Chapter 10

Spectral Processing

X. Amatriain, J. Bonada, A. Loscos, X. Serra

10.1   Introduction

In the context of this book, we are looking for representations of sound signals and signal processing systems that can provide ways to design sound transformations in a variety of music applications and contexts. It should have been clear throughout the book, that several points of view have to be considered, including a mathematical, thus objective perspective, and a cognitive, thus mainly subjective, standpoint. Both points of view are necessary to fully understand the concept of sound effects and to be able to use the described techniques in practical situations.

The mathematical and signal processing points of view are straightforward to present, although not necessarily easily, since the language of the equations and of flow diagrams is suitable for them. However, the top-down implications are much harder to express due to the huge number of variables involved and to the inherent perceptual subjectivity of the music making process. This is clearly one of the main challenges of the book and the main reason for its existence.

The use of a spectral representation of a sound yields a perspective that is *sometimes* closer to the one used in a sound engineering approach. By understanding the basic concepts of frequency domain analysis, we are able to acquire the tools to use a large number of effects processors and to understand many types of sound transformation systems. Moreover, being the frequency domain analysis a somewhat similar process to the one performed by the human hearing system, it yields fairly intuitive intermediate representations.

The basic idea of spectral processing is that we can analyze a sound to obtain alternative frequency domain representations, which can then be transformed and inverted to produce new sounds (see Fig.10.1). Most of the approaches start by developing an analysis/synthesis system from which the input sound is reconstructed

without any perceptual loss of sound quality. The techniques described in chapters 8 and 9 are clear examples of this approach. Then the main issue is what is the intermediate representation and what parameters are available for applying the desired transformations.
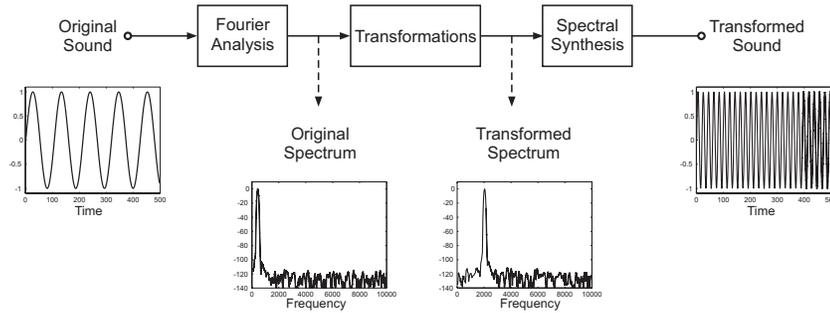

Figure 10.1 Block diagram of a simple spectral processing framework.

Perceptual or musical concepts such as timbre or pitch are clearly related to the spectral characteristics of a sound. Even some very common processes for sound effects are better explained using a frequency domain representation. We usually think on the frequency axis when we talk about equalizing, filtering, pitch shifting, harmonizing ... In fact, some of them are specific to this signal processing approach and do not have an immediate counterpart on the time domain. On the other hand, most (but not all) of the sound effects presented in this book can be implemented in the frequency domain.

Another issue is whether or not this approach is the most efficient, or practical, for a given application. The process of transforming a time domain signal into a frequency domain representation is, by itself, not an immediate step. Some parameters are difficult to adjust and force us to take several compromises. Some settings, such as the size of the analysis window, have little or nothing to do with the high-level approach we intend to favor, and require the user to have a basic signal processing understanding.

In that sense, when we talk about higher level spectral processing we are thinking of an intermediate analysis step in which relevant features are extracted or computed from the spectrum. These relevant features should be much closer to a musical or high-level approach. We can then process the features themselves or even apply transformations that keep the features unchanged. For example, we can extract the fundamental frequency and the spectral shape from a sound and then modify the fundamental frequency without affecting the shape of the spectrum.

Assuming the fact that there is no single representation and processing system optimal for everything, our approach will be to present a set of complementary spectral models that can be combined to be used for the largest possible set of sounds and musical applications.

In section 10.2 we introduce two spectral models: *Sinusoidal and Sinusoidal plus Residual*. These models already represent a step up on the abstraction ladder and
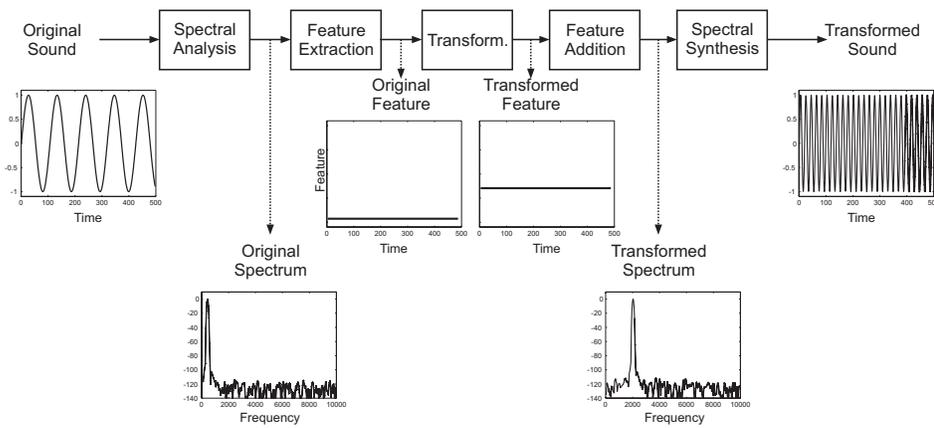
Figure 10.2 Block diagram of a higher-level spectral processing framework.

from either of them, we can identify and extract higher-level information of a sound, such as: harmonics, pitch, spectral shape, vibrato, or note boundaries, that is *Higher Level Features*. This analysis step brings the representation closer to our perceptual understanding of a sound. The complexity of the analysis will depend on the type of feature that we wish to identify and the sound to analyze. The benefits of going to this higher level of analysis are enormous and open up a wide range of new musical applications.

Having set the basis of the *Sinusoidal plus Residual* model, we will then give some details of the techniques used both in its analysis and synthesis process, providing Matlab code to implement an analysis-synthesis framework in section 10.3. This Matlab implementation is based on the *Spectral Modeling Synthesis* framework. SMS [SMS] is a set of spectral based techniques and related implementations for the analysis/transformation/synthesis of an audio signal based on the scheme presented in Fig. 10.2.

In section 10.4 we will provide a set of basic audio effects and transformations based on the implemented *Sinusoidal plus Residual* analysis/synthesis. Matlab code is provided for all of them.

We will finish with an explanation of content dependent processing implementations. In section 10.5.1 we introduce a real-time singing voice conversion application that has been developed for use in Karaoke, and in section 10.5.2 we define the basis of a nearly lossless time scaling algorithm. The complexity and extension of these implementations prevent us from providing the associated Matlab code, so we leave that task as a challenge for advanced readers.

10.2 Spectral Models

The most common approach for converting a time domain signal into its frequency domain representation is the *Short-Time Fourier Transform (STFT)*. It is a general

technique from which we can implement lossless analysis/synthesis systems. Many sound transformation systems are based on direct implementations of the basic algorithm and several examples have been presented in chapter 8.

In this chapter, we will briefly mention the *Sinusoidal Model* and will concentrate, with a Matlab sample code, on the *Sinusoidal plus Residual Model*. Anyhow, the decision as to what spectral representation to use in a particular situation is not an easy one. The boundaries are not clear and there are always compromises to take into account, such as: (1) sound fidelity, (2) flexibility, (3) coding efficiency, and (4) computational requirements. Ideally, we want to maximize fidelity and flexibility while minimizing memory consumption and computational requirements. The best choice for maximum fidelity and minimum computation time is the STFT that, anyhow, yields a rather inflexible representation and inefficient coding scheme. Thus our interest in finding higher-level representations as the ones we present in this section.

10.2.1   Sinusoidal Model

Using the output of the *STFT*, the *Sinusoidal* model represents a step towards a more flexible representations while compromising both sound fidelity and computing time. It is based on modeling the time-varying spectral characteristics of a sound as sums of time-varying sinusoids. The input sound $s(t)$ is modeled by,

$$s(t) = \sum_{r=1}^{R} A_r(t) \cos[\theta_r(t)] \qquad (10.1)$$

where $A_r(t)$ and $\theta_r(t)$ are the instantaneous amplitude and phase of the r^{th} sinusoid, respectively [McA86, Smi87].

To obtain a sinusoidal representation from a sound, an analysis is performed in order to estimate the instantaneous amplitudes and phases of the sinusoids. This estimation is generally done by first computing the STFT of the sound, as described in chapter 8, then detecting the spectral peaks (and measuring the magnitude, frequency and phase of each one), and finally organizing them as time-varying sinusoidal tracks.

It is a quite general technique that can be used in a wide range of sounds and offers a gain in flexibility compared with the direct STFT implementation.

10.2.2   Sinusoidal plus Residual Model

The *Sinusoidal plus Residual* model can cover a wide "compromise space" and can in fact be seen as the generalization of both the *STFT* and the *Sinusoidal* models. Using this approach, we can decide what part of the spectral information is modeled as *sinusoids* and what is left as *STFT*. With a good analysis, the *Sinusoidal plus Residual* representation is very flexible while maintaining a good sound fidelity, and the representation is quite efficient. In this approach, the *Sinusoidal* representation is

used to model only the stable partials of a sound. The residual, or its approximation, models what is left, which should ideally be a stochastic component. This model is less general than either the *STFT* or the *Sinusoidal* representations but it results in an enormous gain in flexibility [Ser89, Ser90, Ser96].

The input sound $s(t)$ is modeled by,

$$s(t) = \sum_{r=1}^{R} A_r(t) \cos[\theta_r(t)] + e(t) \tag{10.2}$$

where $A_r(t)$ and $\theta_r(t)$ are the instantaneous amplitude and phase of the r^{th} sinusoid, respectively, and $e(t)$ is the noise component at time t (in seconds).

The sinusoidal plus residual model assumes that the sinusoids are stable partials of the sound with a slowly changing amplitude and frequency. With this restriction, we are able to add major constraints to the detection of sinusoids in the spectrum and omit the detection of the phase of each peak. The instantaneous phase that appears in the equation is taken to be the integral of the instantaneous frequency $\omega_r(t)$, and therefore satisfies

$$\theta_r(t) = \int_0^t \omega_r(\tau)d\tau \tag{10.3}$$

where $\omega(t)$ is the frequency in radians, and r is the sinusoid number. When the sinusoids are used to model only the stable partials of the sound, we refer to this part of the sound as the deterministic component.

Within this model we can either leave the residual signal, $e(t)$, to be the difference between the original sound and the sinusoidal component, resulting into an identity system, or we can assume that $e(t)$ is a stochastic signal. In this case, the residual can be described as filtered white noise,

$$e(t) = \int_0^t h(t,\tau)u(\tau)d\tau \tag{10.4}$$

where $u(t)$ is white noise and $h(t,\tau)$ is the response of a time varying filter to an impulse at time t. That is, the residual is modeled by the time-domain convolution of white noise with a time-varying frequency-shaping filter.

The implementation of the analysis for the Sinusoidal plus Residual Model is more complex than the one for the Sinusoidal Model. Figure 10.3 shows a simplified block diagram of this analysis.

The first few steps are the same as those in a sinusoidal-only analysis. The major differences start on the peak continuation process since in order to have a good partial-residual decomposition we have to refine this peak-continuation process in such a way as to be able to identify the stable partials of the sound. Several strategies can be used to accomplish this. The simplest case is when the sound is monophonic and pseudo-harmonic. By using the fundamental frequency information in the peak continuation algorithm, we can easily identify the harmonic partials.
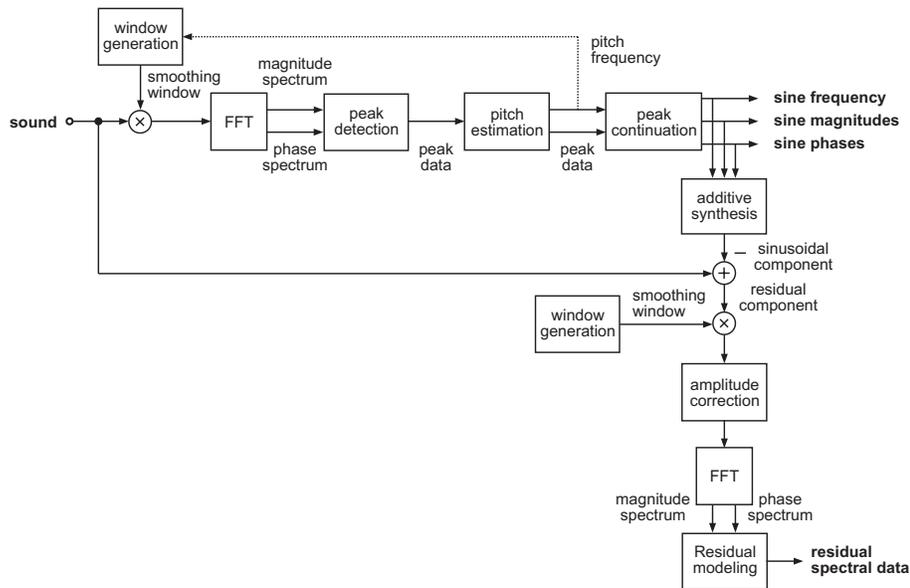
Figure 10.3 Block diagram of the Sinusoidal plus residual analysis.

The residual component is obtained by first generating the sinusoidal component with additive synthesis, and then subtracting it from the original waveform. This is possible because the instantaneous phases of the original sound are matched and therefore the shape of the time domain waveform preserved. A spectral analysis of this time domain residual is done by first windowing it, using a window which is independent of the one used to find sinusoids, and thus we are free to choose a different time-frequency compromise. An amplitude correction step can improve the time smearing produced in the sinusoidal subtraction. Then the FFT is computed and the resulting spectrum can be modeled using several existing techniques. The spectral phases might be discarded if the residual can be approximated as a stochastic signal. Figure 10.4 shows a spectrogram illustrating the sinusoidal and residual components.

The original sinusoidal plus residual model has led to other different spectral models that still share some of its basic principles [Din97, Fit00, Ver00].

10.3   Techniques

It is beyond the scope of this chapter to discuss deeply the whole analysis-synthesis process that results in a Sinusoidal plus Residual representation of the sound, but lets describe in some detail the major steps.
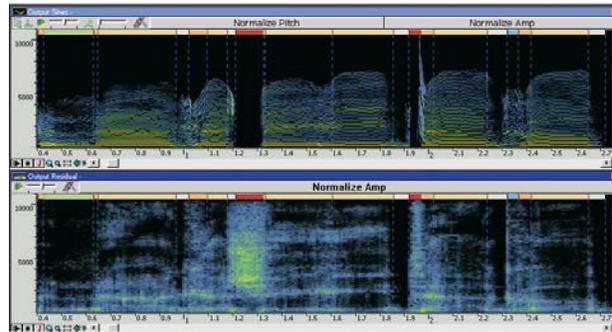
Figure 10.4 Spectrogram: a) sinusoidal component, b) residual spectrum.

10.3.1 Analysis

The analysis step of the Sinusoidal plus Residual model was already presented in the previous section and it is illustrated in Fig. 10.3. Next we will introduce the most important techniques and the basic considerations that need be taken into account when analyzing a sound.

Previous considerations: STFT settings

In this section, we will see that the STFT process is far from being unsupervised, and its settings are indeed critical in order to get a good representation of the sound. The main parameters involved in this step are window size, window type, frame size and hop size.

As has already been mentioned in previous chapters, the first step involved in the process of converting a time domain signal into its frequency domain representation, is the windowing of the sound. This operation involves selecting a number of samples from the sound signal and multiplying their value by a windowing function [Har78].

The number of samples taken in every processing step is defined by the window size. It is a crucial parameter, especially if we take into account that the number of spectral samples that the DFT will yield at its output, corresponds to half the number of samples of its input spread over half of the original sampling rate. We will not go into the details of the DFT mathematics that lead to this property, but it is very important to note that the longer the window, the more frequency resolution we will have. On the other hand, it is straightforward to see the drawback of taking very long windows: the loss of time resolution. This phenomenon is known as the time vs. frequency resolution tradeoff (see Fig. 10.5). A more specific limitation of the window size has to do with choosing windows with odd sample-length in order to guarantee even symmetry about the origin.
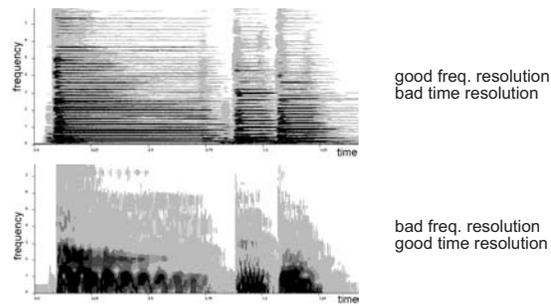
Figure 10.5 Time vs. frequency resolution tradeoff.

The kind of window used also has a very strong effect on the qualities of the spectral representation we will obtain. At this point we should remember that a time domain multiplication (as the one done between the signal and the windowing function), becomes a frequency domain convolution between the Fourier Transforms of each of the signals (see Fig. 10.6). One may be tempted to forget about deciding on these matters and apply no window at all, just taking n samples from the signal and feeding them into the chosen FFT algorithm. Even in this case, though, a rectangular window is being used, so the spectrum of the signal is being convolved with the transform of a rectangular pulse, a *sinc-like* function.
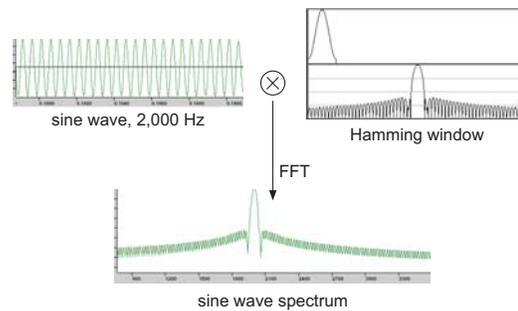

Figure 10.6 Effect of applying a window in the time domain.

Two features of the transform of the window are specially relevant to whether a particular function is useful or not: the width of the main lobe, and the main to highest side lobe relation. The main lobe bandwidth is expressed in bins (spectral samples) and, in conjunction with the window size, defines the ability to distinguish two sinusoidal peaks (see Fig. 10.7). The following formula expresses the relationship the window size M, the main lobe bandwidth B_s and the sampling rate f_s should have in order to distinguish two sinusoids of frequency f_k and f_{k+1}:

$$M \geq B_s \frac{f_s}{|f_{k+1} - f_k|}. \tag{10.5}$$

The amplitude relationship between the main and the highest side lobe explains

two sinusoids of 2.000 Hz and 2.200 Hz

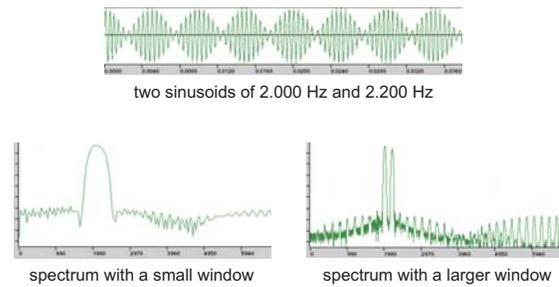spectrum with a small window

spectrum with a larger window

Figure 10.7 Effect of the window size in distinguishing between two sinusoids.

the amount of distortion a peak will receive from surrounding partials. It would be ideal to have a window with an extremely narrow main lobe and a very high main to secondary lobe relation. However, the inherent tradeoff between these two parameters forces a compromise to be taken.

Common windows that can be used in the analysis step are Rectangular, Triangular, Kaiser-Bessel, Hamming, Hanning and Blackmann-Harris. In the Matlab code supplied in this chapter, we have chosen a Blackmann-Harris 92 dB window for the sake of simplicity. This window has a rather wide main lobe (9 bins) but an extremely high main-to-secondary lobe relation of 92 dB. This difference is so close to the dynamic range of a 16-bit representation that, in that case, we need only take into account the influence of the main lobe. The following Matlab function (M-file 10.1) implements the generation of a Blackman-Harris window.

M-file 10.1 (bh92.m)
```
function[bh92SINE2SINE,bh92SINE2SINEsize]=bh92SINE2SINEgeneration;
%function[bh92SINE2SINE,bh92SINE2SINEsize]=bh92SINE2SINEgeneration;
%
% ==> generation of the Blackman-Harris window
% output data:
%     bh92SINE2SINEsize: size of the window
%     bh92SINE2SINE: (sampled) window

bh92SINE2SINEsize = 4096;
bh92SINE2SINE = zeros(bh92SINE2SINEsize,1);
bh92N = 512;
bh92const = [.35875, .48829, .14128, .01168];
bh92Theta = -4*2*pi/bh92N;
bh92ThetaIncr = 8*2*pi/bh92N/bh92SINE2SINEsize;
for i=1:bh92SINE2SINEsize
    for m=0:3
        bh92SINE2SINE(i)=bh92SINE2SINE(i)-bh92const(m+1)/2*...
            (sine2sine(bh92Theta-m*2*pi/bh92N,bh92N)+...
            sine2sine(bh92Theta+m*2*pi/bh92N,bh92N));
```

```
   end;
   bh92Theta = bh92Theta + bh92ThetaIncr;
end;
bh92SINE2SINE = bh92SINE2SINE/bh92SINE2SINE(bh92SINE2SINEsize/2+1);
```

The value of the `sine2sine` function (not included in the basic Matlab package) is computed as follows:

M-file 10.2 (`sine2sine.m`)
```
function x = sine2sine( x , N )
% sine2sine function !!!
x = sin((N/2)*x) / sin(x/2);
```

One may think that a possible way of overcoming the time/frequency tradeoff is to add zeros to the windowed signals in order to have a longer FFT and so increase the frequency resolution. This process is known as *zero-padding* and it represents an interpolation in the frequency domain. Thus, when we zero-pad a signal before the DFT process, we are not adding any information to its frequency representation (we will still not distinguish two sinusoids if Eq. (10.5) is not satisfied), but we are indeed increasing the frequency resolution by adding intermediate interpolated bins. This process can help in the peak detection process, as explained later.

A final step is the circular shift already described in section 8.2.2. This *buffer centering* guarantees the preservation of zero-phase conditions in the analysis process.

Once the spectrum of a *frame* has been computed, the window must move to the next position in the waveform in order to take the next set of samples. The distance between the centers of two consecutive windows is known as *hop size*. If the hop size is smaller than the window size, we will be including some overlap, that is, some samples will be used more than once in the analysis process. In general, the more overlap, the smoother the transitions of the spectrum will be across time, but that is a computationally expensive process. The window type and the hop size must be chosen in such a way that the resulting envelope adds approximately to a constant, following the equation

$$A_w(m) \equiv \sum_{n=-\infty}^{\infty} w(m - nH) \approx \text{constant}. \tag{10.6}$$

A measure of the deviation of A_w from a constant is the difference between the maximum and minimum values for the envelope as a percentage of the maximum value:

$$d_w = 100 \times \frac{\max_m[A_w(m)] - \min_m[A_w(m)]}{\max_m[A_w(m)]} \tag{10.7}$$

This measure is referred to as the *amplitude deviation* of the overlap factor. Variables should be chosen accordingly to keep this factor around or below one per cent.

We have seen that the STFT process that is bound to provide a suitable frequency domain representation of the input signal, is a far from trivial process and is dependent on some low-level parameters closely related to the signal processing domain. A little theoretical knowledge is required but only practice will surely lead to the desired results.

Peak Detection

The sinusoidal model assumes that each spectrum of the STFT representation can be explained by a series of sinusoids. For a given frequency resolution, using enough points in the spectrum, a sinusoid can be identified by its shape. Theoretically, a sinusoid that is stable both in amplitude and in frequency - a partial - has a well-defined frequency representation: the transform of the analysis window used to compute the Fourier transform. It should be possible to take advantage of this characteristic to distinguish partials from other frequency components. However, in practice this is rarely the case, since most natural sounds are not perfectly periodic and do not have nicely spaced and clearly defined peaks in the frequency domain. There are interactions between the different components, and the shapes of the spectral peaks cannot be detected without tolerating some mismatch. Only certain instrumental sounds (e.g., the steady-state part of an oboe sound) are periodic enough and sufficiently free from prominent noise components that the frequency representation of a stable sinusoid can be recognized easily in a single spectrum (see Fig. 10.8). A practical solution is to detect as many peaks as possible, with some small constraints, and delay the decision of what is a "well behaved" partial, to the next step in the analysis: the peak continuation algorithm.

A "peak" is defined as a local maximum in the magnitude spectrum, and the only practical constraints to be made in the peak search are to have a frequency range and a magnitude threshold. Due to the sampled nature of the spectrum returned by the FFT, each peak is accurate only to within half a sample. A spectral sample represents a frequency interval of f_s/N Hz, where f_s is the sampling rate and N is the FFT size. Zero-padding in the time domain increases the number of spectral samples per Hz and thus increases the accuracy of the simple peak detection (see previous section). However, to obtain frequency accuracy on the level of 0.1 % of the distance from the top of an ideal peak to its first zero crossing (in the case of a Rectangular window), the zero-padding factor required is 1000.

A more efficient spectral interpolation scheme is to zero-pad such that quadratic (or other simple) spectral interpolation, using only samples immediately surrounding the maximum-magnitude sample, suffices to refine the estimate to 0.1% accuracy. That is the approach we have chosen and is illustrated in Fig. 10.9. The frequency and magnitude of a peak is obtained from the magnitude spectrum expressed in dB. Then the phase value of the peak is measured by reading the value of the unwrapped phase spectrum at the position resulting from the frequency of the peak.

Although we cannot rely on the exact shape of the peak to decide whether it is a partial or not, it is sometimes useful to have a measure of how close its shape is
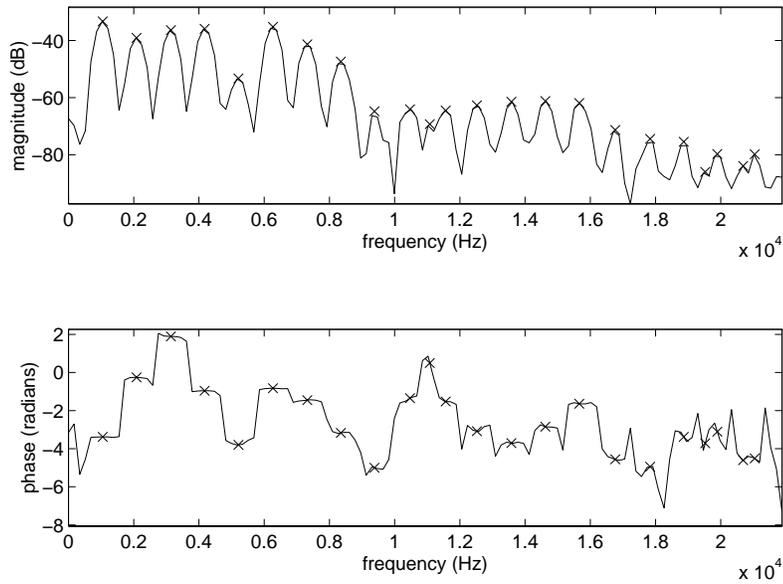
Figure 10.8 Peak detection. a. Peaks in magnitude spectrum. b. Peaks in the phase spectrum.
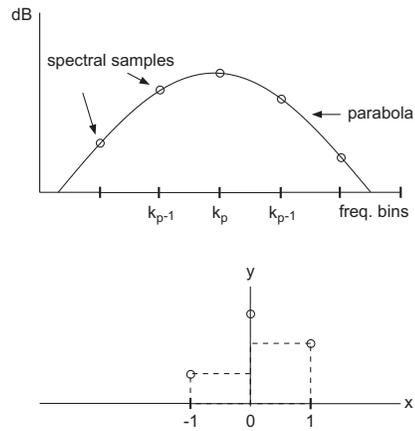
Figure 10.9 Parabolic interpolation in the peak detection process.

to the ideal sinusoidal peak. With this idea in mind, different techniques have been used in order to improve the estimation of the spectral peaks parameters [Dep97].

The following Matlab code (M-file 10.3) implements the peak detection algorithm, the function `PickPeaks` finds the local maximums of the spectrum.

M-file 10.3 (`PickPeaks.m`)

```
function  [loc, val] = PickPeaks(spectrum, nPeaks, minspace)
```

```
%function  [loc, val] = pickpeaks(spectrum, nPeaks, minspace)
%
%==> peaking the nPeaks highest peaks in the given spectrum
%    from the greater to the lowest
% data:
%        loc: bin number of peaks  (if loc(i)==0, no peak detected)
%        val: amplitude of the given spectrum
%        spectrum: spectrum (abs(fft(signal))
%        nPicks: number of peaks to pick
%        minspace: minimum of space between two peaks

[r, c]  = size(spectrum);
rmin = min(spectrum) - 1;

% ---find a peak, zero out the data around the peak, and repeat
val =  ones(nPeaks,c)*NaN;
loc =  zeros(nPeaks,c);

for k=1:c        %--- find all local peaks
   difference = diff([rmin; spectrum(:,k); rmin]); % derivate
   iloc = find(difference(1:r)>= 0 & difference(2:r+1) <= 0);
                                    % peak locations
   ival = spectrum(iloc,k);         % peak values

   for p=1:nPeaks
      [val(p,k),l] = max(ival);     % find current maximum
      loc(p,k) = iloc(l);           % save value and location
      ind = find(abs(iloc(l)-iloc) > minspace);
          % find peaks which are far away
      if (isempty(ind))
         break                      % no more local peaks to pick
      end
      ival = ival(ind);     % shrink peak value and location array
      iloc = iloc(ind);
   end
end
```

The function `interpolatedValues` (M-file 10.4) computes interpolated values for each peak.

M-file 10.4 (`interpolatedValues.m`)
```
function [iftloc, iftphase, iftval] = interpolatedValues ...
   (r, phi, N, zp, ftloc, ftval)
%function [iftloc, iftphase, iftval] = interpolatedValues ...
%   (r, phi, N, zp, ftloc, ftval)
%
```

```
%==> computation of the interpolated values
%    of location and magnitude (parabolic interpolation)
%    and phase (linear interpolation)
%
% data:
%      iftloc:    interpolated location (bin)
%      iftval:    interpolated magnitude
%      iftphase:  interpolated phase
%      ftloc:     peak locations (bin)
%      ftval:     peak magnitudes
%      r:         magnitude of the FFT
%      phi:       phase of the FFT
%      N:         size of the FFT
%      zp:        zero-padding multiplicative coefficient

%--- calculate interpolated peak position in bins (iftloc) ------
leftftval = r((ftloc-1).*((ftloc-1)>0)+((ftloc-1)<=0).*1);
rightftval= r((ftloc+1).*((ftloc+1)<N/2)+((ftloc+1)>=N/2).*(N/2));
leftftval = 20*log10(leftftval);
rightftval= 20*log10(rightftval);
ftval     = 20*log10(ftval);
iftloc    = ftloc + .5*(leftftval - rightftval) ./ ...
    (leftftval - 2*ftval + rightftval);


%--- interpolated ftloc -----------------------------------------

iftloc    = (iftloc>=1).*iftloc + (iftloc<1).*1;
iftloc    = (iftloc>N/2+1).*(zp/2+1) + (iftloc<=N/2+1).*iftloc;


%--- calculate interpolated phase (iphase) ----------------------
leftftphase = phi(floor(iftloc));
rightftphase= phi(floor(iftloc)+1);
intpfactor  = iftloc-ftloc;
intpfactor  = (intpfactor>0).*intpfactor ...
                +(intpfactor<0).*(1+intpfactor);
diffphase   = unwrap2pi(rightftphase-leftftphase);
iftphase    = leftftphase+intpfactor.*diffphase;


%--- calculate interpolate amplitude (iftval) -------------------
iftval = ftval-.25*(leftftval-rightftval).*(iftloc-ftloc);
```

These functions (as well as others that will be introduced later in this chapter) make use of the **unwrap2pi** function given by M-file 10.5.

M-file 10.5 (unwrap2pi.m)
```
function argunwrap = unwrap2pi (arg)
```

```
% function argunwrap = unwrap2pi (arg)
%
%==> unwrapping of the phase, in [-pi, pi]
%   arg: phase to unwrap
arg = arg - floor(arg/2/pi)*2*pi;
argunwrap = arg - (arg>=pi)*2*pi;
```

Pitch Estimation

Although the term *Pitch* should ideally be used to refer only to perceptual issues, the term *Fundamental Frequency* is not suitable to describe the output of techniques that will be explained herein. For that reason we will use both terms without making any distinction to refer to the output of these algorithms that aim to provide an estimation of this psychoacoustical sensation that is often (but not always) explained by the value of the fundamental frequency of a given harmonic series.

Pitch estimation is an optional step used when we know that the input sound is monophonic and pseudo-harmonic. Given this restriction and the set of spectral peaks of a frame, obtained as in the *Sinusoidal Analysis*, with magnitude and frequency values for each one, there are many possible pitch estimation strategies, none of them perfect [Hes83, Mah94, Can98]. The most obvious approach is to define the pitch as the common divisor of the harmonic series that best explains the spectral peaks found in a given frame. For example, in the Two-Way Mismatch procedure proposed by Maher and Beauchamp the estimated F_0 is chosen as to minimize discrepancies between measured peak frequencies and the harmonic frequencies generated by trial values of F_0. For each trial F_0, mismatches between the harmonics generated and the measured peak frequencies are averaged over a fixed subset of the available peaks. This is a basic idea on top of which we can add features and tune all the parameters for a given family of sounds.

Many trade-offs are involved in the implementation of a fundamental frequency detection system and every application will require a clear design strategy. For example, the issue of real-time performance is a requirement with strong design implications. We can add context specific optimizations when knowledge of the signal is available. Knowing, for instance, the frequency range of the F_0 of a particular sound helps both the accuracy and the computational cost. Then, there are sounds with specific characteristics, like in a clarinet where the even partials are softer than the odd ones. From this information, we can define a set of rules that will improve the performance of the used estimator.

In the framework of the sinusoidal plus residual analysis system, there are strong dependencies between the fundamental frequency detection step and many other analysis steps. For example, choosing an appropriate window for the Fourier analysis will facilitate detection of the fundamental and, at the same time, getting a good fundamental frequency will assist other analysis steps, including the selection of an appropriate window. Thus, it could be designed as a recursive process.

The following Matlab code (M-file 10.6) implements an algorithm for pitch detection (note that first, different computations are accomplished in order to decide

if the region being analyzed is harmonic or not).

M-file 10.6 (`pitchDetection.m`)
```
function[pitchvalue,pitcherror,isHarm]=pitchDetection(r,N, ...
                                      SR,nPeaks,iftloc,iftval)
% function[pitchvalue,pitcherror,isHarm]= ...
% pitchDetection(r,N,SR,nPeaks,iftloc,iftval)
%
%==> pitch detection function, using the Two-Way Mismatch
% algorithm (see TWM.m)
%
% data:
%       r:        FFT magnitude
%       N:        size of the FFT
%       SR:       sampling rate
%       nPeaks:   number of peaks tracked
%       iftloc, iftval: location (bin) and magnitude of the peak

%--- harmonicity evaluation of the signal
highenergy = sum(r(round(5000/SR*N):N/2)); % 5000 Hz to SR/2 Hz
lowenergy  = sum(r(round(50/SR*N):round(2000/SR*N)));
                                        % 50 Hz to 2000 Hz
isHarm = max(0,(highenergy/lowenergy < 0.6));

if (isHarm==1)    %-- 2-way mismatch pitch estimation when harmonic
   npitchpeaks = min(50,nPeaks);
   [pitchvalue,pitcherror] = ...
            TWM(iftloc(1:npitchpeaks),iftval(1:npitchpeaks),N,SR);
else
   pitchvalue = 0;
   pitcherror = 0;
end;

%--- in case of too much pitch error,
%    signal supposed to be inhamonic
isHarm = min (isHarm,(pitcherror<=1.5));
```

The Two-way mismatch procedure is implemented as follows:

M-file 10.7 (`TWM.m`)
```
function [pitch, pitcherror] = TWM (iloc, ival, N, SR)
%function [pitch, pitcherror] = TWM (iloc, ival, N, SR)
%
% => Two-way mismatch error pitch detection
%    using Beauchamp & Maher algorithm
%
```

```
% data:
%        iloc:   location (bin) of the peaks
%        ival:   magnitudes of the peaks
%        N:      number of peaks
%        SR:     sampling rate


ifreq = (iloc-1)/N*SR; % frequency in Hertz

%--- avoid zero frequency peak
[zvalue,zindex] = min(ifreq);
if (zvalue==0)
   ifreq(zindex) = 1;
   ival(zindex) = -100;
end

ival2 = ival;
[MaxMag,MaxLoc1] = max(ival2);
ival2(MaxLoc1)   = -100;
[MaxMag2,MaxLoc2]= max(ival2);
ival2(MaxLoc2)   = -100;
[MaxMag3,MaxLoc3]= max(ival2);

%--- pitch candidates
nCand = 10; % number of candidates
pitchc = zeros(1,3*nCand);
pitchc(1:nCand)=(ifreq(MaxLoc1)*ones(1,nCand))./((nCand ...
                                        +1-[1:nCand]));
pitchc(nCand+1:nCand*2)=(ifreq(MaxLoc2)*ones(1,nCand))./ ((nCand ...
                                        +1-[1:nCand]));
pitchc(nCand*2+1:nCand*3)=(ifreq(MaxLoc3)*ones(1,nCand))./((nCand ...
                                        +1-[1:nCand]));
%pitchc=100:300;
harmonic = pitchc;

%--- predicted to measured mismatch error
ErrorPM = zeros(fliplr(size(harmonic)));
MaxNPM = min(10,length(iloc));
for i=1:MaxNPM
   difmatrixPM = harmonic' * ones(size(ifreq))';
   difmatrixPM = abs(difmatrixPM ...
                 -ones(fliplr(size(harmonic)))*ifreq');
   [FreqDistance,peakloc] = min(difmatrixPM,[],2);
   Ponddif     = FreqDistance .* (harmonic'.^(-0.5));
   PeakMag     = ival(peakloc);
   MagFactor   = max(0, MaxMag - PeakMag + 20);
```

```
    MagFactor    = max(0, 1.0 - MagFactor/75.0);
    ErrorPM      = ErrorPM ...
                    +(Ponddif+MagFactor.*(1.4*Ponddif-0.5));
    harmonic     = harmonic+pitchc;
end

%--- measured to predicted mismatch error
ErrorMP = zeros (fliplr(size(harmonic)));
MaxNMP  = min(10,length(ifreq));

for i=1:length(pitchc)
    nharm        = round(ifreq(1:MaxNMP)/pitchc(i));
    nharm        = (nharm>=1).*nharm + (nharm<1);
    FreqDistance = abs(ifreq(1:MaxNMP) - nharm*pitchc(i));
    Ponddif      = FreqDistance.* (ifreq(1:MaxNMP).^(-0.5));
    PeakMag      = ival(1:MaxNMP);
    MagFactor    = max(0,MaxMag - PeakMag + 20);
    MagFactor    = max(0,1.0 - MagFactor/75.0);
    ErrorMP(i)   = sum(MagFactor.*(Ponddif ...
                        +MagFactor.*(1.4*Ponddif-0.5)));
end

%--- total error
Error = (ErrorPM/MaxNPM) + (0.3*ErrorMP/MaxNMP);
[pitcherror, pitchindex] = min(Error);

pitch = pitchc(pitchindex);
```

Peak Continuation

The peak detection process returns the estimated magnitude, frequency, and phase of the prominent peaks in a given frame sorted by frequency. Once the spectral peaks of a frame have been detected, and possibly a fundamental frequency identified, a peak continuation algorithm can organize the peaks into time-varying trajectories.

The output of the Sinusoidal Analysis is a set of spectral peak values (frequency, magnitude and phase) organized into frequency trajectories, where each trajectory models a time-varying sinusoid (see Fig. 10.10). As it will be shown later, from this information we can synthesize a sound using additive synthesis. The less restrictive the peak detection step is, the more faithful the reconstruction of the original sound will be after synthesis.

The sinusoidal model assumes that each of these peaks is part of a frequency trajectory and the peak continuation algorithm is responsible for assigning each peak to a given "track". There are many possibilities for such a process. The original one used by McAulay and Quatiery (see Fig. 10.11) in their sinusoidal representation [McA86] is based on finding, for each peak, the closest one in frequency in the
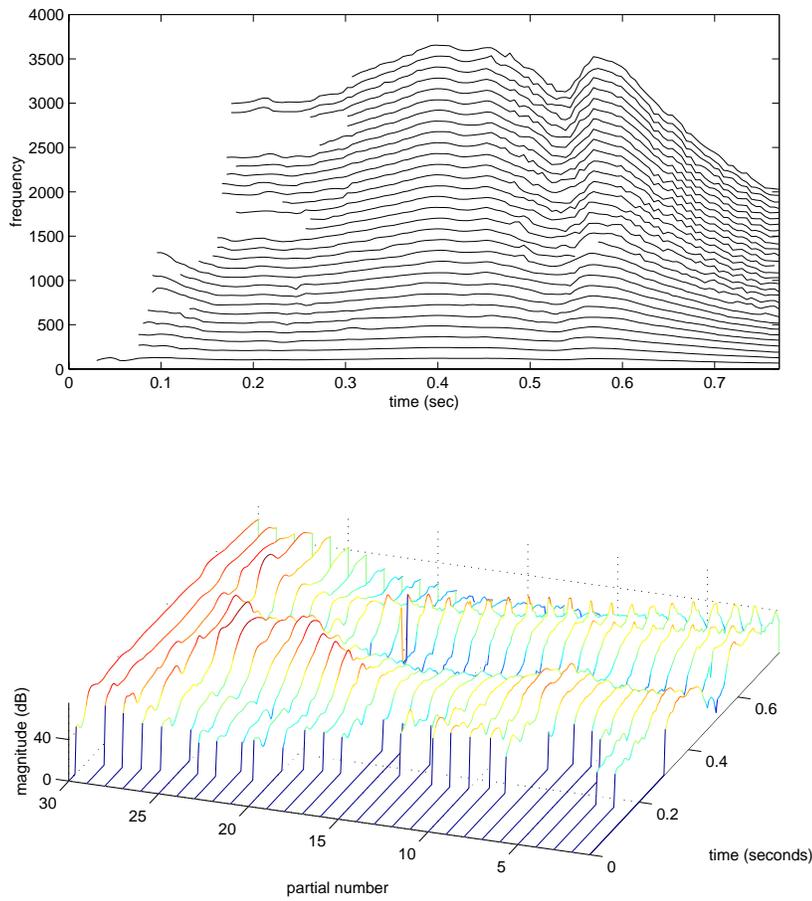
Figure 10.10 Frequency trajectories resulting from the sinusoidal analysis of a vocal sound.
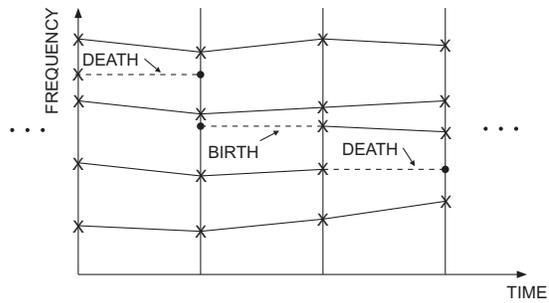
following frame.

Figure 10.11 Traditional peak continuation algorithm [McA86].

The schemes used in the traditional sinusoidal model (as the one just mentioned), incorporate all the spectral peaks into trajectories, thus obtaining a sinusoidal representation for the whole sound. These schemes are not optimal when we want the trajectories to follow just the stable partials, leaving the rest to be modeled as part of the residual component. For example, when the partials significantly change in frequency from one frame to the next, these algorithms easily switch from the partial that they were tracking to another one, which is closer at that point.
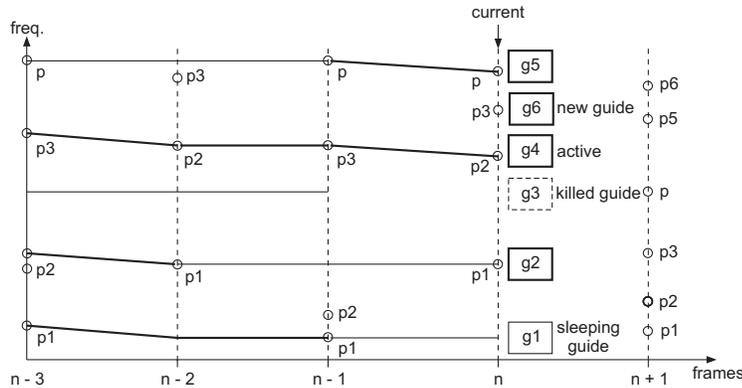

Figure 10.12 Peak continuation process. *g* represent the guides and *p* the spectral peaks.

Here we will describe a basic framework under which we can define rules for specifying the behavior of the partials of musical sounds and thus implement systems for identifying partials out of spectral peaks. The behavior of a partial, and therefore the way to track it, varies depending on the signal. Whether we have speech, a harmonic instrumental tone, a gong sound, a sound of an animal, or any other, the time-varying behavior of the partials will be different. Thus, the algorithm requires some knowledge about the characteristics of the sound that is being analyzed.

The basic idea of the algorithm is that a set of "guides" advances in time through the spectral peaks of each frame, looking for the appropriate peaks to be used (according to the specified constraints) and forming trajectories out of them (see Fig. 10.12). Thus, a guide is an abstract entity employed to create sinusoidal trajectories, which are the actual result of the peak continuation process. The instantaneous state of the guides, including their frequency and magnitude, are continuously updated as the guides are turned on, continued, and finally turned off. For the case of harmonic sounds, these guides are initialized according to the harmonic series of the detected fundamental frequency, and for inharmonic sounds, each guide is created dynamically.

The guides use the peak values and their context, such as surrounding peaks and fundamental frequency, to advance in time and form trajectories. For example, by using the detected fundamental frequency and the "memory" of an incoming trajectory to a given frame, we control the adaptation of the guides to the instantaneous changes in the sound. For a very harmonic sound, since all the harmonics

evolve together, the fundamental should be the main control. Nevertheless, when the sound is less harmonic and we cannot rely on the fundamental frequency as a strong reference for all the harmonics, the information of the incoming trajectory should have a bigger weight.

Each peak is assigned to the guide that is closest to it and that is within a given frequency and amplitude deviation. If a guide does not find a match, it is assumed that the corresponding trajectory must "turn off". In inharmonic sounds, if a guide has not found a continuation peak for a given amount of time the guide is killed. New guides, and therefore new trajectories, are created from the peaks of the current frame that are not incorporated into trajectories by the existing guides. If there are killed or unused guides, a new guide can be started. Searching through the "unclaimed" peaks of the frame for the one with the highest magnitude creates a guide. Once the trajectories have been continued for a few frames, the short ones can be deleted and we can fill the "gaps" encountered in long trajectories. A real-time implementation would not be able to use the rules that make use of the information of "future" frames.

The creation of trajectories from the spectral peaks is compatible with very different strategies and algorithms. A promising approach is to use Hidden Markov models [Dep93]. This type of approach might be very valuable for tracking partials in polyphonic sounds and complex inharmonic tones.

In our Matlab implementation (M-file 10.8), we have chosen to provide a simple tracking algorithm that uses a simplified version of the techniques previously introduced for the case of harmonic and inharmonic sounds.

M-file 10.8 (peakTrackSimple.m)
```
function[iloc,ival,iphase,previousiloc,previousival, ...
      distminindex]=peakTrackSimple(nSines,nPeaks,N, ...
   SR,pitchvalue,iftloc,iftval,iftphase,isHarm, ...
   previousiloc,previousival);
% function[iloc,ival,iphase,previousiloc,previousival, ...
%      distminindex]=peakTrackSimple(nSines,nPeaks,N, ...
%   SR,pitchvalue,iftloc,iftval,iftphase,isHarm, ...
%   previousiloc,previousival);
%
%==> simplest partial tracking
% data:
%   iloc,ival,iphase:  location (bin), magnitude
%                      and phase of peaks (current frame)
%   previousiloc,previousival,previousiphase: idem for
%                                             previous frame
%   iftloc, iftval, iftphase:  idem of all of the peaks in the FT
%   distminindex:   indexes of the minimum distance
%                   between iloc and iftloc
%   nPeaks:         number of peaks detected
%   nSines:         number of peaks tracked
%   N:              size of the FFT
```

```
%   SR:              sampling rate
%   pitchvalue:      estimated pitch value
%   isHarm:          indicator of harmonicity

tmpharm = pitchvalue;    %--- temporary harmonic
iloc = zeros(nSines,1);
MindB = -100;
ival = zeros(nSines,1) + MindB;
iphase = zeros(nSines,1);
distminindex = zeros(nSines,1);
Delta = 0.01;


for i=1:nSines        %--- for each sinus detected
if (isHarm==1)          %--- for a harmonic sound
[closestpeakmag,closestpeakindex]=min(abs((iftloc-1)/N*SR-tmpharm));
tmpharm = tmpharm + pitchvalue;
else                    %--- for an inharmonic sound
[closestpeakmag,closestpeakindex]=min(abs(iftloc-previousiloc(i)));
end
iloc(i)   = iftloc(closestpeakindex); %--- bin of the closest
ival(i)   = iftval(closestpeakindex);
iphase(i) = iftphase(closestpeakindex);
dist      = abs(previousiloc-iloc(i));
[distminval, distminindex(i)] = min(dist);
end
```

You will also need the code for the function CreateNewTrack (M-file 10.9), which implements the initialization of newborn tracks during the peak continuation process.

M-file 10.9 (CreateNewTrack.m)

```
function[newiloc,newival]=CreateNewTrack(iftloc,iftval, ...
         previousiloc,previousival,nSines,MinMag);
% function[newiloc,newival]=CreateNewTrack(iftloc,iftval, ...
%          previousiloc,previousival,nSines,MinMag);
%
%==> creation of a new track by looking for a new significant
%    peak not already tracked
% data: iftlov, iftval: bin number & magnitude of peaks detected
%        previousiloc,
%        previousival:   idem for previous peaks detected
%        nSines:         number of sines
%        MinMag:         minimum magnitude (-100 dB) for
%                        0 amplitude

%--- removing peaks already tracked
```

```
for i=1:nSines
    [min, ind] = min(abs(iftval - previousival(i)));
    iftval(ind) = MinMag;
end
%--- keeping the maximum
[newival, ind] = max(iftval);
newiloc = iftloc(ind);
```

M-file 10.10 implements a visual representation of the sinusoidal tracks.

M-file 10.10 (PlotTracking.m)
```
function PlotTracking(SineFreq, pitch)
%function PlotTracking(SineFreq, pitch)
%
%==> plot the partial tracking
% data:
%       SineFreq: frequencies of the tracks
%       pitch:    frequency of the pitch
[nSines, nFrames] = size(SineFreq);

for n=1:nSines
    f=1;
    while (f<=nFrames)
        while (f<=nFrames & SineFreq(n,f)==0)
            f = f+1;
        end
        iStart = min(f,nFrames);
        while (f<=nFrames & SineFreq(n,f)>0)
            f = f+1;
        end
        iEnd = min(max(1,f-1),nFrames);
        if (iEnd > iStart)
        line((iStart:iEnd), SineFreq(n,iStart:iEnd));
        end
    end
end

h = line((1:nFrames), pitch(1:nFrames));
set(h,'linewidth', 2, 'Color', 'black');
```

Residual Analysis

Once we have identified the stable partials of a sound, we are ready to subtract them from the original signal and obtain the residual component. This subtraction can be done either in the time domain or in the frequency domain. A time domain

approach requires that first the time domain signal is synthesized from the sinusoidal trajectories, while if we stay in the frequency domain, we can perform the subtraction directly in the already computed magnitude spectrum. For the time domain subtraction, the phases of the original sound have to be preserved in the synthesized signal, thus we have to use a type of additive synthesis with which we can control the instantaneous phase. A type of synthesis that is computationally quite expensive. On the other hand, the sinusoidal subtraction in the spectral domain is simpler but not considerably more. Our sinusoidal information from the analysis is very much undersampled, since for every sinusoid we only have the value at the top of the peaks, and thus we have to generate all the frequency samples that belong to the sinusoidal peak to be subtracted.

Once we have either the residual spectrum or the residual time signal, it is useful to study it in order to check how well the partials of the sound were subtracted and therefore analyzed. If partials remain in the residual, the possibilities for transformations will be reduced, mainly because it will not be possible to approximate the residual as a stochastic signal, thus reducing its flexibility. In this case, we should re-analyze the sound until we get a good residual, free of deterministic components. Ideally, the resulting residual should be as close as possible to a stochastic signal.

From the residual signal, we can continue our modeling strategy. To model the stochastic part of sounds, such as the attacks of most percussion instrument, the bow noise in string instruments, or the breath noise in wind instruments, we need a good time resolution and we can give up some frequency resolution. The deterministic component cannot maintain the sharpness of attacks, because, even if a high framerate is used we are forced to use a long enough window, and this size determines most of the time resolution. When the deterministic subtraction is done in the time domain, the time resolution in the stochastic analysis can be improved by redefining the analysis window. The frequency domain approach implies that the subtraction is done in the spectra computed for the deterministic analysis, thus the STFT parameters cannot be changed [Ser89].

Since it is the deterministic signal that is subtracted from the original sound, measured from long windows, the resulting residual signal might have the sharp attacks smeared. To improve the stochastic analysis, we can "fix" this residual so that the sharpness of the attacks of the original sound is preserved. The resulting residual is compared with the original waveform and its amplitude re-scaled whenever the residual has a greater energy than the original waveform. Then the stochastic analysis is performed on this scaled residual. Thus, the smaller the window the better time resolution we will get in the residual. We can also compare the synthesized deterministic signal with the original sound and whenever this signal has a greater energy than the original waveform it means that a smearing of the deterministic component has been produced. This can be fixed somewhat by scaling the amplitudes of the deterministic analysis in the corresponding frame by the difference between the original sound and the deterministic signal.

Sinusoidal subtraction

The first step of the Residual Analysis is the synthesis of the sinusoidal tracks obtained as the output of the Peak Continuation algorithm. For a time domain subtraction (see Fig. 10.13) the synthesized signal will reproduce the instantaneous phase and amplitude of the partials of the original sound. One frame of the sinusoidal part of the sound, $d(m)$, is generated by

$$d(m) = \sum_{r=1}^{R} \hat{A}_r \cos[m\hat{\omega}_r + \hat{\varphi}_r], \qquad m = 0, 1, 2, \ldots, S - 1 \tag{10.8}$$

where R is the number of trajectories present in the current frame and S is the length of the frame. To avoid "clicks" at the frame boundaries, the parameters $\hat{A}_r, \hat{\omega}_r, \hat{\varphi}_r$ are smoothly interpolated from frame to frame.

The instantaneous amplitude $\hat{A}(m)$ is easily obtained by linear interpolation from frame to frame. Frequency and phase values are tied together (frequency is the phase derivative), and both control the instantaneous phase $\hat{\theta}(m)$, defined by

$$\hat{\theta}(m) = m\hat{\omega} + \hat{\varphi}. \tag{10.9}$$

Different approaches are possible for computing the instantaneous phase [McA86]. Thus we are able to synthesize one frame of a sound by

$$d^l(m) = \sum_{r=1}^{R^l} \hat{A}_r^l(m) \cos[\hat{\theta}_r^l(m)] \tag{10.10}$$

which goes smoothly from the previous to the current frame with each sinusoid accounting for both the rapid phase changes (frequency) and the slowly varying phase changes.

Residual Approximation

One of the underlying assumptions of the *Sinusoidal plus Residual* model is that the residual is a stochastic signal. Such an assumption implies that the residual is fully described by its amplitude and its general frequency characteristics (see Fig. 10.14). It is unnecessary to keep either the instantaneous phase or the exact spectral shape information. Based on this, a frame of the stochastic residual can be completely characterized by the output of a filter, which has a noise input signal. The filter encodes the amplitude and general frequency characteristics of the residual. The representation of the residual for the overall sound will be a sequence of these filters, i.e., a time-varying filter.

The filter design problem is generally solved by performing some sort of curve fitting in the magnitude spectrum of the current frame [Str80, Sed88]. Standard techniques are: spline interpolation [Cox71], the method of least squares [Sed88], or straight-line approximations.
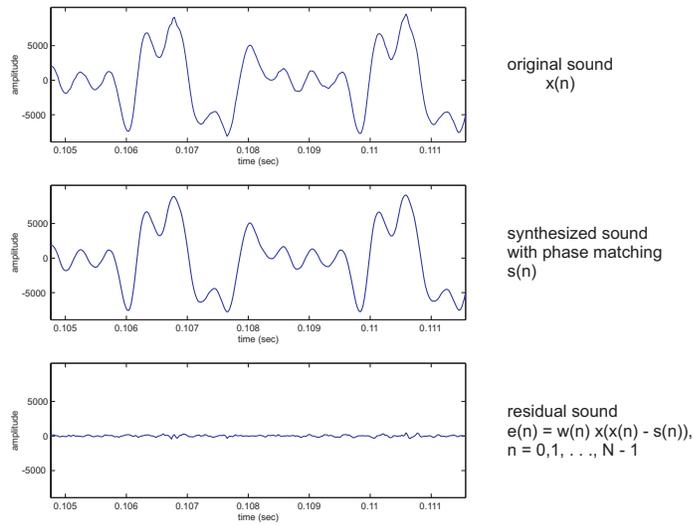
Figure 10.13 Time domain substraction.
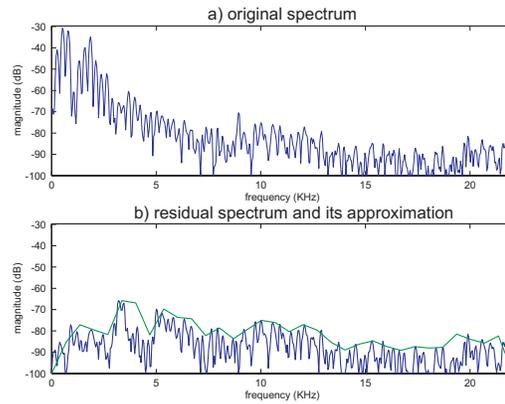
Figure 10.14 (a) Original spectrum, (b) Residual spectrum and approximation.

One way to carry out the line-segment approximation is to step through the magnitude spectrum and find local maxima in each of several defined sections, thus giving equally spaced points in the spectrum that are connected by straight lines to create the spectral envelope. The number of points gives the accuracy of the fit, and can be set depending on the sound complexity. Other options are unequally spaced points, for example, logarithmically spaced, or spaced according to other perceptual criteria.

Another practical alternative, as already seen in chapter 9 (see 9.2.2), is to use a type of least squares approximation called linear predictive coding, LPC [Mak75, Mar75]. LPC is a popular technique used in speech research for fitting an nth-order

polynomial to a magnitude spectrum. For our purposes, the line-segment approach is more flexible than LPC, and although LPC results in less analysis points, the flexibility is considered more important. For a comprehensive collection of different approximation techniques of the residual component see [Goo97].

10.3.2   Feature analysis

The accomplishment of a meaningful parameterization for sound transformation applications is a difficult task. We want a parameterization that offers an intuitive control over the sound transformation process, with which we can access most of the perceptual attributes of a sound. The analysis techniques described so far result in a simple parameterization, appropriate for describing the lower physical characteristics of the sound. In the Sinusoidal plus Residual model, these parameters are the instantaneous frequency, amplitude and phase of each partial and the instantaneous spectral characteristics of the residual signal.

There are other useful instantaneous attributes that give a higher-level abstraction of the sound characteristics. For example we can describe fundamental frequency, amplitude and spectral shape of the sinusoidal component, amplitude and spectral shape of the residual component, and the overall amplitude. These attributes are calculated at each analysis frame from the output of the basic Sinusoidal plus Residual analysis. Afterwards, some of them can be extracted.

From a digital effects design point of view, the extraction of such attributes allows us to implement transformations that modify only one of those features without affecting the rest. A clear example is illustrated in Fig. 10.2 where the fundamental frequency is extracted, multiplied by a scaling factor, and then incorporated back into the original spectral data.

Many other features like the degree of harmonicity, noisiness, spectral tilt, or spectral centroid, can also be computed from the spectral representation of a sound. Some of them are just information attributes that describe the characteristics of the frame and have mainly found applications in sound classification tasks.

Apart from the instantaneous, or frame, values, it is also useful to have parameters that characterize the time evolution of the sound. The time changes can be described by the derivatives of each one of the instantaneous attributes.

Another important step towards a musically useful parameterization is the segmentation of a sound into regions that are homogeneous in terms of its sound attributes. Then we can identify and extract region attributes that will give higher-level control over the sound.

From the basic sinusoidal plus residual representation it is possible to extract some of the attributes mentioned above. The critical issue is how to extract them while minimizing interferences, thus obtaining significant high level attributes free of correlations [Ser98]. The general process will be to first extract instantaneous attributes and their derivatives, then segment the sound based on that information, and finally extract region attributes.

As already indicated, the basic instantaneous attributes are: amplitude of sinusoidal and residual component, overall amplitude, fundamental frequency, spectral shape of sinusoidal and residual component, harmonic distortion, noisiness, spectral centroid, and spectral tilt. These attributes are obtained at each frame using the information that results from the basic Sinusoidal plus Residual analysis and not taking into account the data from previous or future frames. The amplitude of the sinusoidal component is the sum of the amplitudes of all harmonics of one frame expressed in dB,

$$AS = 20\log_{10}\left(\sum_{i=1}^{l} a_i\right) \tag{10.11}$$

where a_i is the linear amplitude of the ith harmonic and l is the total number of harmonics found in the frame.

The amplitude of the residual component is the sum of the absolute values of the residual of one frame expressed in dB. This amplitude can also be computed by adding the frequency samples of the corresponding magnitude spectrum, according to

$$
\begin{aligned}
AR &= 20\log_{10}\left(\sum_{n=0}^{M-1} |x_R(n)|\right) \\
&= 20\log_{10}\left(\sum_{k=0}^{N-1} |X_R(k)|\right),
\end{aligned}
\tag{10.12}
$$

where $x_R(n)$ is the residual sound, M is the size of the frame, $X_R(k)$ is the spectrum of the residual sound, and N is the size of the magnitude spectrum.

The total amplitude of the sound at one frame is the sum of its absolute values expressed in dB. It can also be computed by summing the amplitudes of the sinusoidal and residual components, as given by

$$
\begin{aligned}
A &= 20\log_{10}\left(\sum_{n=0}^{M-1} |x(n)|\right) = 20\log_{10}\left(\sum_{k=0}^{N-1} |X(k)|\right) \\
&= 20\log_{10}\left(\sum_{i=1}^{l} a_i + \sum_{k=0}^{N-1} |X_R(k)|\right)
\end{aligned}
\tag{10.13}
$$

where $x(n)$ is the original sound and $X(k)$ is its spectrum.

The fundamental frequency is the frequency that best explains the harmonics of one frame. Many different algorithms can be used to compute the fundamental frequency (see previous section 10.3.1, for example) but a reasonable approximation, once we have the sinusoidal component, can be the weighted average of all the normalized harmonic frequencies

$$F_0 = \sum_{i=1}^{l} \frac{f_i}{i} \times \frac{a_i}{\sum_{i=1}^{l} a_i} \tag{10.14}$$

where f_i is the frequency of the ith harmonic.

The spectral shape of the sinusoidal component is the envelope described by the amplitudes and frequencies of the harmonics, or its approximation,

$$S_{shape} \quad = \quad \{(f_1, a_1)(f_2, a_2) \dots (f_l, a_l)\}. \tag{10.15}$$

The spectral shape of the residual component is an approximation of the magnitude spectrum of the residual sound of one frame. A simple function is computed as the line segment approximation of the spectrum,

$$R_{shape} \quad = \quad \{e_1, e_2, \dots, e_q, \dots, e_{Ncoef}\} \tag{10.16}$$

Other spectral approximation techniques can be considered depending on the type of residual and the application [Goo96].

The frame-to-frame variation of each attribute is a useful measure of its time evolution, thus an indication of changes in the sound. It is computed in the same way for each attribute,

$$\Delta \quad = \quad \frac{\text{Val}(l) - \text{Val}(l-1)}{H/f_s} \tag{10.17}$$

where $\text{Val}(l)$ is the attribute value for the current frame, $\text{Val}(l-1)$ is the attribute value for the previous one, H is the hop size and f_s is the sampling rate.

As an example, the following function (M-file 10.11) implements the computation of the spectral shape that will be used in some of the effects implemented in the next sections.

M-file 10.11 (sort.m)
```
%--- sorting according to the frequencies iloc
[isortedloc, ind] = sort(iloc);
isortedval = ival(ind);
[indr, indc] = find(isortedloc);
newloc = isortedloc(indr);
newval = isortedval(indr);
%--- computing the spectral shape without redundant values
spectralShape = [];
spectralShape(1,1) = 1;
spectralShape(2,1) = MinMag;
shapePos = 1;
for i=1:length(newloc)
   if newloc(i) > spectralShape(1,shapePos)
      shapePos = shapePos + 1;
      spectralShape(1,shapePos) = newloc(i);
      spectralShape(2,shapePos) = newval(i);
   end
end
%--- adding boudaries values
spectralShape(1,shapePos+1) = N/2;
spectralShape(2,shapePos+1) = MinMag;
```

Segmentation

Sound segmentation has proven important in automatic speech recognition and music transcription algorithms. For our purposes it is very valuable as a way to apply region dependent transformations. For example, a time stretching algorithm would be able to transform the steady state regions, leaving the rest unmodified.

A musically meaningful segmentation process divides a melody into notes and silences and then each note into an attack, a steady state and a release region.

The techniques originally developed for speech [Vid90], such as those based on pattern recognition or knowledge-based methodologies, start to be used in music segmentation applications [Ros98]. Most of the approaches apply classification methods that start from sound features, such as the ones described in this chapter, and are able to group sequences of frames into predefined categories. No reliable and general-purpose technique has been found. Our experience is that they require narrowing the problem to a specific type of musical signal or including a user intervention stage to help directing the segmentation process.

Region attributes

Once a given sound has been segmented into regions we can compute the attributes that describe each one. Most of the interesting attributes are the mean and variance of each of the frame attributes for the whole region. For example, we can compute the spectral shape or the mean and variance for the amplitude of sinusoidal and residual components, the fundamental frequency, or the spectral tilt.

Global attributes that can characterize attacks and releases make use of the average variation of each of the instantaneous attributes, such as average fundamental frequency variation, average amplitude variation, or average spectral shape change. In the steady state regions it is important to extract the average value of each of the instantaneous attributes and measure other global attributes such as time-varying rate and depth of vibrato. Vibrato is a specific attribute present in many steady state regions of sustained instrumental sounds that requires a special treatment [Her98].

Some region attributes can be extracted from the frame attributes in the same way that they were extracted from the Sinusoidal plus Residual data. The result of the extraction of the frame and region attributes is a hierarchical multi-level data structure where each level represents a different sound abstraction.

10.3.3   Synthesis

From the output of the analysis techniques presented we can synthesize a new sound. The similarity with respect to the original sound will depend on how well the input sound fits the implicit model of the analysis technique and the settings of the different variables that the given technique has. In the context of the chapter we are interested in transforming the analysis output in order to produce a specified effect in the synthesized sound.

All these transformations can be done in the frequency domain. Afterwards, the output sound can be synthesized using the techniques presented in this section. The sinusoidal component will be generated using some type of additive synthesis approach and the residual, if present, will be synthesized using some type of subtractive synthesis approach.

Thus, the transformation and synthesis of a sound is done in the frequency domain; generating sinusoids, noise, or arbitrary spectral components, and adding them all to a spectral frame. Then, we compute a single IFFT for each frame, which can yield efficient implementations.

Figure 10.15 shows a block diagram of the final part of the synthesis process. Previous to that we have to transform and add all the *High Level Features*, if they have been extracted, and obtain the lower level data (sine and residual) for the frame to be synthesized. Since the stored data might have a different frame rate, or a variable one, we also have to generate the appropriate frame by interpolating the stored ones. These techniques are presented in the following sections.
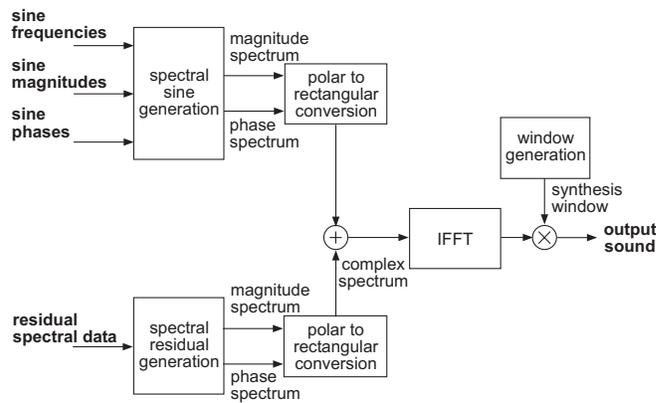
Figure 10.15 Diagram of the spectral synthesis.

Sinusoidal synthesis

The sinusoidal component is generated with additive synthesis, similar to the sinusoidal synthesis that was part of the analysis, with the difference that now the phase trajectories might be discarded.

Additive synthesis is based on the control of the instantaneous frequency and amplitude of a bank of oscillators, as shown in Fig. 10.16. The instantaneous amplitude $\hat{A}(m)$ of an oscillator is obtained by linear interpolation

$$\hat{A}(m) = \hat{A}^{l-1} + \frac{(\hat{A}^l - \hat{A}^{l-1})}{S}m, \tag{10.18}$$

where $m = 0, 1, \ldots, S - 1$ is the time sample in the l^{th} synthesis frame.

Figure 10.16 Additive synthesis block diagram.

The instantaneous phase is taken to be the integral of the instantaneous frequency, where the instantaneous radian frequency $\hat{\omega}(m)$ is obtained by linear interpolation

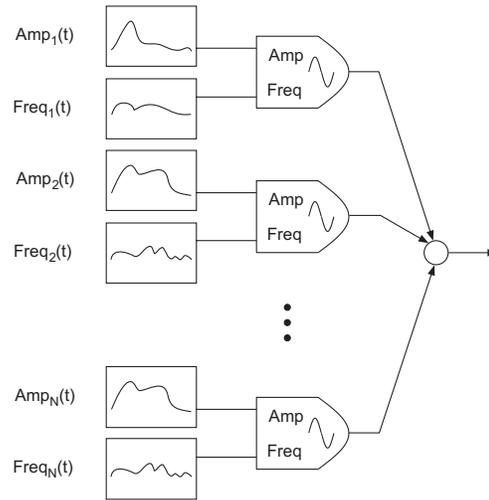$$\hat{\omega}(m) = \hat{\omega}^{l-1} + \frac{(\hat{\omega}^l - \hat{\omega}^{l-1})}{S} m \qquad (10.19)$$

and the instantaneous phase for the r^{th} sinusoid is given by

$$\hat{\theta}_r(m) = \hat{\theta}_r(l-1) + \hat{\omega}_r(m). \qquad (10.20)$$

Finally, the synthesis equation becomes

$$d^l(m) = \sum_{r=1}^{R^l} \hat{A}_r^l(m) \cos[\hat{\theta}_r^l(m)] \qquad (10.21)$$

where $\hat{A}(m)$ and $\hat{\theta}(m)$ are the calculated instantaneous amplitude and phase.

A very efficient implementation of additive synthesis, when the instantaneous phase is not preserved, is based on the inverse FFT [Rod92]. While this approach looses some of the flexibility of the traditional oscillator bank implementation, especially the instantaneous control of frequency and magnitude, the gain in speed is significant. This gain is based on the fact that a sinusoid in the frequency domain is a sinc-type function, the transform of the window used, and on these functions not all the samples carry the same weight. To generate a sinusoid in the spectral domain it is sufficient to calculate the samples of the main lobe of the window transform, with the appropriate magnitude, frequency and phase values. We can then synthesize as many sinusoids as we want by adding these main lobes in the FFT buffer and

performing an IFFT to obtain the resulting time-domain signal. By an overlap-add
process we then get the time-varying characteristics of the sound. In the following
Matlab code (M-file 10.12) we implement a sinusoidal synthesis algorithm based on
this latter approach.

M-file 10.12 (`sinefillspectrum.m`)
```
function padsynthft=sinefillspectrum(iloc,ival,iphase,nSines, ...
    w1Length, zp, bh92SINE2SINE, bh92SINE2SINEsize)
%function padsynthf =sinefillspectrum(iloc,ival,iphase,nSines, ...
%   w1Length, zp, bh92SINE2SINE, bh92SINE2SINEsize)
%
%=> compute the spectrum of all the sines in the frequency
%   domain, in order to remove it from the signal
% data:
%  padsynth:
%  iloc, ival, iphase: location (bin), magnitude value (dB)
%                      and phase of a peak
%  nSines:             number of sines (=length of ival and iloc)
%  w1Length:           size of the analysis window
%  zp:                 zero-padding multiplicative coefficient
%  bh92SINE2SINE:      Blackman-Harris window
%  bh92SINE2SINEsize:  Blackman-Harris window size

peakmag=10.^(ival/20);          % magnitude (in [0;1])
halflobe=8*zp/2-1;              % bin number of the half lobe
firstbin=floor(iloc)-halflobe;  % first bin for filling positive
                                % frequencies
firstbin2=floor(w1Length*zp-iloc+2)-halflobe;
   % idem for negative frequencies
binremainder=iloc-floor(iloc);
sinphase=sin(iphase);
cosphase=cos(iphase);
findex=1-binremainder;
bh92SINE2SINEindexes =zeros(8*zp,1);
sinepadsynthft=zeros(w1Length*zp+halflobe+halflobe+1,1);
padsynthft =zeros(w1Length*zp,1);

%--- computation of the complex value
for i=1:nSines    %--- for each sine
 if (iloc(i)~=0) %--- JUST WORK WITH NON ZEROS VALUES OF iloc !!!
                 % -> tracked sines
  beginindex = floor(0.5 + findex(i)*512/zp)+1;
  bh92SINE2SINEindexes=[beginindex:512/zp:beginindex ...
     +512/zp*(8*zp-1)]';
  if (bh92SINE2SINEindexes(8*zp)>bh92SINE2SINEsize)
      bh92SINE2SINEindexes(8*zp)=bh92SINE2SINEsize;
```

```
  end
  magsin=bh92SINE2SINE(bh92SINE2SINEindexes) ...
      .*sinphase(i)*peakmag(i);
  magcos=bh92SINE2SINE(bh92SINE2SINEindexes) ...
      .*cosphase(i)*peakmag(i);
  %--- fill positive frequency
  sinepadsynthft(firstbin(i)+halflobe:firstbin(i) ...
      +halflobe+8*zp-1)= ...
  sinepadsynthft(firstbin(i)+halflobe:firstbin(i)+ ...
  halflobe+8*zp-1)+(magcos+j*magsin);
  %--- fill negative frequency
  if (firstbin2(i)+halflobe <= w1Length*zp)
  sinepadsynthft(firstbin2(i)+halflobe:firstbin2(i) ...
      +halflobe+8*zp-1)= ...
  sinepadsynthft(firstbin2(i)+halflobe:firstbin2(i)+ ...
  halflobe+8*zp-1)+(magcos-j*magsin);
  end
 end
end

%--- fill padsynthft
padsynthft=padsynthft+sinepadsynthft(halflobe+1:halflobe+1 ...
    +w1Length*zp-1);
padsynthft(1:halflobe) = padsynthft(1:halflobe) + ...
    sinepadsynthft(w1Length*zp+1:w1Length*zp+halflobe);
padsynthft(w1Length*zp-halflobe+1:w1Length*zp) = ...
    padsynthft(w1Length*zp-halflobe+1:w1Length*zp) ...
    + sinepadsynthft(1:halflobe);
```

The synthesis frame rate is completely independent of the analysis one. In the implementation using the IFFT we want to have a frame rate high enough so as to preserve the temporal characteristics of the sound. As in all short-time based processes we have the problem of having to make a compromise between time and frequency resolution. The window transform should have the fewest possible significant bins since this will be the number of points required to generate each sinusoid. A good window choice is the Blackman-Harris 92dB because, as already explained in section 10.3.1, its main lobe includes most of the energy. However the problem is that such a window does not overlap perfectly to a constant in the time domain without having to use very high overlap factors, thus very high frame rates. A solution to this problem [Rod92] is to undo the effect of the window by dividing by it in the time domain and applying a triangular window before performing the overlap-add process. This will give a good time-frequency compromise. The Matlab code for generating the triangular windowis given by M-file 10.13.

M-file 10.13 (triang.m)
```
function w = triang(n)
```

```
% TRIANG Triangular window.
if rem(n,2)
% It's an odd length sequence
w = 2*(1:(n+1)/2)/(n+1);
w = [w w((n-1)/2:-1:1)]';
else
% It's even
w = (2*(1:(n+1)/2)-1)/n;
w = [w w(n/2:-1:1)]';
end
```

Residual synthesis

The synthesis of the residual component of the sound is also performed in the frequency domain (see Fig. 10.17). When the analyzed residual has not been approximated, i.e. it is represented as a magnitude and phase spectrum for each frame, as a STFT, each residual spectrum is added to the spectrum of the sinusoidal component at each frame. But when a magnitude spectral envelope has approximated the residual, an appropriate complex spectrum has to be generated.
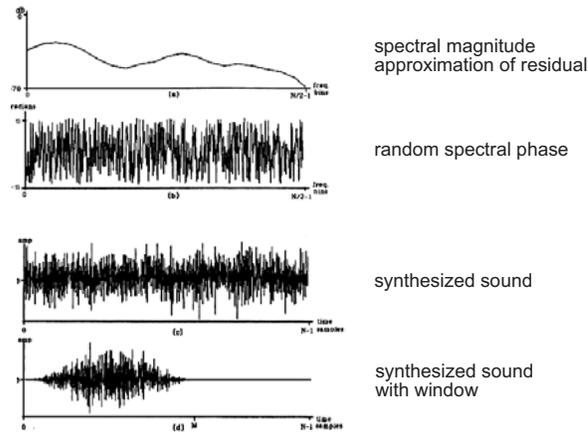
Figure 10.17 Residual synthesis approximation.

The synthesis of a stochastic signal from the residual approximation can be understood as the generation of noise that has the frequency and amplitude characteristics described by the spectral magnitude envelopes. The intuitive operation is to filter white noise with these frequency envelopes, that is, perform a time-varying filtering of white noise, which is generally implemented using the time-domain convolution of white noise with the impulse response corresponding to the spectral envelope of a frame. We do it in the frequency domain by creating a magnitude spectrum from the approximated one, or its transformation, and generating a random phase spectrum with new values at each frame in order to avoid periodicity.

Integration of Sinusoidal and Residual synthesis

Once the two spectral components are generated, we have to add the spectrum of the residual component to that of the sinusoids. In the process of generating the noise spectrum there has not been any window applied, since the data was added directly into the spectrum without any smoothing consideration, but in the sinusoidal synthesis we have used a Blackman-Harris 92dB, which is undone in the time domain after the IFFT. Therefore we should apply the same window to the noise spectrum before adding it to the sinusoidal spectrum. Convolving the transform of the Blackman-Harris 92dB by the noise spectrum accomplishes this, and there is only the need to use the main lobe of the window since it includes most of its energy. This is implemented quite efficiently because it only involves a few bins and the window is symmetric. Then we can use a single IFFT for the combined spectrum (see Fig. 10.18). Finally in the time domain we undo the effect of the Blackman-Harris 92dB and impose the triangular window. By an overlap-add process we combine successive frames to get the time-varying characteristics of the sound.
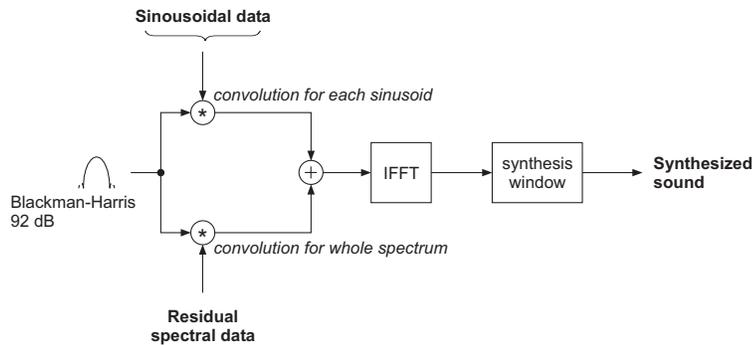
Figure 10.18 Integrating sinusoidal plus residual synthesis.

Several other approaches have been used for synthesizing the output of a sinusoidal plus residual analysis. However, these techniques include modifications to the model as a whole (see [Fit00], for example).

10.3.4   Main Analysis-Synthesis Application

In the following Matlab code (M-file 10.14), we make use of all the previous functions in order to implement a complete analysis-synthesis process. We will use this framework to implement sound effects in the next sections. Note, that no residual approximation is used in this implementation.

M-file 10.14 (SMS.m)

```
%=============================================================
% SMS-Matlab like emulation
```

```
%%==============================================================
clear all
close all
%==== USER DATA =====
DAFx_in  = wavread('love.wav'); % wave file
SR       = 44100;  % sampling rate
w1Length = 2048;   % analysis window size
n1       = 256;    % analysis window hop size
nPeaks   = 100;    % number of peaks detected
nSines   = 50;     % number of sinuosoids to track (and synthetise)
minSpacePeaks = 2; % minimum space (bins) between two picked peaks
zp       = 2;      % zero-padding coefficient
rgain    = 1.;     % gain for the residual component
MaxFreq  = 11000;  % maximum frequency, in Hertz, for plottings
MinMag   = -100;   % minimum magnitude, in dB, for plottings

%--- figure data
%fig1 = 'yes'; % if uncommented, will plot the Blackman-Harris
              % window
%fig2 = 'yes';  % if uncommented, will plot the peaks detection
               % and tracking in one frame
%fig3 = 'yes';  % if uncommented, will plot the peak trackings
               % real-time
%fig4 = 'yes';  % if uncommented, will plot the original and
               % the transformed FFT in one frame
%fig5 = 'yes';  % if uncommented, will plot the peak trackings
               % only at the end of the process
%fig6 = 'yes';  % if uncommented, will plot the original signal,
               % its sine and residual part,
               % and the transformed signal

%=== Definition of the Windows ===

%--- definition of the analysis window
fConst=2*pi/(w1Length+1-1);
w1=[1:w1Length]';
w1=.35875 -.48829*cos(fConst*w1)+.14128*cos(fConst*2*w1) ...
   -.01168*cos(fConst*3*w1);
w1=w1/sum(w1)*2;
N=w1Length*zp; % new size of the window
%--- synthesis window
w2=w1;
n2=n1;
%--- triangular window
wt2=triang(n2*2+1); % triangular window
%--- main lobe table of bh92
```

```
[bh92SINE2SINE,bh92SINE2SINEsize]=bh92SINE2SINEgeneration;
%--- data for the loops
frametime = n1/SR;
pin  = 0;
pout = 0;
TuneLength=length(DAFx_in);
pend=TuneLength-w1Length;

%=== Definition of the data arrays ===

DAFx_in      = [zeros(w1Length/2-n1-1,1); DAFx_in];
DAFx_outsine = zeros(TuneLength,1);
DAFx_outres  = zeros(TuneLength,1);
%--- arrays for the partial tracking
iloc            = zeros(nSines,1);
ival            = zeros(nSines,1);
iphase          = zeros(nSines,1);
previousiloc    = zeros(nSines,1);
previousival    = zeros(nSines,1);
maxSines = 400; % maximum voices for harmonizer
syniloc   = zeros(maxSines,1);
synival   = zeros(maxSines,1);
previoussyniloc = zeros(maxSines,1);
previousiphase  = zeros(maxSines,1);
currentiphase   = zeros(maxSines,1);
%--- arrays for the sinus' frequencies and amplitudes
SineFreq = zeros(nSines,ceil(TuneLength/n2));
SineAmp  = zeros(nSines,ceil(TuneLength/n2));
pitch = zeros(1,1+ceil(pend/n1));
pitcherr = zeros(1,1+ceil(pend/n1));

%--- creating figures ---
if(exist('fig1'))
   h = figure(1); set(h,'position', [10, 45, 200, 200]);
end
if(exist('fig2'))
   h = figure(2); set(h,'position', [10, 320, 450, 350]);
axisFig2 = [0  MaxFreq MinMag  0]; zoom on;
end
if(exist('fig3'))
   h = figure(3); set(h,'position', [220, 45, 550, 200]);
   axisFig3 = [1  1+ceil(pend/n1) 0  MaxFreq]; zoom on;
end
if(exist('fig4'))
   h = figure(4); set(h,'position', [470, 320, 450, 350]);
   axisFig4 = [0  MaxFreq MinMag  0]; zoom on;
```

```
end
if(exist('fig5'))
    h = figure(5); set(h,'position', [220, 45, 550, 200]);
    axisFig5 = [1  1+ceil(pend/n1) 0  MaxFreq]; zoom on;
end

%--- plot the Blackman-Harris window
if(exist('fig1'))
figure(1)
plot(20*log10(abs(fftshift(fft(bh92SINE2SINE)/bh92SINE2SINEsize))))
title('Blackman-Harris window');xlabel('Samples');
ylabel('Amplitude')
end

tic
%UUUUUUUUUUUUUUUUUUUUUUUUUUUUUUUUUUUUUUUUUUUUUUUUUUUUUUUUUUUUU
disp('analazing frame ...');

while pin<pend
%--- windowing
grain = DAFx_in(pin+1:pin+w1Length).*w1(1:w1Length);
%--- zero padding
padgrain = zeros(N,1);
padgrain(1:w1Length/2)     = grain(w1Length/2+1:w1Length);
padgrain(N-w1Length/2+1:N) = grain(1:w1Length/2);
%--- fft computation
f   = fft(padgrain);
r   = abs(f);
phi = angle(f);
ft  = r.*exp(j*phi);

%===== Analysis =====

%--- peak detection (and their plottings)
[ftloc, ftval]=PickPeaks(r(1:N/2),nPeaks,minSpacePeaks);

%--- calculate interpolated values (peak position,phase,amplitude)
[iftloc, iftphase, iftval] = ...
      interpolatedValues (r,phi,N,zp,ftloc,ftval);

%--- pitch detection
[pitchvalue,pitcherror,isHarm] = ...
      pitchDetection (r,N,SR,nPeaks,iftloc,iftval);
pitch(1+pin/n1) = pitchvalue*isHarm;
pitcherr(1+pin/n1) = pitcherror;
```

```
%--- peaks tracking
if (pin==0) %--- for the first frame
   nNewPeaks = nSines;
else  %--- creating new born tracks
for i=1:nSines
   if (previousiloc(i)==0)
   [previousiloc(i), previousival(i)] = CreateNewTrack ...
       (iftloc, iftval, previousiloc, previousival, nSines, MinMag);
   nNewPeaks = nNewPeaks - 1;
   end
end

%--- simple Peak tracker
[iloc,ival,iphase,previousiloc,previousival,distminindex] = ...
  peakTrackSimple(nSines,nPeaks,N,SR,pitchvalue,iftloc, ...
  iftval,iftphase,isHarm,previousiloc,previousival);
end

%--- savings
previousival = ival;
previousiloc = iloc;
SineFreq(:,1+pin/n1)=max((iloc-1)/N*SR,0.);
   % frequency of the partials
SineAmp(:,1+pin/n1)=max(ival, MinMag);
   % amplitudes of the partials

syniloc(1:nSines) = max(1,iloc);
synival(1:nSines) = ival;

if(exist('fig3')) % plot: the trackings of partials
  figure(3); clf; hold on
  PlotTracking(SineFreq(:,1:1+pin/n1), pitch(1:1+pin/n1));
  xlabel('Frame number');ylabel('Frequency (Hz)');
  axis(axisFig3);title('Peak tracking'); drawnow
  end

%--- residual computation
resfft = ft;
if(isHarm==1)
resfft=resfft-sinefillspectrum(iloc,ival,iphase,nSines,...
   w1Length, zp, bh92SINE2SINE, bh92SINE2SINEsize);
end

%--- figures
if(exist('fig2'))
figure(2); clf; hold on
```

```
                        % plot: FFT of the windowed signal (Hz,dB)
plot((1:N/2)/N*SR, 20*log10(r(1:N/2)));
  for l=1:nPeaks      % plot: the peaks detected
  plot([ftloc(l)-1 ftloc(l)-1]/N*SR, ...
  [20*log10(ftval(l)),MinMag-1],'r:x');
  end
for l=1:nSines        % plot: sines tracked and the residual part
plot([iloc(l)-1, iloc(l)-1]/N*SR, [ival(l), MinMag-1],'k')
end
plot((1:N/2)/N*SR, 20*log10(abs(resfft(1:N/2))),'g');
if(isHarm)            % plot: true pitch of each harmonic
for l=1:nSines
  plot([pitchvalue*l, pitchvalue*l],[1, MinMag-1],'y:')
end
end
xlabel('Frequency (Hz)');ylabel('Magnitude (dB)');axis(axisFig2);
title('Peak detection and tracking for one frame'); drawnow
end

nSynSines = nSines;
%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%
%===== Transformations =====
%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%


%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%
%===== Synthesis =====

%--- phase computation
if (pin > 0)
   for i=1:nSynSines
     if (syniloc(i)~=0)
   ifreq = (previoussyniloc(distminindex(i))+ syniloc(i))/2;
       % average bin
   freq = (ifreq-1)/N*SR; % freq in Hz (if loc=1 --> freq=0)
   currentiphase(i)=unwrap2pi(previousiphase(distminindex(i))+...
           2*pi*freq*frametime);
 end
   end
end

previoussynival = synival;
previoussyniloc = syniloc;
previousiphase  = currentiphase;

%--- compute sine spectrum
```

```
padsynthft=sinefillspectrum(syniloc, synival, currentiphase, ...
   nSynSines,w1Length, zp, bh92SINE2SINE, bh92SINE2SINEsize);
if (isHarm==0)
padsynthft = zeros(size(padsynthft));
end

%--- residual computation
respadgrain=real(ifft(resfft));
resgrain=[respadgrain(N-w1Length/2+1:N); ...
respadgrain(1:w1Length/2)]./w2(1:w1Length);
ressynthgrain=wt2(1:n2*2).*resgrain(w1Length/2-n2:w1Length/2+n2-1);
DAFx_outres(pout+1:pout+n2*2)=DAFx_outres(pout+1:pout+n2*2)+ ...
   ressynthgrain;

%--- sinusoidal computation
sinpadgrain=real(ifft(padsynthft));
singrain=[sinpadgrain(N-w1Length/2+1:N); ...
   sinpadgrain(1:w1Length/2)]./w2(1:w1Length);
sinsynthgrain=wt2(1:n2*2).*singrain(w1Length/2-n2:w1Length/2+n2-1);
DAFx_outsine(pout+1:pout+n2*2)=DAFx_outsine(pout+1:pout+n2*2)+ ...
   sinsynthgrain;

%--- figure with original signal and transformed signal FFT
synthr = abs(fft(respadgrain + sinpadgrain));
if(exist('fig4'))
 figure(4); clf; hold on
 plot((1:N/2)/N*SR, 20*log10(r(1:N/2)),'b:'); axis(axisFig4);
 plot((1:N/2)/N*SR, 20*log10(synthr(1:N/2)),'r');
 figure(4);
 xlabel('Frequency (Hz)');ylabel('Magnitude (dB)');axis(axisFig4);
 title('FFT of original (blue) and transformed (red) signals');
drawnow
end

%--- increment loop indexes
   pin  = pin + n1;
   pout = pout + n2;
   disp(pin/n1);
end
%UUUUUUUUUUUUUUUUUUUUUUUUUUUUUUUUUUUUUUUUUUUUUUUUUUUUUUUUUUUU
toc

%===== write output sounds =====
DAFx_in = DAFx_in(w1Length/2-n1:length(DAFx_in));
% remove the zeros added for the process
DAFx_outresynth = DAFx_outsine(1:TuneLength)+ ...
```

```
rgain*DAFx_outres(1:TuneLength);
mm = max(abs(DAFx_outresynth));
wavwrite(DAFx_outresynth/mm, SR, 'DAFx_out.wav');
wavwrite(DAFx_outsine/mm, SR, 'DAFx_outsine.wav');
wavwrite(DAFx_outres/mm, SR ,'DAFx_outres.wav');


if(exist('fig3')==0 & exist('fig5'))% plot: trackings of partials
                              % only at the end of the process
figure(5); clf; hold on
PlotTracking(SineFreq(:,1:1+pend/n1), pitch(1:1+pend/n1));
xlabel('Frame number'); ylabel('Frequency (Hz)'); axis(axisFig5);
title('Peak tracking'); drawnow
end

if(exist('fig6')) % plot the input signal, its sinus
                % and its residual part, and the transformed signal
figure(6)
subplot(4,1,1); plot(DAFx_in); xlabel('input signal');
subplot(4,1,2); plot(DAFx_outsine);xlabel('sinus part');
subplot(4,1,3); plot(DAFx_outres);xlabel('residual part');
subplot(4,1,4); plot(DAFx_outresynth);
xlabel('resynthetized signal');
end
```

10.4   FX and Transformations

In this section we introduce a set of effects and transformations based on the analysis-synthesis framework introduced throughout this chapter. All of them are accompanied by their corresponding Matlab code. In order to use them, you just have to add the code of the effect to use under the "= `Transformation` =" line in the main analysis-synthesis application code of the previous section.

10.4.1   Filtering with Arbitrary Resolution

Filters are probably the paradigm of a "classical" effect. Many different implementations are provided in the general DSP literature and in the previous chapters of this book. Here we introduce a different approach that differs in many aspects from the classical one.

For our "filter" implementation, we take advantage of the sinusoidal plus residual model in order to modify the amplitude of any arbitrary partial present in the sinusoidal component.

For example, we can implement a bandpass filter defined by (x, y) points where x is the frequency value in Hertz and y is the amplitude factor to apply (see Fig. 10.19).

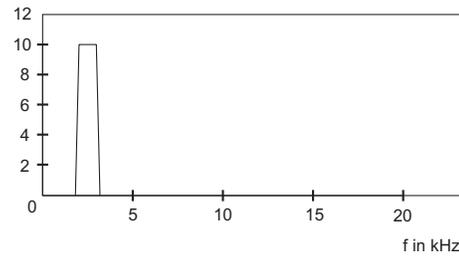In the example code given below, we define a bandpass filter with passband range [2100 3000].

Figure 10.19 Bandpass filter with arbitrary resolution.

M-file 10.15 (`ch10_t_filter_arb.m`)
```
%===== Filtering with arbitrary resolution =====
Filter         = [ 0 2099 2100 3000 3001 22050 ; 0 0 1 1 0 0 ];
[syniloc,ind]  = sort(iloc);
FilterEnvelope = interp1(Filter(1,:)',Filter(2,:)',syniloc/N*SR);
synival        = ival(ind)+(20*log10(max(FilterEnvelope,10^-9)));
synival(ind)   = synival;
syniloc(ind)   = syniloc;
```

As shown, our filter does not need to be characterized by a traditional transfer function, and more complex functions can be defined by summing delta-functions.

For example, the following code filters out the even partials of the input sound. If applied to a sound with a broadband spectrum, like a vocal sound, it will convert it to a clarinet-like sound.

M-file 10.16 (`ch10_t_voice2clar.m`)
```
%=== voice to clarinet ===
syniloc = iloc;
synival = ival;
if (isHarm == 1)
    for i=1:nSines
        harmNum = round((((iloc(i)-1)/w1Length*SR/2)/pitchvalue);
        if (mod(harmNum,2)==0) % case of an even harmonic number
            synival(i) = MinMag;
        end
    end
end
```

10.4.2   Partial Dependent Frequency Scaling

In a similar way, we can apply a frequency scaling to the sinusoidal components of our modeled sound. In that way, we can transpose all the partials in the spectrum or

reproduce pseudo-inharmonicities like frequency stretching of higher partials, which is representative for a piano sound.

In this first example we introduce a frequency shift factor to all the partials of our sound (see Fig. 10.20). Note, though, that if a constant is added to every partial of a harmonic spectrum, the resulting sound will be inharmonic.
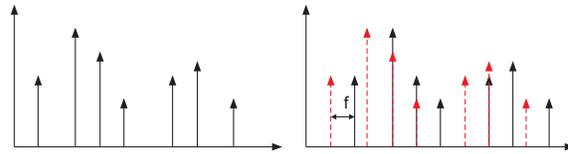

Figure 10.20 Frequency shift of the partials.

M-file 10.17 (ch10_t_freqshift.m)
```
%==== Frequency Shift =====
   fstretch = 300;  % frequency shift in Hz
   syniloc  = iloc + round(fstretch/SR*N);
   syniloc  = syniloc.*(syniloc<=N/2);
```

Another effect we can implement following this same idea is to add a stretching factor to the frequency of every partial. The relative shift of every partial will depend on its original partial index, following the formula

$$f_i = f_i \cdot f_{stretch}^{(i-1)}. \tag{10.22}$$

Figure 10.21 illustrates this frequency stretching.
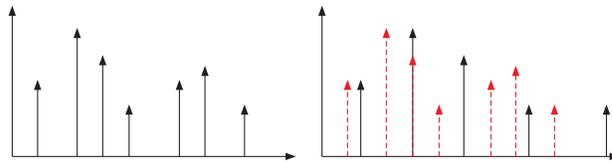
Figure 10.21 Frequency stretching.

M-file 10.18 (ch10_t_freqstretch.m)
```
%===== Frequency Stretch   =====
   fstretch = 1.1;
   [syniloc,ind] = sort(iloc);
   syniloc = syniloc.*((fstretch).^[0:nSines-1]');
   syniloc = syniloc.*(syniloc<=N/2);
```

In the same way, we can scale all the partials multiplying them by a given scaling factor. Note, that this effect will act as a pitch shifter without timbre preservation.

M-file 10.19 (ch10_t_freqscale.m)
```
%==== Frequency Scale =====
    fscale  = 1.6;  % frequency scaling factor
    syniloc = iloc * fscale;
    syniloc = syniloc.*(syniloc<=N/2);
```

10.4.3   Pitch Transposition with Timbre Preservation

In section 9.3.4, a technique was introduce in order to transpose the pitch of a sound without affecting its timbre. Here we use a similar technique in order to preserve the spectral shape of only the sinusoidal component. For that reason we scale the frequency of each partial applying the original spectral shape.

M-file 10.20 (ch10_t_pitchtimbre.m)
```
%===== Pitch transposition with timbre preservation =====
if (isHarm == 1)
  pt = 2.;                    % pitch transposition factor
  [spectralShape,shapePos]=CalculateSpectralShape(iloc,...
  ival,MinMag,N);
  [syniloc, synival]= PitchTransposition(iloc,ival,...
  spectralShape,shapePos,pt,N);
  %--- comb filtering the residual
  CombCoef = 1;
  if (isHarm==1)
  resfft = CombFilter(resfft, N, SR/(pitchvalue*pt), CombCoef);
  end
end
```

The function `PitchTransposition` is given by:

M-file 10.21 (PitchTransposition.m)
```
%===== Pitch Transposition =====
function [syniloc,synival]=PitchTransposition(iloc,ival, ...
        spectralShape,shapePos,pt,N)
syniloc = iloc.*pt;
syniloc = syniloc.*(syniloc<=N/2);
%lin. interpol. of the spectral shape for synival computation
if shapePos > 1
synival=interp1(spectralShape(1,:)',spectralShape(2,:)',syniloc);
else
synival = ival;
end
```

The function `CombFilter` is implemented as:

M-file 10.22 (CombFilter.m)
```
function combFT = CombFilter(FT, N, delay, ampl)
```

```
%===> Comb filter in the frequency domain
% data:
%     combFT: FT of the signal comb filtered
%     FT:     FT of the signal to filtered
%     N:      size of the FT
%     delay:  delay to apply, in samples
%     ampl:   amplitude of the multiplying coefficient (in [0,1])
coef = ampl * exp(-2*j*pi*delay*(0:N-1)/N)';
combFT = FT .* (1 + coef + coef.^2);
```

Pitch Discretization to Temperate Scale

An interesting effect can be accomplished by forcing the pitch to take the nearest frequency value of the temperate scale. It is indeed a very particular case of pitch transposition where the pitch is quantified to one of the 12 semitones of an octave. This effect is widely used on vocal sounds for *dance* music and is many times referred to with the misleading name of *vocoder effect*.

M-file 10.23 (ch10_t_PitchDiscrete.m)
```
%===== Pitch discretization to temperate scale =====
if (pitchvalue ~= 0)
  nst = round(12*log(pitchvalue/55)/log(2));
  discpitch = 55*((2^(1/12))^nst); % discretized pitch
  pt = discpitch/pitchvalue ;      % pitch transposition factor
  [spectralShape,shapePos]=CalculateSpectralShape(iloc,ival, ...
  MinMag,N);
  [syniloc, synival]=PitchTransposition(iloc,ival,spectralShape, ...
  shapePos,pt,N);
  %--- comb filtering the residual
  CombCoef = 1;
  if (isHarm==1)
  resfft = CombFilter(resfft, N, SR/(pitchvalue*pt), CombCoef);
  end
end;
```

10.4.4   Vibrato and Tremolo

Vibrato and tremolo are common effects used in different kinds of acoustical instruments, including the human voice. Both are low frequency modulations: vibrato is applied to the frequency and tremolo to the amplitude of the partials. Note, though, that in this particular implementation, both effects share the same modulation frequency.

M-file 10.24 (ch10_t_vibtrem.m)
```
%===== vibrato and tremolo =====
if (isHarm == 1)
```

```
vtf = 5;        % vibrato-tremolo frequency in Hz
va  = 10;       % vibrato depth in percentil
td  = 3;        % tremolo depth in dB
synival = ival + td*sin(2*pi*vtf*pin/SR);   % tremolo
pt=1+va/200*sin(2*pi*vtf*pin/SR);% pitch transposition factor
[spectralShape, shapePos] = CalculateSpectralShape(iloc, ...
ival,MinMag,N);
[syniloc,synival]=PitchTransposition(iloc,ival,spectralShape, ...
shapePos,pt,N);
    %--- comb filtering the residual
    CombCoef = 1;
    resfft = CombFilter(resfft, N, SR/(pitchvalue*pt), CombCoef);
end
```

10.4.5   Spectral Shape Shift

As already seen in the previous chapter, many interesting effects can be accomplished by shifting the spectral shape or spectral envelope of the sinusoidal components of a sound. This shift is performed in such a way that no new partials are generated, just the amplitude envelope of the sinusoidal components is modified (see Fig. 10.22). In the following code we implement a shift of the spectral envelope by just modifying the amplitude of the partials according to the values of the shifted version of the spectral shape.
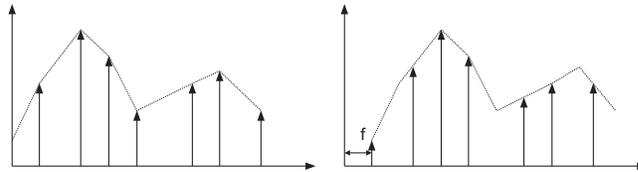
Figure 10.22 Spectral shape shift of value Δf.

M-file 10.25 (ch10_t_SpectSS.m)
```
%===== Spectral Shape Shift (positive or negative) =====
sss  = -200;                   % spectral shape shift value in Hz
%--- spectral shape computation
[spectralShape,shapePos]=CalculateSpectralShape(iloc, ...
ival,MinMag,N);
%--- spectral shape shift
syniloc = zeros(nSines,1);
if shapePos > 1
[shiftedSpectralShape,shapePos]=SpectralShapeShift(sss, ...
iloc, ival, spectralShape, shapePos, N, SR);
end
syniloc = iloc;
```

```
%linear interpol. of the spectral shape for synival computation
if shapePos > 1
synival = interp1(shiftedSpectralShape(1,1:shapePos+1)', ....
shiftedSpectralShape(2,1:shapePos+1)', syniloc, 'linear');
else
synival = ival;
end
```

The function **SpectralShapeShift** is implemented as follows:

M-file 10.26 (SpectralShapeShift.m)
```
%===== Spectral Shape Shift =====
function [shiftedSpectralShape,shapePos]=SpectralShapeShift(sss, ...
iloc, ival, spectralShape, shapePos, N, SR)
shiftedSpectralShape = spectralShape;
sssn = round (sss*N/SR);% spectral shape shift in number of bins
if sssn > 0
   shiftedSpectralShape(1,2:shapePos)=min(N/2, ...
   spectralShape(1,2:shapePos) + sssn);
   for i=shapePos:-1:1
      if shiftedSpectralShape(1,i) < N/2
         shapePos = i;
         break;
      end;
   end;
else
   shiftedSpectralShape(1,2:shapePos)= ...
   max(1,spectralShape(1,2:shapePos)+ sssn);
   for i=1:shapePos
      if shiftedSpectralShape(1,i) > 1
         shiftedSpectralShape(1,2:2+shapePos+1-i) = ...
 shiftedSpectralShape(1,i:shapePos+1);
         shapePos = shapePos-(i-2);
         break;
      end;
   end;
end;
```

10.4.6 Gender Change

Using the results of 10.4.3 and 10.4.5 we can change the gender of a given vocal sound. Note how by combining different "basic" effects we are able to step higher in the level of abstraction and get closer to what a naive user could ask for in a sound transformation environment, such as having a gender control on a vocal processor.

In this implementation, we apply two transformations in order to convert a male voice into a female one (variable tr='m2f'). The first one is a pitch transposition

an octave higher. The other one is a shift in the spectral shape . The theoretical
explanation to this effect is that women change their formant (resonant filters)
frequencies depending on the pitch. That is, when a female singer rises up the
pitch, the formants move along with the fundamental frequency.

To convert a female into a male voice (variable `tr='f2m'`) we also apply a pitch
transposition and a shift in the spectral shape. This shifting has to be applied in a
way the formants of the female voice remain stable along different pitches.

M-file 10.27 (`ch10_t_gender.m`)
```
%===== gender change: woman to man =====
tr='m2f'; %male to female
%tr='f2m'; %female to male
if (isHarm == 1)
pitchmin=100;
pitchmax=500;
sssmax = 50;
if (pitchvalue<pitchmin)
sss = 0;
elseif (pitchvalue>pitchmax)
sss = sssmax;
else
sss = (pitchvalue-pitchmin)/((pitchmax-pitchmin)/sssmax);
end
if(tr=='f2m')
sss=-sss;
pt=0.5;
else
pt=2;
end
%--- spectral shape computation
[spectralShape,shapePos]=CalculateSpectralShape(iloc,ival,MinMag,N);
% --- spectral shape shift
syniloc = zeros(nSines,1);
if shapePos > 1
[shiftedSpectralShape,shapePos]=SpectralShapeShift(sss,iloc,...
 ival,spectralShape,shapePos,N,SR);
end
syniloc = iloc;
%linear interpol. of the spectral shape for synival computation
if shapePos > 1
synival = interp1(shiftedSpectralShape(1,1:shapePos+1)', ...
shiftedSpectralShape(2,1:shapePos+1)', syniloc, 'linear');
else
synival = ival;
end
%--- pitch transposition
```

```
pt = 0.5;
[syniloc, synival] = PitchTransposition(iloc,ival,spectralShape, ...
shapePos,pt,N);
%--- comb filtering the residual
CombCoef = 1;
if (isHarm==1)
resfft = CombFilter(resfft, N, SR/(pitchvalue*pt), CombCoef);
end
end
```

10.4.7 Harmonizer

In order to create the effect of a harmonizing vocal chorus, we can add pitch-shifted versions of the original voice (with the same timbre) and force them to be in tune with the original melody.

M-file 10.28 (ch10_t_harmonizer.m)
```
%===== harmonizer =====
nVoices   = 2;
nSynSines = nSines*(1+nVoices);
[spectralShape, shapePos] = CalculateSpectralShape(....
   syniloc(1:nSines), synival(1:nSines), MinMag, N);
synival(1:nSines) = synival(1:nSines) - 100;
pt = [1.3 1.5]; % pitch transposition factor
ac = [-1 -2]; % amplitude change factor in dB
for i=1:nVoices
   [tmpsyniloc, tmpsynival] = PitchTransposition(...
  syniloc(1:nSines), synival(1:nSines),...
      spectralShape, shapePos, pt(i), N);
   tmpsynival = tmpsynival + ac(i);
   syniloc(nSines*i+1:nSines*(i+1)) = tmpsyniloc;
   synival(nSines*i+1:nSines*(i+1)) = tmpsynival;
   if (pin > 0)
      distminindex(nSines*i+1:nSines*(i+1))= ...
         distminindex(1:nSines)+nSines*i;
   end
end
```

10.4.8 Hoarseness

Although hoarseness is sometimes thought of as a symptom of some kind of vocal disorder [Chi94], this effect has sometimes been used by singers in order to resemble the voice of famous performers (Louis Armstrong or Tom Waits, for example). In this elemental approximation, we accomplish a similar effect by just applying a gain to the residual component of our analysis.

M-file 10.29 (ch10_t_hoarse.m)
```
rgain = 2; % gain factor applied to the residual
```

10.4.9   Morphing

Morphing is a transformation with which, out of two or more elements, we can generate new ones with hybrid properties.

With different names, and using different signal processing techniques, the idea of audio morphing is well known in the computer music community [Ser94, Tel95, Osa95, Sla96]. In most of these techniques, the morph is based on the interpolation of sound parameterizations resulting from analysis/synthesis techniques, such as the Short-time Fourier Transform (STFT), Linear Predictive Coding (LPC) or Sinusoidal Models (see cross synthesis and spectral interpolation in sections 9.3.1 and 9.3.3, respectively).

In the following Matlab code we introduce a morphing algorithm based on the interpolation of the frequency, phase, and amplitude of the sinusoidal component of two sounds. The factor `alpha` controls the amount of the first sound we will have in the resulting morph. Different controlling factors could be introduced for more flexibility. Note, that if the sounds have different durations, the sound resulting from the morphing will have the duration of the shortest one.

Next, we include the code lines that have to be inserted in the transformation part. However, the morph transformation requires two or more inputs we have not included in order to keep the code short and understandable. Therefore the code will have to include the following modifications:

1. Read two input sounds:

   ```
   DAFx_in1  = wavread('source1.wav');
   DAFx_in2  = wavread('source2.wav');
   ```

2. Analyze both sounds. This means every analysis code line will have to be duplicated using the variable names:

   ```
   iloc1, iloc2, ival1, ival2,
   iphase1, iphase2, syniloc1, syniloc2, synival1, synival2,
   syniphase1, syniphase2, distminindex1, distminindex2,
   previousiloc1, previousiloc2, previousival1, previousival2,
   pitch1, pitch2, and pitcherror1, pitcherror2.
   ```

M-file 10.30 (ch10_t_morph.m)
```
%=== Morphing ===
%--- sorting the frequencies in bins; iloc1, iloc2
[syniloc1, ind1] = sort(iloc1);
synival1   = ival1(ind1);
syniphase1 = iphase1(ind1);
```

```
distminindex1 = distminindex1(ind1);
[syniloc2, ind2] = sort(iloc2);
synival2   = ival2(ind2);
syniphase2 = iphase2(ind2);
distminindex2 = distminindex2(ind2);
%--- interpolation -----
alpha = 0.5;        % interpolation factor
syniloc     = alpha*syniloc1 + (1-alpha)*syniloc2;
synival     = alpha*synival1 + (1-alpha)*synival2;
%--- pitch computation
isHarmsyn = isHarm1*isHarm2;
if (isHarmsyn ==1)
npitchpeaks = min(50,nPeaks);
[pitchvalue,pitcherror]=
TWM(syniloc(1:npitchpeaks),synival(1:npitchpeaks),N,SR);
else
pitchvalue = 0;
pitcherror = 0;
end
if (pin==0) %--- for the first frame
nNewPeaks = nSines;
else
%--- creation of new born tracks
for i=1:nSines
  if (previoussyniloc(i)==0)
  [previoussyniloc(i),previoussynival(i)]=CreateNewTrack ...
  (syniloc,synival,previoussyniloc,previoussynival,nSines,MinMag);
  nNewPeaks = nNewPeaks - 1;
 end
end
%--- peak tracking of the peaks of the synthetized signal
 [syniloc,synival,syniphase,previoussyniloc,previoussynival, ...
 distminindex]=peakTrackSimple (nSines,nPeaks,N,SR,pitchvalue,...
 syniloc,synival,syniphase,isHarmsyn,previoussyniloc, ...
 previoussynival);
end
```

10.5   Content Dependent Processing

The hierarchical data structure that includes a complete description of a given sound offers many possibilities for sound transformations. Modifying several attributes at the same time and at different abstraction levels achieve, as it has already been pointed out in the previous section, most musically or end-user meaningful transformations.

Higher-level transformations can refer to aspects like sound character, articula-

tion or expressive phrasing. These ideas lead to the development of front ends such as graphical interfaces or knowledge-based systems [Arc97, Arc98] that are able to deal with the complexity of this sound representation.

In this section we introduce two applications that have been developed with these ideas in mind: a singing voice conversion and a time scaling module.

10.5.1   Real-Time Singing Voice Conversion

Here we present a very particular case of audio morphing. We want to morph, in real-time, two singing voice signals in such a way that we can control the resulting synthetic voice by mixing some characteristics of the two sources. Whenever this control is performed by means of modifying a reference voice signal matching its individual parameters to another, we can refer to it as voice conversion [Abe88].

In such a context, a karaoke-type application, in which the user can sing like his/her favorite singers, was developed [Can00]. The result is basically an automatic impersonating system that allows the user to morph his/her voice attributes (such as pitch, timbre, vibrato and articulations) with the ones from a pre-recorded singer, which from now on we will refer to as *target*.

In this particular implementation, the target's performance of the complete song to be morphed is recorded and analyzed beforehand. In order to incorporate the corresponding characteristics of the target's voice to the user's voice, the system first recognizes what the user is singing (phonemes and notes), looks for the same sounds in the target performance (i.e. synchronizing the sounds), interpolates the selected voice attributes, and synthesizes the morphed output voice. All this is accomplished in real-time.

Figure 10.23 shows the general block diagram of the voice impersonator system. The system relies on two main techniques that define and constrict the architecture: the SMS framework (see 10.2.2) and a Hidden Markov Model based Speech Recognizer (SR). The SMS implementation is responsible for providing a suitable parameterization of the singing voice in order to perform the morph in a flexible and musically-meaningful way. On the other hand, the SR is responsible for matching the singing voice of the user with the target.

Let's take an overview of the whole process. Before we can morph a particular song, we have to supply information about the song to be morphed and the song recording itself (Target Information and Song Information). The system requires the phonetic transcription of the lyrics, the melody as MIDI data, and the actual recording to be used as the target audio data. Thus, a good impersonator of the singer that originally sang the song has to be recorded. This recording has to be analyzed with SMS, segmented into morphing units (phonemes), and each unit labeled with the appropriate note and phonetic information of the song. This preparation stage is done semi-automatically, using a non-real time application developed for this task.

Once we have all the required inputs set we can start processing the user's voice. The first module of the running system includes the real-time analysis and
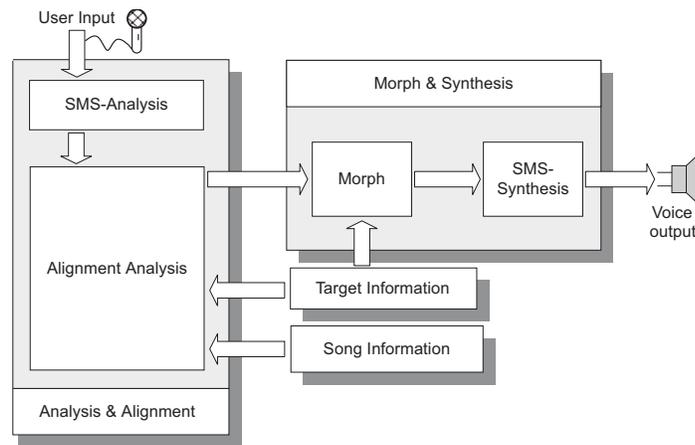
Figure 10.23 System block diagram.

the recognition/alignment steps. Each analysis frame, with the appropriate param-
eterization, is associated with the phoneme of a specific moment of the song and
thus with a target frame. Once a user frame is matched with a target frame, we
morph them by interpolating data from both frames and we synthesize the output
sound. Only voiced phonemes are morphed and the user has control over which
parameters are interpolated, and by how much. The frames belonging to unvoiced
phonemes are left untouched, thus always having the user's unvoiced consonants in
the output.

Several modifications are done to the basic SMS procedures to adapt them to the
requirements of the impersonator system. The major changes include the real-time
implementation of the whole analysis/synthesis process with a processing latency
of less than 30 milliseconds and the tuning of all parameters to the particular
case of the singing voice. These modifications include the extraction of higher-level
parameters meaningful in the case of the singing voice and that will be later used
in the morphing process.

The system includes an Automatic Speech Recognizer (ASR) based on phoneme-
base discrete HMM's in order to solve the matching problem. This ASR has been
adapted to handle musical information and works with very low delay [Los99] since
we cannot wait for a phoneme to be finished before it is recognized, moreover,
we have to assign a phoneme to each frame. This would be a rather impossi-
ble/impractical situation if it was not for the fact that the lyrics of the song are
known beforehand. This reduces a big portion of the search problem: all the possible
paths are restricted to just one string of phonemes, with several possible pronun-
ciations. The problem is cut down to the question of locating the phoneme in the
lyrics and positioning the start and end points.

Besides knowing the lyrics, musical information is also available. The user is
singing along with the music, and hopefully according to a tempo and melody

already specified in the score. Thus, we also know the time at which a phoneme is supposed to be sung, its approximate duration, its associated pitch, etc. All this information is used to improve the performance of the recognizer and also to allow resynchronization, for example in the case of a singer skipping a part of the song.

Depending on the phoneme the user is singing, a unit from the target is selected. Each frame from the user is morphed with a different frame from the target, advancing sequentially in time. Then the user has the choice of interpolating the different parameters extracted at the analysis stage, such as amplitude, fundamental frequency, spectral shape, residual signal, etc. In general, the amplitude will not be interpolated, thus always using the amplitude from the user and the unvoiced phonemes will not be morphed either, thus always using the consonants from the user. This will give the user the feeling of being in control. This recognition and matching process is illustrated in Fig. 10.24.
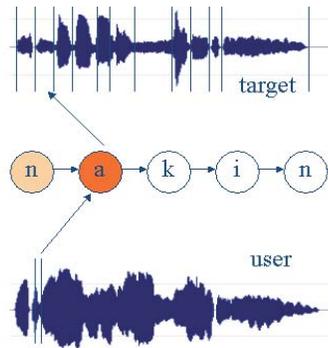
Figure 10.24 Recognition and matching of morphable units.

Whenever the spectral shape is interpolated, and the morph factor is set around 50%, the resulting spectral shape is smoothed and looses much of its timbre characteristic. This problem can be solved if formants are included in the spectral shape model and they are taken into account in the interpolation step.

In most cases, the durations of the user and target phonemes to be morphed will be different. If a given user's phoneme is shorter than the one from the target, the system will simply skip the remaining part of the target phoneme and go directly to the articulation portion. In the case when the user sings a longer phoneme than the one present in the target data, the system enters in the loop mode. Each voiced phoneme of the target has a loop point frame, marked in the pre-processing, non-real time stage. The system uses this frame for loop-synthesis in case the user sings beyond that point in the phoneme. Once we reach this frame in the target, the rest of the frames of the user will be interpolated with that same frame until the user ends the phoneme. This process is illustrated in Fig. 10.25.

The frame used as a loop frame requires a good spectral shape and, if possible, a pitch very close to the note that corresponds to that phoneme. Since we keep
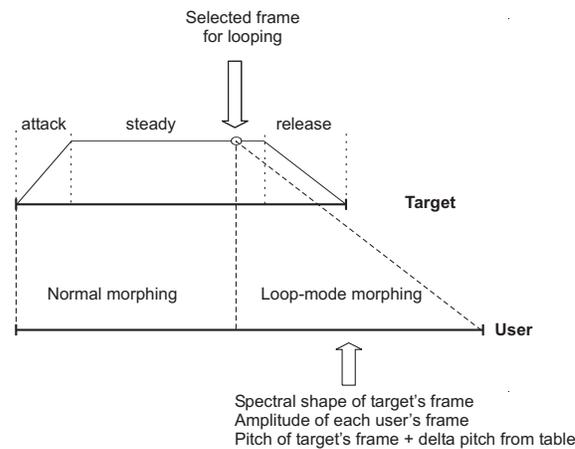
Figure 10.25 Loop synthesis diagram.

a constant spectral shape, we have to do something to make the synthesis sound natural. The way we do it is by using some "natural" templates obtained from the analysis of a longer phoneme that are then used to generate more target frames to morph with the loop frame. For example, one feature that adds naturalness is pitch variations of a steady state note sung by the same target. These delta pitches are kept in a lookup table whose first access is random and consecutive values are read afterwards. Two tables are kept, one with variations of steady pitch and another one with vibrato to generate target frames.

Once all the chosen parameters have been interpolated in a given frame, they are added back to the basic synthesis frame of the user. The synthesis is done with the standard synthesis procedures of SMS.

10.5.2   Time Scaling

Time-scaling an audio signal means changing the length of the sound without affecting other perceptual features, such as pitch or timbre. Many different techniques, both in time and frequency domain, have been proposed to implement this effect. Some frequency domain techniques yield high-quality results and can work with large scaling factors. However, they are bound to present some artifacts, like phasiness, loss of attack sharpness and loss of stereo image. In this section we will present a frequency domain technique for near loss-less time-scale modification of a general musical stereo mix [Bon00].

The basic system

The general block diagram of the system is represented in Fig. 10.26. First, the input sound is windowed and applied to the FFT which gives the analysis frame AF_n,

that is, the spectrum bins and the amplitude and phase envelopes (n is the analysis frame index). Then the time scaling module generates the synthesis frame SF_m that is fed to the inverse FFT (IFFT, m is the synthesis frame index). Finally, the windowing & overlap-add block divides the sound segment by the analysis window and multiplies it by the overlap-add window, to reconstruct the output sound. The basics of the FFT/IFFT approach are detailed in chapter 8.
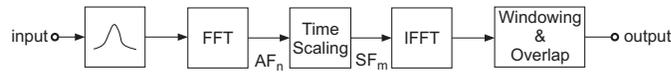

Figure 10.26 General diagram.

It is important to remark that the frame rate used in both the analysis and synthesis modules is the same, as opposed to the most broadly used time scaling techniques in which a change of frame rate in synthesis is used in order to achieve the effect. The window size and type must also be the same in both processes.

Figure 10.27 illustrates the operations for a time-scale stretching factor $TS > 1$, and a time compression factor $TS < 1$. The horizontal axis corresponds to the time of the center of the frame in the input audio signal. Therefore, when $TS > 1$, the time increments relative to the input audio signal will be shorter in the synthesis than in the analysis frames, but the actual frame rate will be exactly the same. Each synthesis frame points to the nearest look-ahead analysis frame. In some
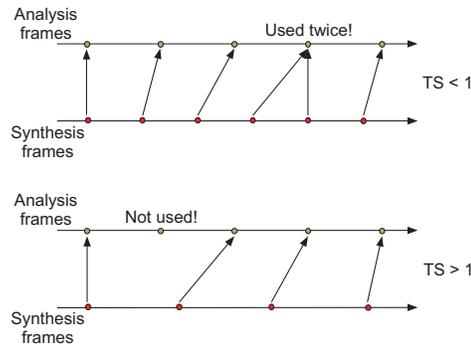

Figure 10.27 Analysis and synthesis frames.

cases, as shown in Fig. 10.27, an analysis frame is used twice (or more) while in other cases some frames are never used. This technique will not add any artifacts, provided the frame size we use is small enough and the sound does not present abrupt changes in that particular region. In the case of a percussive attack, though, a frame repetition or omission can be noticed regardless of the analysis frame size. Therefore, some knowledge of the features of a sound segment is needed to decide where this technique can or cannot be applied.

Figure 10.28 shows a detailed block diagram of the time scaling module. The analysis frames AF_n, containing the spectrum amplitude and phase envelopes, are

fed to the time-scaling module. This module performs a peak detection and a peak continuation algorithm (see 10.3.1) on the current and previous z^{-1} amplitude envelopes. Then, only the peaks that belong to a sinusoidal track are used as inputs to the spectrum phase generation module. Note that the time scaling module only changes the phase, leaving the spectral amplitude envelope as it is.
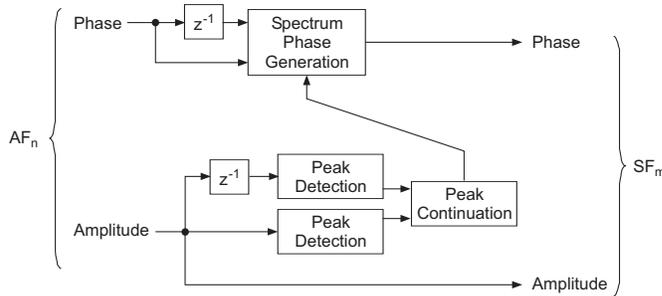
Figure 10.28 The time scaling module.

The phase of each peak is computed supposing that the frequency varies linearly between two consecutive frames and that there is some phase deviation $\Delta\varphi$ (see Fig. 10.29). The usage of the same frame rate in analysis and synthesis allows us to suppose that the phase variation between two consecutive frames is also the same.
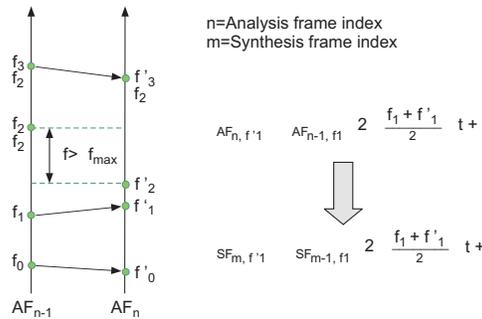
Figure 10.29 Peak continuation and phase generation.

Common problems and solutions in time-scaling

Chapter 8 already introduced a spectral technique for time scaling based on the phase vocoder approach. This kind of implementation presents very well known artifacts. In this section we will describe some of these problems and the solution that the implementation we are proposing can provide.

- Phasiness

In the phase vocoder implementation, the original frame has a flat phase envelope around the peak because of the circular convolution of the analysis window with the sinusoid. But after time scaling is applied, the phase looses its original behavior. This artifact is introduced due to the fact that the phase of each bin advances at different speed (see section 8.4.3). This loss of peak's phase coherence is known as *phasiness*. To avoid this problem we can apply the original relative behavior of the phase around the peak. As pointed out in [Lar97], each peak subdivides the spectrum into a different region, with a phase related to that of the peak. The phase around each peak is obtained applying the delta phase function of the original spectrum phase envelope (see Fig. 10.30).
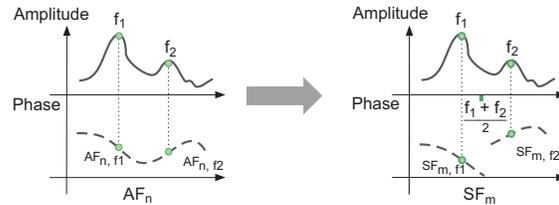

Figure 10.30 Original delta phase function around each peak.

- Loss of attack transients

Another typical artifact of the phase vocoder approach is the smoothing of the attack transients. A possible solution is to modify the sinusoidal plus residual model in order to model these transients [Ver98]. Another possible approach is not to time-scale the input signal on this kind of regions so that the original timing is respected (see Fig. 10.31). Consequently, and in order to preserve the overall scaling factor, a greater amount of scaling should be applied to surrounding regions.
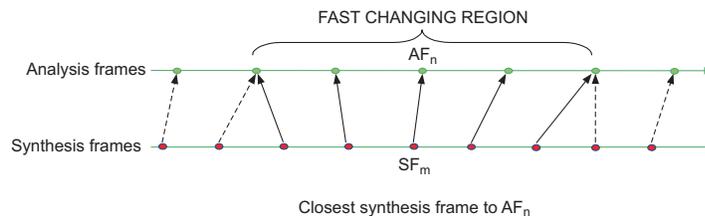

Figure 10.31 Attack transient region.

In order to apply the previous technique, it is necessary to detect attack transients of the sound in an unsupervised manner. The computation of relative changes of energy along several frequency bands can be used for that purpose. A low fre-

quency band could, for example, detect sharp bass notes, while a high frequency band could be set to detect hits of a crash cymbal.

The spectrum of the input signal is given by

$$X(sR_a, k) = |X(sR_a, k)| \cdot e^{j\varphi(sR_a, k)} \tag{10.23}$$

where the FFT has been sampled every R_a samples in time, and s is the time index of the short-term transform. If we define a set of frequency bands $B_i(k)$, then the energy of the i^{th} band can be computed as

$$E(s, i) = \sum_{i=0}^{N-1} B_i(k) \cdot X^2(sR_a, k) \tag{10.24}$$

and the relative change of energy $C(s, i)$ at frame s as,

$$C(s, i) = \frac{-2E(s-2, i) - E(s-1, i) + E(s+1, i) + 2E(s+2, i)}{E(s, i)}. \tag{10.25}$$

The maxima of $C(s, i)$ over some threshold should then indicate the attack transients of the input signal at the desired band.

- Frequency versus time resolution

As explained in 10.3.1, it is desirable to have long windows in order to achieve a high frequency resolution, but also to have short windows so to achieve a better time resolution. If the audio signal presents an important low frequency component, the use of a long window is a must, because the low frequency peaks will be too close to be distinguishable if we use a short window. On the other hand if we apply a very long window, the time-scaling process will add reverb and will smooth the sound.

The solution proposed is to use parallel windowing, that is, several analysis channels (see Fig. 10.32). Each channel is the result of an FFT with a specific window size, window type and zero-padding. Obviously, the window should be longer for low frequencies than for high frequencies. The peak detection process is applied to each of the channels while the peak continuation takes care of the desired channel frequency cuts, so it can connect peaks of different channels. Then the time scaling module fills the spectrum of all the channels (*amplitude and phase envelopes*) and applies a set of parallel filters $H_k(f)$ that must add up to a constant (*allpass filter*).

If the cutoff frequency of a channel was close to a spectral peak, this would be broken apart into two different channels and we would be introducing some kind of artifacts. For that reason, and in order to guarantee that phase and amplitude envelopes around the peak behave the way we expect, we need to provide our system with time-varying frequency cuts. Each frequency cut is computed as the middle point between the two closest peaks to the original frequency cut (see Fig. 10.33).
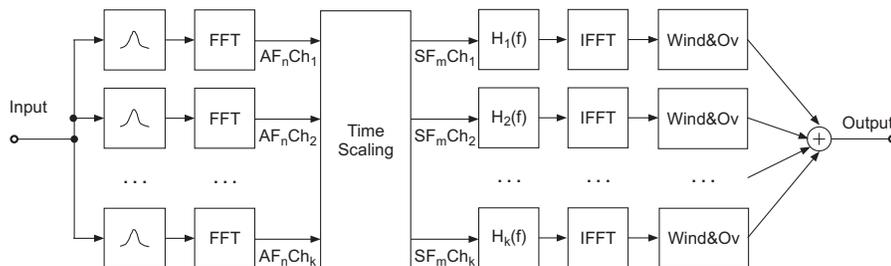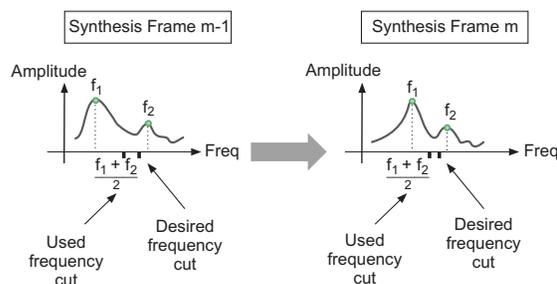
Figure 10.32 Multiple parallel windowing.

Figure 10.33 Variable phase frequency cut.

- Loss of stereo image

In the case of stereo signals, if we process each one of the channels independently, most of the stereo image is bound to be lost. This artifact is mainly due to the fact that the time scaling process changes the phase relationship between the two channels. Therefore, if we want to keep the stereo image, it is necessary to preserve the phase and amplitude relationship between left and right channels.

The fact that the system does not change the amplitude envelope of the spectrum guarantees that the amplitude relationship between channels will be preserved, provided we always use frames with identical time tags for both channels. For that purpose, we need to synchronize the attack transients between the two channels.

Figure 10.34 shows the simplified block diagram of the stereo time scaling system. Notice that the number of FFT and IFFT operations is multiplied by two and, as a consequence, the same happens to the processing time.

Time-varying time scaling

The system presented can deal with time-varying scaling factors with no loss of quality tradeoff. The only significant change is that the time increments of the synthesis frames in the input signal are not constant.

The application of time-varying tempo variations opens up many new and interesting perspectives. The system could be easily adapted and used for alignment and
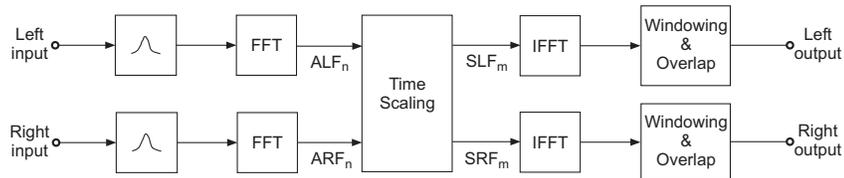
Figure 10.34 Stereo time scaling.

synchronization of two sound sources. Also, the amount of time scaling could be used in a wise way to inspire emotions. For example, to increase the climax or the suspense of a musical piece, by slowing or increasing the tempo during certain fragments. Another interesting application could be to control the time scaling factor in the same way as the orchestra conductor does and play in real-time a previously recorded background with a live performance.

10.6   Conclusion

Throughout this chapter, we have shown how the use of higher-level spectral models can lead to new and interesting sound effects and transformations. We have also seen that it is not easy nor immediate to get a good spectral representation of a sound, so the usage of this kind of approach needs to be carefully considered bearing in mind the application and the type of sounds we want to process.

For example, most of the techniques presented here work well only on monophonic sounds and some rely on the pseudo-harmonicity of the input signal.

Nevertheless, the use of spectral models for musical processing has not been around too long and it has already proven useful for many applications, as the ones presented in this chapter. Under many circumstances, higher-level spectral models, such as the Sinusoidal plus Residual, offer much more flexibility and processing capabilities than more immediate representations of the sound signal.

In general, higher-level sound representations will offer more flexibility at the cost of a more complex and time-consuming analysis process. It is important to remember that the model of the sound we choose will surely have great effect on the kind of transformations we will be able to achieve and on the complexity and efficiency of our implementation. Hopefully, the reading of this chapter, and the book as a whole, will guide the reader towards making the right decisions in order to get the desired results.

Bibliography

[Abe88]   M. Abe, S. Nakamura, K. Shikano and H. Kuwabara, Voice Conversion through Vector Quantization. Proc. *ICASSP-1988*, pp. 655-658, 1988.

[Arc97]   J.L. Arcos, R. López de Mántaras and X. Serra. Saxex: a Case-Based
          Reasoning System for Generating Expressive Musical Performances. Proc.
          International Computer Music Conference, pp. 25-30, 1997.

[Arc98]   J.L. Arcos, R. López de Mántaras and X. Serra. Saxex: a Case-Based Rea-
          soning System for Generating Expressive Musical Performances. *Journal
          of New Music Research*, Vol. 27, No. 3, pp. 194-210, Sept. 1998.

[Bon00]   J. Bonada. Automatic technique in frequency domain for near-lossless
          time-scale modification of audio. Proc. *International Computer Music
          Conference*, pp. 396-399, 2000.

[Can98]   P. Cano. Fundamental Frequency Estimation in the SMS Analysis. Proc.
          Digital Audio Effects Workshop (DAFX98), Barcelona, pp. 99-102, Novem-
          ber 1998.

[Can00]   P. Cano, A. Loscos, J. Bonada, M. de Boer and X. Serra. Voice Morphing
          System for Impersonating in Karaoke Applications. Proc. *International
          Computer Music Conference*, pp. 109-112, 2000.

[Chi94]   D.G. Childers. Measuring and Modeling Vocal Source-Tract Interaction.
          IEEE Transactions on Biomedical Engineering, Vol. 41, No. 7, pp. 663-671,
          July 1994.

[Cox71]   M.G. Cox. An algorithm for approximating convex functions by means of
          first-degree splines. *Computer Journal*, Vol. 14, pp. 272–275, 1971.

[Dep93]   Ph. Depalle, G. Garcia and X. Rodet. Analysis of Sound for Additive
          Synthesis: Tracking of Partials Using Hidden Markov Models. Proc. *In-
          ternational Computer Music Conference*, pp. 94-97, 1993.

[Dep97]   Ph. Depalle and T. Hélie. Estraction of Spectral Peak Parameters using a
          Short-Time Fourier Transform Modeling and no Sidelobe Windows. *Pro-
          ceedings of the 1997 IEEE Workshop on Applications of Signal Processing
          to Audio and Acoustics*, Monhonk, pp. 298-231, 1997.

[Din97]   Y. Ding and X. Qian. Sinusoidal and Residual Decomposition and Resid-
          ual Modeling of Musical Tones Using the QUASAR Signal Model. Proc.
          International Computer Music Conference, pp. 35-42, 1997.

[Fit00]   K. Fitz, L. Haken and P. Christensen. A New Algorithm for Bandwidth
          Association in Bandwidth-Enhanced Additive Sound Modeling. Proc. *In-
          ternational Computer Music Conference*, pp. 384-387, 2000.

[Goo96]   M. Goodwin. Residual Modeling in Music Analysis-Synthesis. *Proc. IEEE
          ICASSP-1996*, pp. 1005-1008, Atlanta, 1996.

[Goo97]   M. Goodwin. *Adaptive Signal Models: Theory, Algorithms and Audio
          Applications*. Ph.D. Dissertation, University of California. Berkeley, 1997.

[Har78]   F.J. Harris. On the use of windows for harmonic analysis with the discrete Fourier transform. *Proceedings IEEE*, Vol. 66, pp. 51-83, 1978.

[Her98]   P. Herrera and J.Bonada. Vibrato Extraction and Parameterization in the Spectral Modeling Synthesis Framework. Proc. *Digital Audio Effects Workshop (DAFX98)*, pp. 107-110, Barcelona, November 1998.

[Hes83]   W. Hess. *Pitch Determination of Speech Signals*. New York, Springer-Verlag, 1983.

[Lar97]   J. Laroche and M. Dolson. About this phasiness business. Proc. *International Computer Music Conference*, pp. 55-58, 1997.

[Los99]   A. Loscos, P. Cano and J. Bonada. Low-delay Singing Voice Alignment to Text. Proc. *International Computer Music Conference*, pp. 437-440, 1999.

[Mah94]   R.C. Maher and J.W. Beauchamp. Fundamental Frequency Estimation of Musical Signals using a two-way Mismatch Procedure. *Journal of the Acoustical Society of America*, 95(4):2254-2263, 1994.

[Mak75]   J. Makhoul. Linear Prediction: A Tutorial Review. *Proceedings of the IEEE*, Vol. 63, No. 4, pp. 561-580, 1975.

[Mar75]   J.D. Markel and A.H. Gray. *Linear Prediction of Speech*. Springer-Verlag, 1975.

[McA86]   R.J. McAulay and T.F. Quatieri. Speech Analysis/Synthesis based on a Sinusoidal Representation. *IEEE Transactions on Acoustics, Speech and Signal Processing*, 34(4):744–754, 1986.

[Osa95]   N. Osaka. Timbre Interpolation of sounds using a sinusoidal model. Proc. *International Computer Music Conference*, pp. 408-411, 1995.

[Rod92]   X. Rodet and Ph. Depalle. *Spectral Envelopes and Inverse FFT Synthesis*. Proc. *93rd AES Convention*. San Francisco, AES Preprint No. 3393 (H-3), October 1992.

[Ros98]   S. Rossignol and others. Feature Extraction and Temporal Segmentation of Acoustic Signals. Proc. *International Computer Music Conference*, 1998.

[Sed88]   R. Sedgewick. *Algorithms*. Reading, Massachusetts: Addison-Wesley, 1988.

[Ser89]   X. Serra. *A System for Sound Analysis/Transformation/Synthesis based on a Deterministic plus Stochastic Decomposition*. Ph.D. Dissertation, Stanford University, 1989.

[Ser90]   X. Serra and J. Smith. Spectral Modeling Synthesis: A Sound Analysis/Synthesis System based on a Deterministic plus Stochastic Decomposition. *Computer Music Journal*, Vol. 14, No. 4, pp. 12-24, 1990.

[Ser94]   X. Serra. Sound Hybridization Techniques based on a Deterministic plus Stochastic Decomposition Model. Proc. *International Computer Music Conference*, pp. 348-351, 1994.

[Ser96]   X. Serra. Musical Sound Modeling with Sinusoids plus Noise. in G. D. Poli, A. Picialli, S. T. Pope, and C. Roads, editors, *Musical Signal Processing*. pp. 91-122, Swets & Zeitlinger Publishers, 1996.

[Ser98]   X. Serra and J. Bonada. Sound Transformations Based on the SMS High Level Attributes. Proc. *Digital Audio Effects Workshop (DAFX98)*, pp. 138-142, Barcelona, November 1998.

[Sla96]   M. Slaney, M. Covell and B. Lassiter. Automatic audio morphing. *Proc. IEEE ICASSP-1996*, pp. 1001-1004, 1996.

[Smi87]   J.O. Smith and X. Serra. PARSHL: An Analysis/Synthesis Program for Non-Harmonic Sounds based on a Sinusoidal Representation. Proc. *International Computer Music Conference*, pp. 290-297, 1987.

[SMS]   `www.iua.upf.es/~sms`

[Str80]   J. Strawn. Approximation and Syntactic Analysis of Amplitude and Frequency Functions for Digital Sound Synthesis. *Computer Music Journal*, 4(3):3–24, 1980.

[Tel95]   E. Tellman, L. Haken and B. Holloway. Timbre morphing of sounds with unequal numbers of features. *Journal of the Audio Engineering Society*, 43(9), 678-89, 1995.

[Ver98]   T.S. Verma and T.H.Y. Meng. Time Scale Modification Using a Sines+Transients+Noise Signal Model. Proc. *Digital Audio Effects Workshop (DAFX98)*, pp. 49-52, Barcelona, November 1998.

[Ver00]   T.S. Verma and T.H.Y. Meng. Extending Spectral Modeling Synthesis With Transient Modeling Synthesis. *Computer Music Journal*, 24:2, pp.47-59, 2000.

[Vid90]   E. Vidal and A. Marzal. A Review and New Approaches for Automatic Segmentation of Speech Signals. L. Torres and others (eds.), *Signal Processing V: Theories and Applications*, pp. 43-53, Elsevier Science Publishers, 1990.