination thereof designed to perform the functions described herein. A general purpose processor may be a microprocessor, but in the alternative, the processor may be any conventional processor, controller, microcontroller, or state machine. A processor may also be implemented as a combination of computing devices, e.g., a combination of a DSP and a microprocessor, a plurality of microprocessors, one or more microprocessors in conjunction with a DSP core, or any other such configuration.

[0065] The steps of a method or algorithm described in connection with the embodiments disclosed herein may be embodied directly in hardware, in a software module executed by a processor, or in a combination of the two. A software module may reside in Random Access Memory (RAM), flash memory, Read Only Memory (ROM), Electrically Programmable ROM (EPROM), Electrically Erasable Programmable ROM (EEPROM), registers, hard disk, a removable disk, a CD-ROM, or any other form of storage medium known in the art. An illustrative storage medium is coupled to the processor such the processor can read information from, and write information to, the storage medium. In the alternative, the storage medium may be integral to the processor. The processor and the storage medium may reside in an ASIC. The ASIC may reside in a user terminal. In the alternative, the processor and the storage medium may reside as discrete components in a user terminal.

[0066] The previous description of the disclosed embodiments is provided to enable any person skilled in the art to make or use the present invention. Various modifications to these embodiments will be readily apparent to those skilled in the art, and the generic principles defined herein may be applied to other embodiments without departing from the spirit or scope of the invention. Thus, the present invention is not intended to be limited to the embodiments shown herein but is to be accorded the widest scope consistent with the principles and novel features disclosed herein.

CLAIMS

Comment [r3]: One thing n
is the Upper Band warping for
I only saw claims for Upper B:
warping for CELP

1. A method of communicating speech, comprising:
time-warping a residual low band speech signal to an expanded or compressed version of the residual low band speech signal;
time-warping a high band speech signal to an expanded or compressed version of the high band speech signal; and
merging the time-warped low band and high band speech signals to give an entire time-warped speech signal.
2. The method of claim 1, further comprising synthesizing the time-warped residual low band speech signal.
3. The method of claim 2, further comprising synthesizing the high band speech signal before time-warping it.
4. The method of claim 3, further comprising:
classifying speech segments; and
encoding the speech segments.
5. The method of claim 4, wherein encoding the speech segments comprises using code-excited linear prediction, noise-excited linear prediction or 1/8 frame coding.
6. The method of claim 4, wherein the encoding is code-excited linear prediction encoding.
7. The method of claim 4, wherein the encoding is noise-excited linear prediction encoding.
8. The method of claim 7, wherein the encoding comprises encoding linear predictive coding information as gains of different parts of a speech frame.

9. The method of claim 8, wherein the gains are encoded for sets of speech samples.

10. The method of claim 9, further comprising generating a residual low band signal by generating random values and then applying the gains to the random values.

11. The method of claim 9, further comprising representing the linear predictive coding information as 10 encoded gain values for the residual low band speech signal, wherein each encoded gain value represents 16 samples of speech.

12. The method of claim 7, further comprising producing 140 samples of the high band speech signal from an unwarped low band excitation signal.

13. The method of claim 7, wherein the time-warping of the low band speech signal comprises generating a higher/lower number of samples and applying some function of the decoded gains of the parts of a speech frame to the residual and then synthesizing it.

14. The method of claim 13, wherein the applying of some function of the decoded gains of parts of the speech frame to the residual comprises applying the gain of the last speech segment to the additional samples when the lower band is expanded.

15. The method of claim 7, wherein the time-warping of the high band speech signal comprises:

overlap/adding the same number of samples as were compressed in the lower band if the high band speech signal is compressed; and

overlap/adding the same number of samples as were expanded in the lower band if the high band speech signal is expanded.

16. The method of claim 6, wherein the time-warping of the residual low band speech signal comprises:

estimating at least one pitch period; and

adding or subtracting at least one of the pitch periods after receiving the residual low band speech signal.

17. The method of claim 16, wherein the time-warping of the high band speech signal comprises:

using the pitch periods from the low band speech signal;

overlap/adding one or more pitch periods if the high band speech signal is compressed; and

overlap/adding or repeating one or more pitch periods if the high band speech signal is expanded.

18. The method of claim 6, wherein the time-warping of the residual low band speech signal comprises:

estimating pitch delay;

dividing a speech frame into pitch periods, wherein boundaries of the pitch periods are determined using the pitch delay at various points in the speech frame;

overlap/adding the pitch periods if the residual low band speech signal is compressed; and

overlap/adding or repeating one or more pitch periods if the residual low band speech signal is expanded.

19. The method of claim 18, wherein the time-warping of the high band speech signal comprises:

using the pitch periods from the low band speech signal;

overlap/adding the pitch periods if the high band speech signal is compressed;

and

overlap/adding or repeating one or more pitch periods if the high band speech signal is expanded.

20. The method of claim 18, wherein the estimating of the pitch delay comprises interpolating between a pitch delay of an end of a last frame and an end of a current frame.

21. The method of claim 18, wherein the overlap/adding or repeating one or more of the pitch periods comprises merging the speech segments.

22. The method of claim 18, wherein the overlap/adding or repeating one or more of the pitch periods if the residual low band speech signal is expanded comprises adding an additional pitch period created from a first pitch segment and a second pitch period segment.

23. The method of claim 21, further comprising selecting similar speech segments, wherein the similar speech segments are merged.

24. The method of claim 21, further comprising correlating the speech segments, whereby similar speech segments are selected.

25. The method of claim 22, wherein the adding of an additional pitch period created from a first pitch segment and a second pitch period segment comprises adding the first and second pitch segments such that the first pitch period segment's contribution increases and the second pitch period segment's contribution decreases.

26. The method of claim 1, wherein the low band represents the band up to and including 4 kHz.

27. The method of claim 1, wherein the high band represents the band from about 3.5 kHz to about 7 kHz.

28. A vocoder having at least one input and at least one output, comprising:
an encoder comprising a filter having at least one input operably connected to the input of the vocoder and at least one output; and
a decoder comprising a synthesizer having at least one input operably connected to the at least one output of the encoder and at least one output operably connected to the at least one output of the vocoder.

29. The vocoder of claim 28, wherein the decoder comprises:
a memory, wherein the decoder is adapted to execute software instructions stored in the memory comprising:

time-warping a residual low band speech signal to an expanded or compressed version of the residual low band speech signal;

time-warping a high band speech signal to an expanded or compressed version of the high band speech signal; and

merging the time-warped low band and high band speech signals to give an entire time-warped speech signal.

30. The vocoder of claim 29, wherein the synthesizer comprises means for synthesizing the time-warped residual low band speech signal.

31. The vocoder of claim 30, wherein the synthesizer further comprises means for synthesizing the high band speech signal before time-warping it.

32. The vocoder of claim 28, wherein the encoder comprises a memory and the encoder is adapted to execute software instructions stored in the memory comprising classifying speech segments as 1/8 frame, code-excited linear prediction or noise-excited linear prediction.

33. The vocoder of claim 31, wherein the encoder comprises a memory and the encoder is adapted to execute software instructions stored in the memory comprising encoding speech segments using code-excited linear prediction encoding.

34. The vocoder of claim 31, wherein said encoder comprises a memory and the encoder is adapted to execute software instructions stored in the memory comprising encoding speech segments using noise-excited linear prediction encoding.

35. The vocoder of claim 34, wherein the encoding of the speech segments using noise-excited linear prediction encoding software instruction comprises encoding linear predictive coding information as gains of different parts of a speech segment.

36. The vocoder of claim 35, wherein the gains are encoded for sets of speech samples.

37. The vocoder of claim 36, wherein the time-warping instruction of the residual low band speech signal further comprises generating a residual low band speech signal by generating random values and then applying the gains to the random values.

38. The vocoder according to claim 36, wherein the time-warping instruction of the residual low band speech signal further comprises representing the linear predictive coding information as 10 encoded gain values for the residual low band speech signal, wherein each encoded gain value represents 16 samples of speech.

39. The vocoder of claim 34, further comprising producing 140 samples of the high band speech signal from an unwarped low band excitation signal.

40. The vocoder of claim 34, wherein the time-warping software instruction of the low band speech signal comprises generating a higher/lower number of samples and applying some function of the decoded gains of parts of a speech frame to the residual and then synthesizing it.

41. The vocoder of claim 40, wherein the applying of some function of the decoded gains of parts of the speech frame to the residual comprises applying the gain of the last speech segment to the additional samples when the lower band is expanded.

42. The vocoder of claim 33, wherein the time-warping software instruction of the high band speech signal comprises:

overlap/adding the same number of samples as were compressed in the lower band if the high band speech signal is compressed; and

overlap/adding the same number of samples as were expanded in the lower band if the high band speech signal is expanded.

43. The vocoder of claim 33, wherein the time-warping software instruction of the residual low band speech signal comprises:

estimating at least one pitch period; and

adding or subtracting the at least one pitch period after receiving the residual low band speech signal.

44. The vocoder of claim 43, wherein the time-warping software instruction of the high band speech signal comprises:

using the pitch period from the low band speech signal;

overlap/adding one or more pitch periods if the high band speech signal is compressed; and

overlap/adding or repeating one or more pitch periods if the high band speech signal is expanded.

45. The vocoder of claim 33, wherein the time-warping software instruction of the residual low band speech signal comprises:

estimating pitch delay;

dividing a speech frame into pitch periods, wherein boundaries of the pitch periods are determined using the pitch delay at various points in the speech frame;

overlap/adding the pitch periods if the residual speech signal is compressed; and

overlap/adding or repeating one or more pitch periods if the residual speech signal is expanded.

46. The vocoder of claim 45, wherein the time-warping software instruction of the high band speech signal comprises:

using the pitch periods from the low band speech signal;

overlap/adding the pitch periods if the high band speech signal is compressed;

and

overlap/adding or repeating one or more pitch periods if the high band speech signal is expanded.

47. The vocoder of claim 45, wherein the overlap/adding instruction of the pitch periods if the residual low band speech signal is compressed comprises:

segmenting an input sample sequence into blocks of samples;

removing segments of the residual signal at regular time intervals;

merging the removed segments; and

replacing the removed segments with a merged segment.

48. The vocoder of claim 45, wherein the estimating instruction of the pitch delay comprises interpolating between a pitch delay of an end of a last frame and an end of a current frame.

49. The vocoder of claim 45, wherein the overlap/adding or repeating one or more of the pitch periods instruction comprises merging the speech segments.

50. The vocoder of claim 45, wherein the overlap/adding or repeating one or more of the pitch periods instruction if the residual low band speech signal is expanded comprises adding an additional pitch period created from a first pitch period segment and a second pitch period segment.

51. The vocoder of claim 47, wherein the merging instruction of the removed segments comprises increasing a first pitch period segment's contribution and decreasing a second pitch period segment's contribution.

52. The vocoder of claim 49, further comprising selecting similar speech segments, wherein the similar speech segments are merged.

53. The vocoder of claim 49, wherein the time-warping instruction of the residual low band speech signal further comprises correlating the speech segments, whereby similar speech segments are selected.

54. The vocoder of claim 50, wherein the adding instruction of an additional pitch period created from the first and second pitch period segments comprises adding the first and second pitch period segments such that the first pitch period segment's contribution increases and the second pitch period segment's contribution decreases.

55. The vocoder of claim 29, wherein the low band represents the band up to and including 4 kHz.

56. The vocoder of claim 29, wherein the high band represents the band from about 3.5 kHz to about 7 kHz.