

METHOD AND APPARATUS FOR CHANGING THE TIMBRE AND/OR PITCH OF AUDIO SIGNALS

5 FIELD OF THE INVENTION

The present invention relates generally to electronic audio effects and in particular to musical effects that shift the timbre and/or pitch of audio signals.

BACKGROUND OF THE INVENTION

10 In any periodic musical note, there is a fundamental frequency that determines the particular pitch of the note, as well as numerous harmonics which provide character or timbre to the note. It is the particular combination of the harmonic frequencies with the fundamental frequency that make, for example, a guitar and a violin playing the same note sound different from one another. The relationship of the amplitude of the fundamental
15 frequency component to the amplitudes of the harmonics created by an instrument or voice, is referred to as the spectral envelope. In a musical instrument such as a guitar, flute, or saxophone, the spectral envelope of a note played by the instrument expands and contracts more or less proportionally as the pitch of the note is shifted up or down.

Electronic pitch shifters are musical effects that receive an input note and produce an
20 output note with a different pitch. Such effects are often used to allow a single musician to sound like several. For musical instruments, one can change the pitch of a note by sampling and recording the sound from the instrument and playing back the sampled and recorded sounds at a rate that is faster or slower than the rate at which the samples were recorded. The output notes created by this technique sound fairly natural because the spectral envelope
25 of the pitch shifted sounds mimics how the spectral envelope of the sounds produced by the instrument varies with pitch.

In contrast to notes produced by musical instruments, the spectral envelope of vocal notes or sounds does not vary proportionally as the pitch of the vocal note varies. However, the relative magnitudes of the individual frequencies that make up this spectral envelope may
30 change. Shifting the pitch of a vocal note by sampling a note as it is sung or spoken and playing the samples back at a different speed, does not sound natural because the process varies the shape of the spectral envelope in proportion to the amount of pitch shift. To shift realistically the pitch of a vocal sound, a method is required for varying the fundamental frequency while only slightly varying the overall shape of the spectral envelope.

A device that shifts the pitch of vocal notes to create harmonies in real time is described in our prior U.S. Patent No. 5,231,671 (the " '671 patent"). The method of pitch shifting described in the '671 patent was adapted from an article, Lent, K. "An Efficient Method for Pitch Shifting Digitally Sampled Sounds," *Computer Music Journal*, Volume 13, No. 4, (1989) (the "Lent method"). The Lent method allows the pitch of a digitally sampled sound to be shifted without changing the spectral envelope. Briefly stated, the Lent method can be used to shift the pitch of a vocal note by replicating portions of a stored input signal at a rate that is faster or slower than the fundamental frequency input note. While this method of shifting the pitch of vocal notes works well, the pitch shifted notes do not sound completely natural, because the spectral envelope remains fixed as the pitches of the notes are varied.

As described above, there are two methods of electronically shifting the pitch of a note. The first method, referred to as resampling, modifies the spectral envelope in proportion to the amount of pitch shift. The second, Lent method more or less maintains the spectral envelope regardless of the amount of pitch shift. Neither of these two methods allow the spectral envelope to be varied in a controllable manner. Therefore, there is a need for a method of altering the spectral envelope of a musical note that is not dependent on the pitch of a note. With such a method, more realistic harmonies can be created. In addition, by changing the timbre of the note with or without changing the output pitch, it is possible to make one instrument sound like another, or one person's voice sound like another.

SUMMARY OF THE INVENTION

To shift the timbre of both vocal notes and notes produced by musical instruments, the present invention uses a novel combination of pitch shifting by altering the sampling rate of a signal and pitch shifting according to the Lent method. In the preferred embodiment, the input signal is sampled at a first rate, and the resulting digital representation is stored in a memory buffer. The stored digital input signal is then resampled at a second rate that is determined by a user. The resampled input signal is then stored in a second memory buffer. The pitch of the resampled input signal is then shifted by scaling the resampled input signal with a window function at a rate equal to the fundamental frequency of the output note desired. If it is desired to only shift the timbre of a note and not the pitch of a note, then

the rate at which the resampled input signal is scaled with the window function is the same as the fundamental frequency of the input note. If it is desired to change the pitch of the output note as well as its timbre, then the rate at which the resampled input signal is scaled with the window function differs from the fundamental frequency of the input note.

5 According to another aspect of the invention, an effect generator is disclosed that can modify the timbre and/or pitch of an input audio signal to match a pitch received on a MIDI channel. Preferably, the effect generator is used with a MIDI karaoke system that provides a stream of melody or harmony notes to the effect generator. The effect generator reads the notes on the MIDI channel and automatically assigns the note an amount of timbre shift. The
10 assignment can be made by comparing the pitch of the harmony note with one or more thresholds or with the pitch of an input audio signal received from a user of the karaoke system. The amount of timbre assigned to each note can make the harmony notes sound different from the input audio signal or can mimic how the input audio signal would change if raised or lowered in pitch.

15 BRIEF DESCRIPTION OF THE DRAWINGS

The foregoing aspects and many of the attendant advantages of this invention will become more readily appreciated as the same becomes better understood by reference to the following detailed description, when taken in conjunction with the accompanying drawings, wherein:
20

FIGURES 1A-1D are representative graphs of the spectra of vocal signals showing how the spectral envelopes change as a result of prior art timbre/pitch shifting techniques as well as the timbre/pitch shifting technique of the present invention;

25 FIGURE 2A is a flow chart of the steps performed by the present invention to shift the timbre and/or pitch of an input note;

FIGURE 2B is a flow chart of the steps performed by the present invention to create timbre shifted, harmony notes from an input vocal note;

FIGURE 3 is a block diagram of a musical effect generator for producing vocal harmonies according to the method of the present invention;

FIGURE 4A and FIGURE 4B are graphs and corresponding diagrammatic memory charts showing how an input vocal signal is resampled according to step of the method of the present invention;

FIGURE 5 is a block diagram showing the functions performed by a digital signal processor that is programmed according to the method of the present invention;

FIGURE 6 is a block diagram showing the functions performed by a windowed audio generator unit within the digital signal processor;

FIGURES 7A and 7B are a graphic representations of the method of shifting the pitch of a digitally sampled vocal signal according to the present invention;

FIGURES 8A and 8B show how a Hanning window is created and stored in memory in the method of the present invention; and

FIGURES 9A and 9B are block diagrams of music effects that dynamically select the amount of timbre shift that is applied to a note.

DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENT

The present invention provides a system for shifting the timbre of a note in a way that sounds more realistic than timbre shifts produced by known systems. In its simplest form, the method can be used to shift the timbre of a note but not the pitch of a note. For example, the method can be used to make a vocal signal sung or spoken by a man sound as if the same note were sung or spoken by a woman. In addition to shifting the timbre of a note, the method of the present invention can be used to change the pitch and timbre of a note. For example, the present invention can be used to make a note sung by a woman sound like another note sung by a man. Finally, the presently preferred embodiment of the invention is used to create timbre shifted, harmony notes from an input note. Although the following description is primarily directed to producing harmony notes from an input vocal note, it will be realized that the note need not be a vocal note but may be produced from any source, and the output note need not be different from or harmonious with the input pitch.

FIGURES 1A-1D compare how the spectral envelope of a vocal note changes when the pitch of the note is shifted according to prior art techniques and by the method of the present invention. FIGURE 1A shows a frequency spectrum 30a that is representative of a typical vocal note. The overall shape of the spectrum is defined by one or more formants

or peaks 32a. The character or timbre of the vocal note is defined by the relative magnitude and position of the fundamental frequency of the note and the harmonics (represented by the arrows 34a).

To shift realistically the pitch of a vocal note, it is necessary to shift the fundamental frequency of a note while maintaining the formants of the spectrum close to those of the original vocal note. FIGURE 1B shows a spectrum 30b of a pitch shifted vocal note that is a musical fifth below the note associated with the spectrum shown in FIGURE 1A. The note associated with the spectrum 30b was created by slowing the playback rate of the sampled original vocal note. As can be seen, the entire spectral envelope defined by the formants 32b as well as the individual harmonics 34b is compressed and shifted to a lower frequency. The result of shifting the formants makes the pitch shifted vocal note sound unnatural.

FIGURE 1C shows a spectrum 30c of a pitch shifted vocal note that is a musical fifth below the note associated with spectrum shown in FIGURE 1A and which was generated in accordance with the method set forth in the '671 patent. The pitch shifted vocal note associated with the spectrum 30c was created by replicating a portion of the input vocal note at a rate that is slower than the fundamental frequency of the original input vocal note. In the spectrum 30c, only the frequencies of the harmonics 34c have changed, as described in the '671 patent. The overall shape of the spectrum remains the same as the spectrum shown in FIGURE 1A. The pitch shifted vocal note associated with the spectrum 30c sounds more natural than the pitch shifted vocal note produced by the note associated with the spectrum 30b shown in FIGURE 1B. However, the pitch shifted vocal note does not sound completely natural. Pitch shifted vocal notes produced by the method described in the '671 patent tend to have timbres that are very similar to the input vocal signal from which they are created. Therefore, all the pitch shifted vocal notes sound like altered variations of the original.

To alter the timbre of a note in a manner that sounds realistic, the present invention uses a novel combination of resampling pitch shifting, whereby the playback rate of the vocal note is altered, and the method described in the '671 patent. The result is a timbre shifted note that can be made to sound deeper and more masculine, or higher and more feminine.

FIGURE 1D shows a spectrum 30d of a pitch shifted vocal note having a frequency that is a musical fifth below the input vocal note associated with the spectrum shown in

FIGURE 1A, and which was generated in accordance with the present invention. As will be described in further detail, the pitch shifted vocal note corresponding to the spectrum 30d was obtained by resampling the previously stored input vocal note at a rate that is slightly slower than the original sampling rate and storing the resampled data in a memory buffer. A portion of the resampled data is then replicated at a rate equal to the fundamental frequency of the musical fifth below the pitch of the input note. As can be seen, the spectrum 30d is slightly compressed but similar to the original spectrum 30a. The result is a pitch shifted vocal note that sounds natural but not like a replicated version of the original input note.

The major steps of the present invention to create a timbre and/or pitch shifted output signal from an input signal are set forth in the flowchart shown in FIGURE 2A. The method begins at step 50 where an input signal is sampled at a first rate by a analog-to-digital converter. The input signal may be produced from a musical instrument such as a flute, guitar, etc., may be a vocal note that is spoken or sung by a user, or may be produced by a digital source such as a synthesizer. After sampling the input signal, the corresponding digital representation of the input signal is stored in a digital memory at step 52. Next, the stored input signal is resampled at a second rate that differs from the first rate at which the input signal was originally sampled. The resampling rate may be fixed at some percentage greater than or less than the original sampling rate. Alternatively, the resampling rate may be selected by the user.

The resampled data is stored in a digital memory at step 56. Finally, the timbre shifted output signal is produced at step 58 by replicating a portion of the resampled data at a rate equal to the fundamental frequency of the desired output signal. For example, if it is only desired to change the timbre of an input signal, then the rate at which the portion of the resampled data is replicated is equal to the fundamental frequency of the input signal. Alternatively, it may be desired to change the timbre and pitch of the input signal, in which case the rate at which the portion of the resampled data is replicated is not the same as the fundamental frequency of the input signal. Finally, for the case in which the method of the present invention is used in harmony effect generators, the rate at which the portion of the resampled data is replicated is set to a fundamental frequency that is harmonically related to the fundamental frequency of the input signal.

In the current implementation of the invention, the timbre shifting technique is used to create harmony notes from input vocal notes sung by a user. Therefore, although the following description is directed to producing timbre shifted, vocal harmony notes, it will be appreciated that the method of the present invention can also be used to vary only the timbre of an input signal or to vary the timbre and pitch of an input signal in a way that is not harmonically related to the pitch of the input signal.

FIGURE 2B is a flow chart of the major steps performed in the present invention to produce timbre shifted, vocal harmonies. The method begins at step 60 wherein the analog input vocal note is sampled and digitized at a first rate. At step 62, the digital samples are stored in a first memory buffer. At step 64, the stored samples are analyzed to determine the pitch of the input vocal note. After the pitch has been determined, the harmony notes to be produced with the input vocal note are selected at step 66. The particular harmony notes produced for a given input note may be preprogrammed, individually selected by a user, or may be received from an external source such as a synthesizer, a sequencer, or an external storage device such as a computer disk, a laser disk, etc.

After the harmony notes are selected, the percent increase or decrease of the sampling rate that has been selected by a user is determined at step 68. The sampling rate may be increased to give the harmony notes a more feminine quality, or decreased to produce harmony notes with a more masculine sound.

At step 70, the digitized input vocal note that was stored in step 62 is resampled at the new rate selected by the user. The resampled data is stored in a second memory buffer. For example, if the user has selected to decrease in the sampling rate, then there will be fewer data samples in the second memory buffer, thereby decreasing the amount of memory required to store the digitized input vocal note. Similarly, if the user has selected to increase the sampling rate, the data of the first buffer will be resampled at a higher rate than the rate at which the data were originally sampled, thereby requiring more samples and increasing the amount of memory required to store the digitized input vocal note in the second buffer. With the data occupying more memory space, the pitch of the note will be lowered, assuming it at the rate at which the samples are read from memory remains the same.

The resampled data is stored in a second memory buffer at step 72. Finally, the harmony notes are created at step 74 by replicating portions of the resampled input vocal

note at rates that are equal to the fundamental frequencies of the harmony notes selected in step 66.

Turning now to FIGURE 3, a musical effect generator 100 that produces timbre shifted, harmony notes according to the method of the present invention receives an input vocal note 105 that is sung by a user. In general, the effect generator has a microprocessor or CPU 138 that is interfaced with a digital signal processor (DSP) 180 and random access memory (RAM) 121 to produce a number of harmony notes 105a, 105b, 105c, and 105d that are combined with the input vocal note to produce a multi-voice output, as described in detail below.

The microprocessor 138 includes its own read only memory (ROM) 140 and random access memory (RAM) 144. A set of input controls 148 are coupled to the microprocessor to allow a user to vary the operating parameters of the musical effect generator. These parameters include selecting which harmony notes will be produced for a given input note and the distribution of the harmony notes between a right and left stereo channel.

A set of displays 150 are operated by the microprocessor. The displays provide a visual indication of how the effect generator is operating and what options have been selected by the user. One or more MIDI ports 154 are coupled to the microprocessor to allow the effect generator to receive MIDI data from other MIDI-compatible instruments or effects. The details of a MIDI port are well known to those of ordinary skill in the art and therefore need not be discussed in further detail.

Finally, the effect generator includes a pair of "gender shift" controls 156. The gender shift controls allow a user to select the amount of resampling pitch shift that will be applied to each harmony note produced. The operation of the gender shift controls is more fully discussed below.

The digital signal processor 180 is a specialized computer chip that performs a variety of functions. The program code to operate the digital signal processor resides in a ROM 141 that is part of the ROM 140 coupled to the microprocessor. Upon startup of the effect generator, the microprocessor 138 loads the digital signal processor with the appropriate computer program to generate the harmony notes according to the method of the present invention.

The effect generator 100 includes a microphone 110 that receives the user's input vocal note and converts it to a corresponding analog electrical vocal signal. The input vocal signal is also referred to as the "dry" audio signal. The input vocal signal is supplied to a low pass filter 114 that removes any high frequency, extraneous noise. The filtered input vocal signal is transmitted to an analog-to-digital (A/D) converter 118 that periodically samples the input vocal signal and converts it to digital form. Each time the A/D converter has a new sample ready, it interrupts the DSP 180 causing the DSP to read the sample and store it in a first memory buffer 122 that is part of the effect generator's random access memory.

Once the input vocal signal has been sampled and stored in the first memory buffer 122, the digital signal processor 180 implements a pitch recognition routine 188 that analyzes the data stored in the memory buffer 122 and determines its pitch. The method used to determine the pitch of a note is fully described in our U.S. Patent No. 4,688,464, which is herein incorporated by reference. For the purposes of this specification, the terms "pitch" and "fundamental frequency" of a note are interchangeable. From the pitch of the input vocal note, the period of the note is calculated.

Conventionally, the period of a note is simply the inverse of its fundamental frequency expressed in seconds. However, in the present embodiment of the invention, the period is calculated and stored in terms of the number of memory locations required to store a complete cycle of the input vocal signal. For example, one complete cycle of the note A 440Hz occupies 109 memory locations if sampled at 48 KHz ($1/440 \times 48,000$). Therefore, the period of A 440Hz is stored as 109. In addition to determining the pitch and period of a note, the digital signal processor also calculates a period marker which is a pointer to a location in memory where a new cycle of the input vocal signal begins. Initially, the period marker is set to point to the beginning of the memory buffer in which the input vocal is stored. Subsequent period markers are calculated by adding the number of data samples in a single cycle of the input vocal signal (i.e. one period), plus the previous period marker. The period marker is updated when a write pointer that points to the next available memory location minus a small delay is beyond where the new period marker will point. The period markers are used by the DSP 180 to produce the harmony notes, as will be described.

The results of the pitch recognition routine 188 are supplied to the microprocessor 138, i.e., a signal of the pitch of the input vocal signal stored in the first buffer 122. Within the ROM 140 of the microprocessor is a look up table that correlates the pitch of an input vocal signal with a MIDI note. In the presently preferred embodiment of the invention, each MIDI note is assigned a number between 0 and 127. For example, the note A 440 Hz is the MIDI note number 69. If an input signal is not exactly on pitch, then the note can either be rounded to the closest MIDI note or assigned a fractional number. For example, a note that is slightly flat of A 440 Hz might be assigned a number such as 68.887 by the microprocessor.

Once the microprocessor has assigned a note to the input vocal signal, the microprocessor determines which harmony notes are to be produced. The particular harmony notes produced can be individually programmed by the user or selected from one or more predefined harmony "rules." For example, a user may program the microprocessor to produce four harmony notes that are a musical third above the input note, a musical fifth above the input note, a musical seventh above the input note, and a musical third below the input note. Alternatively, the user may select a rule such as a "chordal harmony" rule that always produces harmony notes that are the chord tones above and below the input melody line. As will be appreciated, to use a rule such as the chordal harmony rule, the user inputs the chords to be sung, thereby allowing the microprocessor to determine the correct chord tones. The predefined harmony rules are stored within the ROM 140 and are actuated by the user with the input controls 148.

Another way of selecting the harmony notes to be produced is by using the MIDI port 154. Using the port, the microprocessor can receive an indication of which harmony notes to produce from an external source. These notes can be received from a synthesizer, a sequencer, or any other MIDI-compatible device. The effect generator 100 shifts the input vocal signal to have a pitch equal to the pitch of the harmony notes received. Alternatively, the instructions of which harmony notes to produce may be stored on a computer or as a subcode on a laser disk. The laser disk may operate with a karaoke or other entertainment type machine such that, as a user sings the words of a karaoke song, the karaoke machine supplies an indication of the harmony notes to be produced to the musical effect generator 100.

Once the harmony notes have been determined, the digital signal processor 180 implements a resampling subroutine 192 that resamples the input vocal signal stored in the memory buffer 122 at a rate determined by the position of the gender shift controls 156. The resampled data is stored in two memory buffers 128 that are associated with each gender shift control. By sampling at a lower rate, the timbre of the harmony notes will sound more feminine. Alternatively, if the sampling rate is raised, the harmony notes will sound more masculine.

FIGURE 4A shows how the stored input vocal data is resampled by the digital signal processor to compress the spectral envelope and make the input vocal signal sound more masculine. The analog input vocal signal 105 is sampled by the A/D converter 118 at a plurality of equal time intervals 0, 1, 2, 3, . . . , 11. Each sample has a corresponding value a, b, c, . . . , l. The samples are sequentially stored as elements of a circular array within the memory buffer 122. The circular array has a write pointer (wp) that always points to the next available memory location to be filled with new sample data. In addition, the digital signal processor also calculates the last period marker (pm) 122b that indicates where in the memory buffer a new cycle of the input vocal signal begins. As will be appreciated, the number of samples between the last period marker 122b and a previous period marker 122a define one cycle of the input vocal signal.

To compress the spectral content of the input vocal signal, the stored signal is resampled and stored in one of the two memory buffers 128 (shown in FIGURE 3) at a rate slightly higher than the rate at which it was originally sampled. The resampling rate is determined by the setting of the gender shift controls 156. In the example shown in FIGURE 4A, the input vocal signal is slowed by 25 percent. This is accomplished by resampling the data that are stored in the memory buffer 122 at a time period equal to 0.75 times the original sampling period. For example, samples a', b', c', d', . . . are taken at times 0, 0.75, 1.5, 2.25, etc., and stored in the second memory buffer 128.

To calculate values for the data at times between the samples stored in the first memory buffer 122, an interpolation method is used. In the presently preferred embodiment of the invention, linear interpolation is used. For example, to fill in the data for a sample at time 0.75, the digital signal processor reads the value of the sample obtained at time 1 from memory buffer 122, multiplies it by 0.75, and adds to that 0.25 times the value of the

sample obtained at time 0. Although linear interpolation is used in the currently preferred embodiment of the present invention, other more accurate interpolation methods, such as splines, could be used given sufficient computing power within the digital signal processor 180.

5 Once the data have been resampled and stored in the second memory buffer 128, the digital signal processor calculates a period marker 128b to point to the location in the memory buffer 128 where a new cycle of the resampled input vocal signal begins. The period marker 128b is calculated by multiplying the period marker 122b by the percent change in the sampling rate. Thus, the new period marker 128b is calculated by multiplying
10 the period marker 122b by 1.33 ($1/0.75$) and adding the result to the previous period marker 128a in the second memory buffer 128. As can be seen by comparing the two memory buffers 122 and 128 shown in FIGURE 4A, the effect of increasing the sampling rate of the input vocal signal increases the total number of samples required to hold a full cycle of the input vocal signal. For example, the number of samples between the two period
15 markers 122a and 122b in the memory buffer 122 is twelve. By increasing the sampling rate by 33 percent, the number of samples required to hold an entire cycle of the input vocal signal, i.e., the number of samples between period markers 128a and 128b, increases to 16.

20 FIGURE 4B shows how the input vocal signal is resampled by the digital signal processor at a rate that is slower than the rate at which the input vocal signal was originally sampled by the A/D converter 118 and stored in the memory buffer 122. Again, the analog input vocal signal 105 is sampled at a plurality of equal time intervals 0, 1, 2, 3, . . . , 11. Each sample has a corresponding value a, b, c, . . . , l that is stored in the first memory buffer 122. The period marker 122b is calculated to point to the memory location that marks
25 the beginning of a new cycle in the input vocal signal.

30 In FIGURE 4B, the sampling period is shown as being increased by 25 percent. Therefore, the input vocal signal is resampled at times 0, 1.25, 2.5, 3.75, etc., times the original sampling interval. Each sample has a new value a', b', c', . . . , i'. If the sample interval does not exactly align with a one of the previously stored samples, interpolation is used to determine a value for the resampled data. For example, to calculate the value for a sample d' at time 3.75, the digital signal processor calculates the sum of 0.75 times the

value of the data obtained at time 4, and 0.25 times the value of the data obtained at time 3, etc.

Again, once the data has been resampled and stored in the second memory buffer 128, the digital signal processor recalculates the last period marker 128b for the resampled data in the same manner as described above. As can be seen in FIGURE 4B, the number of samples between the period markers 122a and 122b of the original input vocal signal is 12. When the sampling period is increased by 25 percent, only 9.6 samples exist between the period markers 128a and 128b. Therefore, the total number of samples required to store a complete cycle of the input vocal signal has decreased by 20 percent.

In the presently preferred embodiment of the present invention, a user can increase or decrease the sampling rate by $\pm 33\%$. More or less resampling shift could be provided. However, for vocal applications it has been determined that the most realistic sounding time shifts are obtained when the resampling rate is set between -18 and +18%.

Once the input vocal signal has been resampled at a rate indicated by the gender shift controls and stored in the data buffers 128, the DSP 180 recalculates the period of the resampled data. For example, the user may be singing an A note at 440 Hz which has a period of 2.27 milliseconds (109 samples at 48 KHz) and have one of the gender controls set to $\pm 10\%$. When resampled at the new rate, the period of the resampled vocal signal will be 2.043 milliseconds (98 samples at 48 KHz). This new period is used by a window generation routine 196 and to a pitch shifting routine 200 (represented in FIGURE 3) that are implemented by the digital signal processor to create the harmony notes.

With reference to FIGURE 7, the pitch shifting routine operates by scaling a portion of the resampled input vocal signal 400 stored in the memory buffer with a window function 402 to reduce the magnitude of the samples at the beginning and end of the portion, and to maintain the value of the samples in the middle of the portion. The window function 402 is a smoothly varying, bell-shaped function that, in the preferred embodiment of the invention, is a Hanning window. The result of a point-by-point multiplication of the window function 402 and the portion of the resampled vocal signal 400 is a signal segment 406. As can be seen, the resampled vocal signal 400 contains a series of peaks 401a, 401b, 401c etc. The signal segment 406 contains a complete cycle (i.e. one

peak) of the resampled data but has a beginning and an end that are relatively small in magnitude.

Referring now to FIGURE 7B, a harmony note 408 is created by concatenating a series of signal segments 406a, 406b, 406c and 406d together. Comparing the harmony note 408 to the resampled vocal signal 400 (shown in FIGURE 7A), it can be seen that the harmony note has half the number of peaks 408a, 408b, 408c as compared to the resampled data. Therefore, the harmony note 408 will sound an octave below the resampled vocal signal. As will be appreciated, the pitch of the harmony note to be created depends on the rate at which the signal segments, obtained by scaling the resampled vocal signal by the window function, are added together. As described in the '671 patent and in the Lent article, to shift the pitch of a note to any value higher than an octave below the original pitch requires that overlapping signal segments be added together. As will be appreciated, the reason for reducing the magnitude of the samples at the beginning and end of the signal segment is to prevent large variations in the harmony note as a result of adding overlapping signal segments together.

FIGURES 8A and 8B show how the digital signal processor calculates the Hanning windows used in creating the harmony notes. The window generation routine 196 described above stores mathematical representations of four Hanning windows in four memory buffers 134a, 134b, 134c, and 134d (FIGURE 5). Each memory buffer 134a, 134b, 134c and 134d is associated with one of four harmony generators 220, 230, 240, and 250 (FIGURE 5). Within the ROM 140 is a memory buffer 141 that stores a standard Hanning window in 256 memory locations. The values of the data a, b, c, d, etc. stored in the buffer are calculated by the raised cosine formula:

$$(1 - \cos (2\pi x/256))$$

where x represents each sample stored in the buffer. To create a window function within one of the memory buffers 134 that is used to create the harmony notes, the length of the window is first determined and then the window is filled with new data points a', b', c', etc., by interpolating the values of the Hanning window stored in the memory buffer 141.

FIGURE 8B is a flow chart of the steps performed by the window generation routine 196 (FIGURE 3). Beginning at step 420, it is determined which resampled input vocal signal is to be used to create the harmony note. For example, assume a user has set

the gender controls to +10% and -10%. When using the musical effect 100, the user selects which resampled input vocal signal will be used to create a harmony note. The user can specify that the input vocal signal that is resampled at a rate of +10% is used to create a first harmony note, and the input vocal signal that is resampled at a rate of -10% is used to create the other harmony notes, etc. Once the DSP has determined which resampled

input vocal signal is to be used in creating the harmony notes, the length of the window function is initially set to equal twice the period of the associated resampled input signal (expressed in samples) at step 422. Next, the pitch of the harmony note to be produced is compared with the pitch of the resampled input signal at step 424. If the pitch of the harmony note is greater than the pitch of the resampled input note, the DSP proceeds to step 426. At step 426, the DSP determines the number of semitones (x) the harmony note is above a positive threshold. In the presently preferred embodiment of the invention, the positive threshold is set to zero semitones. At step 428, the length of the memory buffer that stores the Hanning window used to create the harmony note is reduced by multiplying the length calculated at step 422 by the results of the equation

$$2^{-x/12}$$

where x is the number of semitones the harmony note is above the positive threshold. For example, if the harmony note is five semitones above the threshold, the length of the memory buffer is reduced by a factor of 0.75.

If the pitch of the harmony note to be created is below the pitch of the resampled input note, the length of the window may be expanded. At step 430, the DSP determines the number of semitones (x) the harmony note is below a negative threshold. In the presently preferred embodiment, the negative threshold is 24 semitones below the pitch of the input note. If the harmony note is below the threshold, the length of the memory buffer that holds the window function is increased by an amount equal to the results of the equation:

$$2^{+x/12}$$

where x is the number of semitones below the threshold. For example, if the harmony note to be created was 29 semitones below the pitch of the input note, then $x = 5$ and the length of the memory buffer that holds the window function is increased by a factor of 1.33.

At step 434, it is determined whether the length of the window function has been increased to an amount that is greater than the amount of memory available to store the

window function. If so, the length of the window function is set to the maximum amount of memory available to store the window function.

If the harmony note to be created is not below the negative threshold, the length of the window function remains the same as was calculated in step 422.

5 After the length of the memory buffer that holds the window function has been calculated, the memory buffer 134 is filled with the values of the window data. This is accomplished by determining, at step 438, a ratio of the length of the buffer 141 (which is currently 256) to the length of the buffer as determined by steps 428 or 432. This ratio is used in step 440 to interpolate the window data. For example, if the new buffer has a length
10 of 284 samples, the buffer 134 is completed by interpolating the data at points 0, .9, 1.8, 2.7 in the same manner as the input vocal signal is resampled as shown in FIGURES 4A, 4B and described above.

A user can also specify a volume ratio for each harmony note produced. This volume ratio affects the magnitude of the samples stored in the memory buffer 134. If the user
15 wants full volume for the harmony notes, the ratio is set to one. If the user wants half the volume, the ratio is set to 0.5. The volume ratio is determined at step 440 and each value in the memory buffers 134 is multiplied by the volume ratio at step 442.

Returning to FIGURE 3, the output of the pitch shifting routine 200 is supplied to a summation block 210 where the output is added to the dry audio signal stored in the memory
20 buffer 122. The combination of the dry audio signal and harmony signals is supplied to a digital-to-analog converter 215 that produces a multi-voice analog signal that is the combination of the input note and harmony notes. As is described in the '671 patent, the output harmony notes are not produced if the pitch recognition routine detects that a user has sung a sibilant sound. Sibilant sounds are sounds such as "s," "ch," "sh," etc. In order for
25 the harmony notes to sound realistic, the pitch of these signals is not shifted. If the pitch recognition routine detects that the user has sung a sibilant sound, the microprocessor sets all the harmonies to be produced to be the same pitch of the input vocal signal. Thus, the harmony notes will all have the same pitch as the input vocal signal, but they will sound slightly different than the input signal due to the timbre shift that occurs due to the combined
30 operation of the resampling and the operation of the pitch shifting routine 200.

To produce more natural sounding harmonies than could be obtained using prior art pitch shifting techniques, the present invention replicates a portion of the resampled input vocal signal that is already pitch and timbre shifted as a result of the resampling. Turning now to FIGURE 5, the pitch shifting routine 200 performed by the digital signal processor 180 is accomplished using the series of harmony generators 220, 230, 240 and 250. Each harmony generator produces one harmony note that is mixed with the dry audio signal stored in the memory buffer 122. The harmony notes to be created are supplied to the digital signal processor on a lead 162 and stored in a look up table 260. The look up table within the digital signal processor is used to determine the fundamental frequency for each of the harmony notes.

Each harmony generator within the digital signal processor produces one of the harmony notes stored in the look up table 260. As described above, the harmony generators scale one of the resampled input vocal signals with the Hanning window stored in the harmony generator's associated memory buffer 134a, 134b, 134c, or 134d, at a rate equal to the fundamental frequency of the harmony note to be created.

The dry audio signal and the output signal of each of the harmony generators 220, 230, 240 and 250 is supplied to the summation block 210 that divides the signals between left and right channels. For example, the output of harmony generator 220 is supplied to a mixer 224. The mixer allows the user to direct the harmony produced to either a left or right audio channel or to a mix of the right and left audio channels. Similarly, the outputs of the harmony generators 230, 240 and 250 are fed to corresponding mixers 234, 244 and 254. Each of the mixers feeds a summation block 270 that combines all the harmony signals for the left channel. Similarly, each of the mixers 224, 234, 244 and 254 feeds a summation block 272 that combines all the harmony signals for the right audio channel.

The digital signal processor also reads the dry audio signal from the memory buffer 122 and applies it to a mixer 284 that can be operated by the user to direct the dry audio to the some combination of the left and/or right audio channels.

Although the digital signal processor 180 is shown including four harmony generators, those skilled in the art will recognize that more or less harmony generators could be provided depending upon the memory available and processing speed of the digital signal processor.

Turning now to FIGURE 6, the details of the functions performed by each of the harmony generators are shown. Each of the harmony generators includes a plurality of windowed audio generators 300, 310, 320 and 330. Each windowed audio generator operates to scale the resampled input vocal signal by the Hanning window as described above. A timer 340 within the windowed audio generator is supplied with a value equal to the fundamental frequency of the harmony note to be produced. The fundamental frequency is determined from the look up table 260 (shown in FIGURE 5) that correlates each harmony note with its corresponding fundamental frequency. When the timer 340 counts down to zero, a signal is sent to a windowed audio generator allocation block 350 that looks for one of the windowed audio generators 300, 310, 320 or 330 to begin the scaling process. For example, if the windowed audio generator 300 is not in use, a buffer pointer 302 is first loaded with the value of the period marker that marks the location in the memory buffer 128 where a complete cycle of the resampled input vocal signal that is to be used in creating the harmony signal begins. Next a window pointer 304 is loaded with a pointer to the beginning of the harmony generator's associated memory buffer 134a, 134b, 134c, or 134d (FIGURE 5). Finally a counter 306 is loaded with the number of samples that are used to store the selected window function. The number of samples in the window function is supplied by the digital signal processor to the harmony generators and is stored in a memory location 370 for use by all the windowed audio generators.

After the buffer pointer 302, the window pointer 304, and counter 306 are initialized, the windowed audio generator then begins a point-by-point multiplication of the resampled input vocal signal stored in the associated memory buffer 128 and the Hanning window stored in associated memory buffer. The result of the multiplication is applied to a summation block 372 that adds the output from all the windowed audio generators 300, 310, 320 and 330. After the multiplication is completed, the pointers 302 and 304 are advanced and the counter 306 is decremented. When the counter 306 reaches zero and all the multiplications have been performed, the windowed audio generator signals the windowed audio generator allocation block 350 that it is available to be used again. The windowed audio generators 310, 320 and 330 operate in the same manner as the windowed audio generator 300.

The timer 340, the period markers stored in the memory location 262 (FIGURE 5), the number of points in the window function stored in the memory location 370, and the Hanning windows stored in the memory locations 134 are all dynamically updated as the user sings different notes into the microphone.

5 As described above, for harmony notes having a pitch below the pitch of the input vocal signal, the Hanning window is calculated to have a length equal to, or longer than, twice the period of the input signal used to create the harmony signal. Therefore, to create a harmony signal that is an octave below the input vocal signal, only one windowed audio generator is needed. However, to create harmony notes having a pitch greater than the pitch
10 of the input vocal note, the length of the Hanning window is shortened. Therefore, to produce an output signal that is above the pitch of the resampled input vocal signal requires only two windowed audio generators.

The musical effect generator described above applies a fixed amount of timbre shift to a pitch shifted note. However, it is possible to dynamically vary the amount of timbre shift
15 to further increase the realism of a digitally processed note.

As indicated above, the musical effect generator of the present invention can be used with a karaoke system that has a prerecorded melody and/or harmony track. Alternatively, the melody or harmony notes can be received from a keyboard or from a computer. Typically, the prerecorded melody or harmony notes are transmitted to the effect generator
20 over a MIDI channel. If only one harmony voice is to be produced, the effect generator can read the desired harmony notes from the MIDI port, look up the amount of timbre shift that is to be applied to a note and create the harmony note by replicating portions of the resampled input note in the manner previously described. However, if more than one harmony voice is to be produced then it is usually required that the notes for each voice be
25 transmitted on their own MIDI channel.

In most instances, the MIDI controller supplying the harmony notes does not have enough free channels to allow a separate channel to be used for each voice. A single MIDI channel could be used to define each melody or harmony note to be produced. However, there is no practical way to inform the effect generator how much timbre shift should be
30 applied to an individual melody or harmony note. Conceptually, it would be possible to code the MIDI file that describes the harmony or melody notes with a MIDI message that precedes

each note and defines how much timbre shift to apply. However, such a file would be difficult to construct and could not be constructed in real time if the melody/harmony notes were being coded by a keyboard as the user sang. Therefore, there is a need for an effect generator that can receive the melody or harmony notes on a single MIDI channel and assign different amounts of timbre shift to the notes that comprise the various voices.

A first alternative embodiment of the present invention is shown in FIGURE 9A. In this embodiment, all the melody or harmony notes that are to accompany a given song are encoded on a single MIDI channel. The effect generator is programmed to read the notes and dynamically assign the amount of timbre shift to the notes in real time. The hardware required to implement this embodiment of the invention is the same as shown in FIGURE 3 and described above. However, the digital signal processor 180 is programmed in a slightly different manner.

The effect generator 500 receives a stream of melody or harmony notes on a single MIDI channel 505 from a MIDI karaoke system, a keyboard or computer system as the user sings. The melody or harmony notes are read by the digital signal processor and are automatically assigned an amount of timbre shift in a process block 515. Preferably, the automatic timbre assignment block 515 is implemented by programming the digital signal processor to compare the pitch of the melody or harmony note to be produced with one or more pitch thresholds.

Depending on where the pitch of a melody or harmony note is relative to the thresholds, the timbre of the note is set according to some predefined or preprogrammed rule. For example, if there are two thresholds, notes having a pitch higher than both thresholds may be resampled at a rate of -10%, while harmony notes between the thresholds may be resampled at a rate of -2% and harmony notes below both thresholds may be resampled at a rate of +5% etc. Of course, the amount of timbre shift may be the same for notes above or below the one or more pitch thresholds. Alternatively, the musical effect generator may be programmed so that no timbre shift is applied to the notes. The one or more pitch thresholds may be predefined or may be programmed for each song by including the one or more threshold notes as MIDI messages at the beginning of the MIDI file that accompanies the song.

As an alternative to comparing the pitch of the melody or harmony notes with a pitch threshold, the automatic timbre assignment block 515 may be implemented by programming the digital signal processor to compare the pitch of the harmony note to the pitch of a desired melody note that is stored in a separate MIDI file and transmitted to the effect generator on a MIDI channel 510. By reading the desired melody notes, the effect generator can look ahead to determine an expected amount of pitch shift required to produce the harmony note (assuming the singer is close to singing on key). The effect generator may then modify the amount of timbre shift for each harmony note depending on the expected amount of pitch shift.

As yet another alternative, the automatic timbre assigning block 515 may be implemented by programming the digital signal processor to compare the pitch of the harmony notes with the pitch of the input vocal note to determine if the harmony note is above or below the melody line. The timbre of the harmony note can be modified as a function of the difference in pitch between the pitch of the harmony note and the pitch of the input vocal note. Because the harmony notes produced have timbres that differ from the input vocal note, they do not sound like pitch shifted versions of the input note, thereby adding realism to the composite sound.

A second alternative embodiment of the effect generator according to the present invention is shown in FIGURE 9B. Here the timbre of a harmony note is not modified in a manner to differentiate the harmony voices from the input voice but is modified in a way that mimics how a singer's voice changes as the singer sings higher or lower notes.

The musical effect generator 520 receives an input vocal signal from a singer and analyzes the signal to determine its pitch. The effect generator receives a stream of desired melody or harmony notes on a MIDI channel 530 that indicate the pitch to which the input vocal signal should be shifted. The digital signal processor within the effect generator dynamically assigns an amount of timbre shift to a note to be produced as represented by the block 540. Preferably the digital signal processor compares the pitch of the desired note with the pitch of the input vocal signal to select how much timbre shift should be applied to the pitch shifted output note. For example, the amount of timbre shift may vary linearly with the difference in pitch between the input vocal signal and the desired harmony or melody note. Alternatively, step function may be used whereby the timbre doesn't change until the

pitch of the desired note differs from the pitch of the input vocal signal by more than some predetermined amount. Once the amount of timbre shift has been determined, the digitized input vocal signal is resampled and the output note is created by replicating portions of the resampled input note at a rate equal to the fundamental frequency of the desired output note as described above.

To achieve a realistic timbre shift that mimics the physical changes that take place in a singer's vocal tract, the resampling rate should be slower than the original sampling rate for notes that have pitches higher than the input vocal note. Conversely, the resampling rate should be faster than the original sampling rate for notes having a pitch below the input vocal note. As an alternative to changing the timbre of a note based on the amount of pitch shift required, it is also possible to vary the timbre based on changes in the loudness of the input vocal signal. The digital signal processor analyzes the magnitude of the digitized input vocal signal and selects an amount of timbre shift as a function of the magnitude. Furthermore, the timbre could be changed depending upon the length of time the input vocal signal has been sung. Once the effect generator has determined the pitch of the input vocal signal, the digital signal processor starts an internal timer that keeps track of the length of time the pitch remains within some redefined limits. The amount of timbre shift is selected as a function of the length of time recorded by the timer. As will be appreciated by those skilled in the art, many different criteria could be used for controlling the amount of timbre shift to be applied to note.

Using the effect generator shown in FIGURE 9B, the composite output signal sounds more realistic because the notes simulate the way the timbre of the note changes naturally in a singer's voice as the pitch of a sung note is varied.

Although the present invention has been described with respect to vocal harmony generators, the present invention also has other uses. One example is as a voice disguiser, where a user would speak into a microphone and an output signal having a different timbre and/or pitch would be produced. If the output signal had a frequency one octave below the input signal, a device could be built wherein the amount of pitch shift used in data resampling is fixed and that requires only one windowed audio generator. Such a device would be useful for law enforcement to disguise the voice of witnesses or as part of an answering machine to conceal the voice of the user. Alternatively, the present invention

could be used by radio announcers who want their voice to sound deeper. In addition, the invention can be used with input notes that are received from musical instruments. The result of the timbre shifting combined with pitch shifting allows one instrument to sound like another.

5 Additionally, the preferred embodiment of the invention first employs the resampling pitch shifting followed by the pitch shifting according to the Lent method. It will be appreciated that the reverse process could also be used, whereby the output signals created using the Lent method are stored in a memory buffer and resampled at a new rate to further shift the pitch. Each of the methods, Lent and pitch shifting by resampling, operate as
10 previously described. There are two issues to keep in mind when implementing the steps in the reverse order. First, the output of the pitch shifter that operates according to the Lent method no longer directly controls the fundamental frequency of the overall output signal. Therefore, it is necessary to compensate for the pitch shift which occurs as a result of the resampling. For example, if the timbre shift control was set to make a singer sound more
15 female, the resampling pitch shifter might adjust the pitch upwards by, say, 12%. If it was desired to produce a timbre shifted output signal at a frequency of 440 Hz, then the pitch shifter that operates according to the Lent method would have to be set to output a signal with a fundamental frequency of $440/1.12 = 392.86$ Hz. In general, the relation is:

$$\text{TSF} = \text{LF} * \text{PSR}$$

20 where:

TSF = the frequency of the fundamental pitch of the timbre shifted output signal;

25 LF = the frequency of the fundamental pitch of the output signal of the pitch shifter that operates according to the Lent method; and

PSR = the Pitch Shift Ratio of the resampling pitch shifter. This is the ratio of (input sample rate)/(resampled sample rate).

30 The second issue is that the clock source for the harmony timer 340 as shown in Figure 6 will be different. When the Lent method pitch shifter is the last step in the process, this timer is decremented at the sample rate of the system, for example 44.1 KHz in a system

providing CD quality audio. This guarantees that the Lent method pitch shifter can provide a continuous stream of pitch shifted audio at that rate. When the Lent method pitch shifter passes its output to the resampling pitch shifter, rather than directly to the output, the timer 340 is clocked at the resampling rate. This ensures that the two processes operate in synchrony. If the resampling is occurring at a higher rate, as in Figure 4A, the Lent method must be producing replicated pitch periods at a higher rate to keep the resampling pitch shifter continuously supplied with data. Similarly, if the resampling is occurring at a lower rate, as in Figure 4B, the Lent method need only produce replicated pitch periods at a lower rate to keep the resampling pitch shifter continuously supplied with data.

While the preferred embodiment of the invention has been illustrated and described, it will be appreciated that various changes can be made therein without departing from the spirit and scope of the invention. Therefore, the scope of the invention is to be determined solely from the following claims.



The embodiments of the invention in which an exclusive property or privilege is claimed are defined as follows:

1. A method of creating a timbre shifted output signal from an input signal, comprising the steps of:

receiving a digital representation of an input signal that has been sampled at a first rate;

resampling said digital representation of the input signal at a second rate that differs from the first rate; and

creating a digital representation of the timbre shifted output signal by periodically extracting a segment of the resampled input signal and replicating the extracted segments at a rate equal to a fundamental frequency of the output signal.

2. A method of creating a timbre, pitch shifted output signal from an input signal, comprising the steps of:

receiving a digital representation of an input signal that has been sampled at a first rate;

creating a digital representation of a pitch shifted output signal by periodically extracting a segment of the input signal and replicating the extracted segments at a rate equal to a fundamental frequency of the pitch shifted output signal; and

creating a digital representation of the timbre shifted output signal by resampling the digital representation of the pitch shifted output signal at a second rate that differs from the first rate.

3. The method of Claim 1 or 2, further comprising the step of applying the digital representation of the timbre shifted output signal to an digital-to-analog converter to convert the digital representation of the timbre shifted output signal to an analog representation of the timbre shifted output signal.

4. The method of Claim 1, wherein the input signal has a fundamental frequency and wherein the timbre shifted output signal has a fundamental frequency that is the same as the fundamental frequency of the input signal.

5. The method of Claim 1, wherein the input signal has a fundamental frequency and wherein the timbre shifted output signal has a fundamental frequency that is different than the fundamental frequency of the input signal.



6. The method of Claims 1 or 2, wherein the step of receiving a digital representation of the input signal comprises the steps of:

receiving an analog representation of the input signal; and

applying the analog representation of the input signal to an analog-to-digital converter to convert the analog representation of the input signal to a digital representation of the input signal.

7. The method of Claims 1 or 2, wherein the input signal is a note produced by a musical instrument.

8. The method of Claims 1 or 2, wherein the input signal is a vocal note.

9. A method of generating a timbre shifted output vocal signal from an input vocal signal, comprising the steps of:

receiving a digital representation of the input vocal signal that has been sampled at a first rate and resampling the digital representation of the input vocal signal at a second sampling rate different from the first sampling rate to create a resampled input vocal signal; and

creating a digital representation of the timbre shifted output vocal signal by periodically extracting a segment of the resampled input vocal signal with a window function and replicating the extracted segments at a rate equal to a fundamental frequency of the output vocal signal.

10. A method of generating a timbre, pitch shifted output vocal signal from an input vocal signal, comprising the steps of:

receiving a digital representation of the input vocal signal that has been sampled at a first rate and creating a digital representation of the pitch shifted output vocal signal by periodically extracting a segment of the input vocal signal with a window function and replicating the extracted segments at a rate equal to a fundamental frequency of the output vocal signal; and

creating a digital representation of the timbre shifted output vocal signal by resampling the digital representation of the pitch shifted output vocal signal at a second sampling rate that is different from the first sampling rate.

11. The method of Claims 9 or 10, wherein the step of receiving a digital representation of the input vocal signal comprises the steps of:



receiving an analog representation of the input vocal signal; and

applying the analog representation of the input vocal signal to an analog-to-digital converter to convert the analog representation of the input vocal signal to a digital representation of the input vocal signal.

12. The method of Claims 9 or 10, further comprising the steps of applying the digital representation of the timbre shifted output vocal signal to a digital-to-analog converter to convert the digital representation of the timbre shifted output vocal signal to an analog representation of the timbre shifted output vocal signal.

13. The method of Claims 9 or 10, wherein the input vocal signal has a fundamental frequency and the timbre shifted output vocal signal has a fundamental frequency that is the same as the fundamental frequency of the input vocal signal.

14. The method of Claims 9 or 10, wherein the input vocal signal has a fundamental frequency and the timbre shifted output vocal signal has a fundamental frequency that is different from the fundamental frequency of the input vocal signal.

15. The method of Claim 9, wherein the input vocal signal and the timbre shifted output vocal signal have a fundamental frequency and wherein the step of extracting a segment of the resampled input vocal signal further comprises the steps of:

generating a window function having a duration that is a function of a difference between the fundamental frequency of the input vocal signal and the fundamental frequency of the timbre shifted output vocal signal; and

multiplying the window function and the digital representation of the resampled input vocal signal together.

16. The method of Claim 10, wherein the input vocal signal and the pitch shifted output vocal signal have a fundamental frequency and wherein the step of extracting a segment of the input vocal signal further comprises the steps of:

generating a window function having a duration that is a function of a difference between the fundamental frequency of the input vocal signal and the fundamental frequency of the pitch shifted output vocal signal; and

multiplying the window function and the digital representation of the pitch shifted output vocal signal together.



17. The method of Claim 9, wherein the digital representation of the input vocal signal is stored in a digital memory prior to resampling and wherein the digital representation of the input vocal signal comprises a number of cycles, each cycle occupying a number of memory locations, the step of resampling the digital representation of the input vocal signal further comprising the steps of:

storing the resampled input vocal signal in a larger number of memory locations per cycle than are occupied by the digital representation of the input vocal signal if the second sampling rate is faster than the first sampling rate; and

storing the resampled input vocal signal in a fewer number of memory locations per cycle than are occupied by the digital representation of the input vocal signal if the second sampling rate is slower than the first sampling rate.

18. The method of Claim 10, wherein the digital representation of the pitch shifted output vocal signal is stored in a digital memory prior to resampling and wherein the digital representation of the pitch shifted output signal comprises a number of cycles, each cycle occupying a number of memory locations, the step of resampling the digital representation of the pitch shifted output vocal signal further comprises the steps of:

storing the resampled pitch shifted output vocal signal in a larger number of memory locations per cycle than are occupied by the digital representation of the pitch shifted output vocal signal if the second sampling rate is faster than the first sampling rate; and

storing the resampled pitch shifted output vocal signal in a fewer number of memory locations per cycle than are occupied by the digital representation of the pitch shifted output vocal signal if the second sampling rate is slower than the first sampling rate.

19. The method of Claim 9, wherein the step of resampling the input vocal signal is performed by interpolating the digital representation of the input vocal signal.

20. The method of Claim 19, where the step of interpolating the digital representation of the input vocal signal is performed using a linear interpolation.

21. The method of Claim 10, wherein the step of resampling the pitch shifted output vocal signal is performed by interpolating the digital representation of the pitch shifted output vocal signal.

22. The method of Claim 21, where the step of interpolating the digital representation of the pitch shifted output vocal signal is performed using a linear interpolation.

23. An apparatus for producing a timbre shifted output signal from an input signal, comprising:

a digital memory;

a digital signal processor for receiving a digital representation of the input signal that has been sampled at a first rate and for storing the digital representation of the input signal in the digital memory;

means for resampling the digital representation of the input signal that is stored in the digital memory at a second rate that differs from the first rate, and for storing the resampled input signal in the digital memory; and

a pitch shifter for creating a digital representation of the timbre shifted output signal by periodically extracting a segment of the resampled input signal and replicating the extracted segments at a rate equal to a fundamental frequency of the timbre shifted output signal.

24. The apparatus of Claim 23, further comprising:

a microphone for converting the input signal into a corresponding electrical input signal; and

an analog-to-digital converter for sampling the electrical input signal at the first rate and converting the electrical input signal into the digital representation of the input signal.

25. The apparatus of Claim 23, further comprising:

a digital-to-analog converter to convert the digital representation of the timbre shifted output signal to an analog representation of the timbre shifted output signal.

26. The apparatus of Claim 23, further comprising a control for varying the second rate at which the input signal is resampled.

27. The apparatus of Claim 23, wherein the pitch shifter extracts a segment of the resampled input signal by scaling the resampled input signal with a window function.

28. The apparatus of Claim 27, wherein pitch shifter scales the resampled input signal with the window function at a rate that is harmonic related to the fundamental frequency of the input signal.

29. The apparatus of Claim 27, wherein the input signal has a fundamental frequency and the timbre shifted output signal has a fundamental frequency and wherein the pitch shifter further comprises:

SUBSTITUTE SHEET





means for adjusting a duration of the window function based upon a difference between the fundamental frequency of the input signal and the fundamental frequency of the timbre shifted output signal.

30. The apparatus of Claim 29, wherein the means for adjusting the duration of the window function decreases the duration of the window if the fundamental frequency of the timbre shifted output signal is greater than the fundamental frequency of the input signal and increases the duration of the window function if the fundamental frequency of the timbre shifted output signal is less than the fundamental frequency of the input signal.

31. A system for creating a timbre shifted and/or pitch shifted output signal from an input signal, comprising:

means for receiving a digital representation of the input signal that has been sampled at a first rate;

means for receiving a first reference note that defines a first desired fundamental frequency of the timbre shifted output signal;

a comparator that performs an analysis of said reference note and selects a resampling rate as a function of said analysis;

a digital signal processor that resamples the digital representation of the input signal at the selected resampling rate; and

a pitch shifter for creating the timbre shifted output signal by periodically extracting a segment of the resampled input signal and replicating the segments at a rate equal to the fundamental frequency of the reference note.

32. The system of Claim 31, wherein the comparator analyzes the reference note by comparing the fundamental frequency of the reference note with one or more thresholds.

33. The system of Claim 31, further comprising:

means for determining the fundamental frequency of the input signal;

wherein the comparator analyzes the reference note by comparing said fundamental frequency of the reference note with said fundamental frequency of the input signal and selects the resampling rate as a function of the difference between said fundamental frequency of the reference note and said fundamental frequency of the input signal.



34. The system of Claim 31, further comprising:

means for receiving a second reference note that defines a second fundamental frequency;

wherein the comparator analyzes the reference note by comparing said fundamental frequency of said first reference note with said fundamental frequency of the second reference note and selects the resampling rate as a function of the difference between the fundamental frequency of the reference note and the fundamental frequency of the second reference note.

35. A system for creating a timbre shifted and/or pitch shifted output signal from an input signal, comprising:

means for receiving a digital representation of the input signal that has been sampled at a first rate;

means for receiving a reference note that defines a desired fundamental frequency of the timbre shifted output signal;

means for calculating the length of time the input signal has been received;

a comparator that analyzes said length of time for which the input signal has been received and selects a resampling rate as a function of said length of time;

a digital signal processor for resampling the digital representation of the input signal at said selected resampling rate; and

a pitch shifter for creating the timbre shifted output signal by extracting a segment of the resampled input signal and replicating the segments at a rate substantially equal to the fundamental frequency of said reference note.

36. A system for creating a timbre shifted and/or pitch shifted output signal from an input signal, comprising:

means for receiving a digital representation of an input signal that has been sampled at a first rate;

means for receiving a reference note that defines a desired fundamental frequency of the timbre shifted output signal;

a comparator that analyzes a magnitude of the digital representation of the input signal and selects a resampling rate as a function of the magnitude;



WO 96/22592

PCT/CA96/00026

-32-

a digital signal processor that resamples the digital representation of the input signal at the selected resampling rate; and

a pitch shifter for creating the timbre shifted output signal by periodically extracting a segment of the resampled input signal and replicating the segments at a rate substantially equal to said fundamental frequency of the reference note.

SUBSTITUTE SHEET