

Figure 6.25 Waveforms in Casio phase distortion synthesis. The audio waveform changes by varying the rate at which a sine lookup table is read out. (a) A constant rate of readout generates a sine wave. (b) A readout whose rate changes twice per cycle, distorting the sine wave into a quasi-sawtooth waveform.

Waveshaping Synthesis

Jean-Claude Risset, working at the Bell Telephone Laboratories in New Jersey, carried out the first experiments with the method now known as waveshaping synthesis (Risset 1969). Daniel Arfib (1979) and Marc LeBrun (1979) independently developed theoretical and empirical elaborations of the basic method. Waveshaping is musically interesting because, as in FM synthesis, it gives us a simple handle on the time-varying bandwidth and spectrum of a tone in a computationally efficient way.

The fundamental idea behind waveshaping (also known as *nonlinear distortion*) is to pass a sound signal x through a “distortion box.” In digital form, the distortion box is a function in a stored table (or array) in computer memory. The function w maps any input value x in the range $[-1, +1]$ to an output value $w(x)$ in the same range.

In the simplest case, x is a sinusoidal wave generated by an oscillator. However, x can be any signal, not just a sinusoid. For each output sample to be computed we use the value of x to index table w . Table w contains the *shaping function* (also called the *transfer function*). We then simply take the value in w indexed by x as our output value $w(x)$.

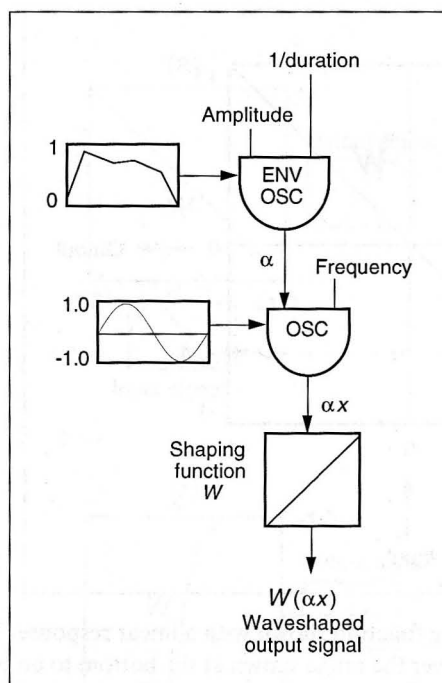


Figure 6.26 Simple waveshaping instrument. A sinusoidal oscillator, whose amplitude is controlled by the amplitude envelope signal α , indexes a value in the shaping function table w . As in other example instruments, the input $1/\text{duration}$ that is fed into the frequency input of the envelope oscillator indicates that it goes through one cycle over the duration of the note.

Simple Waveshaping Instrument

An instrument for simple waveshaping synthesis is shown in figure 6.26. Here an envelope oscillator controls the amplitude of a sinusoidal oscillator that is fed into a shaping function table. The amplitude envelope α is important because it has the effect of scaling the input signal, making it reference different regions of the shaping function w . We look at the implications of this next.

Example Shaping Functions

As figure 6.27 shows, if the shaping function in table w is a straight diagonal line from -1 to $+1$, the output of w is an exact replica of its input x . This is because the table maps an input of -1 (shown at the bottom of the function) to -1 in the output (shown at the right of the function), 0 maps




Figure 6.27 Shaping function shown with a linear response. The function maps an input signal scaled over the range shown at the bottom to an output function whose scale is shown at the right. To see how the function maps an input to an output value, read vertically from the bottom and then look to the right to see the corresponding output value. Thus an input value of -0.4 on the bottom maps to an output value of -0.4 on the right. This equivalence between the input and the output is only true for a linear shaping function.

to 0, 1 maps to 1, and so on. Because this simple relationship between the input and the output occurs only when the shaping function is a straight line, we say that in this case the output is a *linear function* of its input.

If the shaping table contains anything other than a straight diagonal line from -1 to $+1$, x is *distorted* by the shaping function in w . Figure 6.28 shows the effects of several shaping functions on an input sinusoid. Figure 6.28a shows an *inverting* shaping function. For every positive value of the input amplitude the waveshaper emits a correspond negative value and vice versa. Figure 6.28b is a straight line but with a narrower angle than the curve in figure 6.27. It maps to a narrower range on the right-hand (output) side of the shaping function, meaning that it attenuates the input signal. Figure 6.28c expands low-level signals and sends high-level signals into clipping distortion. The amplitude-sensitive nature of waveshaping is demonstrated well in figure 6.28d. The shaping function is a straight line around zero, which is the low-amplitude portion of the grid. Such a function passes a low-amplitude input signal with no distortion. When the amplitude of the




Figure 6.28 Four shaping functions. (a) Inversion of the input signal. (b) Attenuation. (c) Amplification of low-level signals (expansion) and clipping of high-level signals. (d) Complicated amplitude-sensitive distortion.

input signal increases, the extreme ends of the shaping function subject the input signal to a complicated form of distortion.

Amplitude Sensitivity of Waveshaping Spectrum

It is easy to see that the amplitude sensitivity of waveshaping can model a characteristic of acoustic instruments. That is, when one plays an acoustic instrument “harder,” for example, by strumming a guitar forcefully, blowing a saxophone stridently, or striking a drum intensely, this enriches the spectrum. In waveshaping we can emulate this effect by passing a signal whose overall amplitude varies with time through the shaping function. As the amplitude of the input signal varies, one obtains a correspondingly

time-varied spectrum at the output. Put another way, a variation in the time domain at the input is manifest as a variation in the frequency domain at the output. This is an important feature. Given a single shaping function (precomputed and stored in memory), a variety of output waveforms can be obtained simply by varying the amplitude or offset of the input signal in order to apply various regions of the shaping function. Hence waveshaping is an especially efficient synthesis technique. Arfib (1979) gives practical examples of the waveshaping technique in specific musical applications.

Chebyshev Shaping Functions

Research by LeBrun (1979) and Arfib (1979) demonstrated that it is possible to predict exactly the output spectrum of the waveshaping technique under mathematically controlled conditions. By restricting the signal x to an unvarying cosine wave and using a family of smooth polynomials called *Chebyshev functions*, which take values in the range $[-1, +1]$ to construct the shaping function w , one can produce easily any desired combination of harmonics in a steady-state spectrum. This derives from the following identity:

$$T_k \times (\cos[\theta]) = \cos(k \times \theta)$$

where T_k is the k th Chebyshev function. In other words, by applying the k th Chebyshev polynomial to an input sine wave, we obtain a cosine wave at the k th harmonic. This means that each separate Chebyshev polynomial, when used as the shaping function, produces a particular harmonic of x . By summing a weighted combination of Chebyshev polynomials and putting the result in the shaping table, a corresponding harmonic mixture is obtained as the output of the waveshaping technique. For example, to obtain a steady-state waveform with a first harmonic (fundamental), a second harmonic that is 0.3 the amplitude of the first harmonic, and a third harmonic that is 0.17 of the first harmonic, we add the equations

$$T_0 + (0.3 \times T_2) + (0.17 \times T_3),$$

and we put the result into the transfer function wavetable. If a cosine wave is passed through this table, then the output spectrum contains the desired harmonic ratios.

An advantage of using the Chebyshev functions is that we can guarantee that the output of the waveshaper is bandlimited. That is, it does not contain frequencies above the Nyquist rate, and therefore it is free of foldover distortion. Table 6.1 lists the equations for T_0 through T_8 where $x = \cos(q)$.

Table 6.1 Chebychev functions T_0 through T_8

$T_0 = 1$
$T_1 = x$
$T_2 = 2x^2 - 1$
$T_3 = 4x^3 - 3x$
$T_4 = 8x^4 - 8x^2 + 1$
$T_5 = 16x^5 - 20x^3 + 5x$
$T_6 = 32x^6 - 48x^4 + 18x^2 - 1$
$T_7 = 64x^7 - 112x^5 + 56x^3 - 7x$
$T_8 = 128x^8 - 256x^6 + 160x^4 - 32x^2 + 1$

Amplitude Normalization

The main drawback of waveshaping synthesis is that the output amplitude of the simple waveshaping instrument shown in figure 6.28 varies considerably, even using only one shaping function. This variance is the result of different parts of the shaping function being applied. That is, it depends on the amplitude of the input signal to the shaping function.

In waveshaping the amplitude of x is actually used to control timbre, not the overall loudness of the sound. If we want full independence between timbre and the output amplitude, some form of amplitude normalization is required. At least three kinds of normalization are possible: loudness normalization, power normalization, and peak normalization.

For musical purposes, our ideal would be loudness normalization, in which the perceived loudness of the instrument is constant for all values of α . However, this involves complicated psychoacoustic interactions and context-dependencies, so it is too difficult and computationally expensive for most implementations. Power normalization is based on division by the *root mean square* (RMS) of the harmonic amplitudes generated by a particular shaping function. LeBrun (1979) gives details on this technique. Peak normalization is probably the least complicated and most practical of the three. It is accomplished by scaling the output in relation to the maximum value. *Peak normalization ensures that the output amplitude of different tones will at least have the same peak value, and will therefore not overload the digital-to-analog converters with a value out of their range.*

Figure 6.29 shows a waveshaping instrument with a peak normalization. The easiest way to do this is to prepare a table containing normalization factors for all values of α , since the envelope a determines the amplitude of x . For example, if the value of α input to the normalization table is 0.7, we

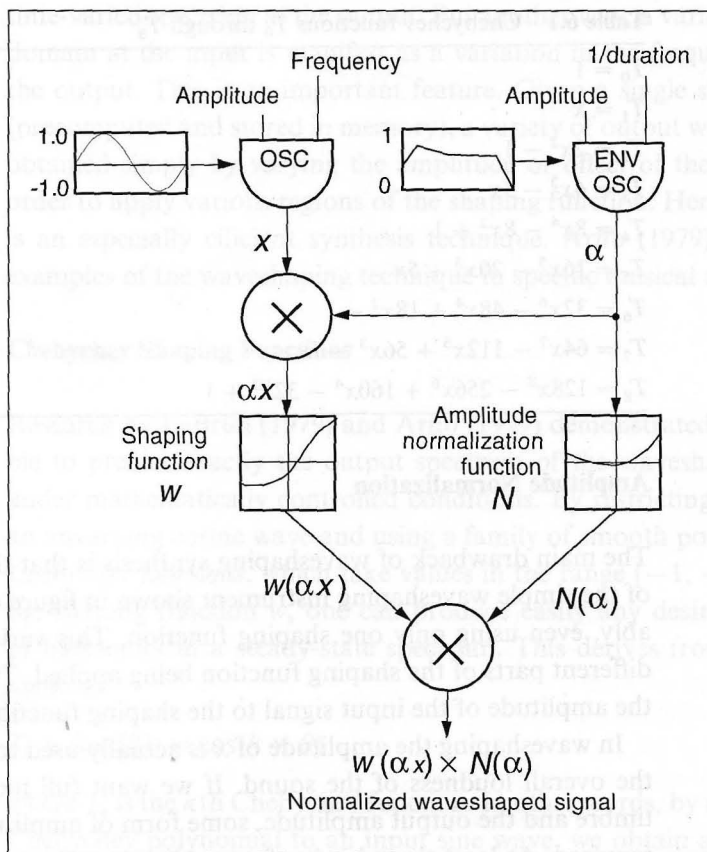


Figure 6.29 Waveshaping instrument with a normalization section. The value of α indexes a value in the normalization table that scales the output of the waveshaper.

multiply the output by the entry in the normalization table corresponding to α .

Variations on Waveshaping

The classic waveshaping technique—sending a cosine wave through a Chebychev polynomial shaping function—produces a range of harmonic spectra. We can extend this class of waveshaping spectra by changing the input or the shaping function. Another possibility is modifying the signal coming out of the waveshaping instrument by another signal-processing device, such as a filter.

As mentioned previously, the input x to the waveshaper can be any signal, not just a cosine wave. Reinhard (1981), for example, details what happens when x is the sum of two cosine waves at different frequencies.

Another variation is to use a frequency-modulated signal for the input x . The benefit of this is that one can obtain inharmonic combinations of partials and formant structures (Arfib 1979).

The signal x can also be a sampled or concrete sound. When the shaping function w is a simple and smooth polynomial, the effect is not unlike phasing, since the harmonics of the input undulate in a time-varying way. Hence a waveshaping instrument can generate an interesting hybrid of natural and electronic sound. If w contains any straight horizontal or vertical lines the effect is a strong distortion, like the distortion of transistorized electric guitar amplifiers turned to maximum.

Neither does w have to be a Chebychev polynomial. The main benefit of using Chebychev polynomials as shaping functions is that the output is bandlimited and is therefore not subject to foldover distortion. But if this benefit is not paramount, w can be constructed out of other kinds of equations. It can also be drawn by hand (Buxton et al. 1982). See chapter 8 for an account of waveshaping with a noise-modulated shaping function.

Movable Waveshaping

Another variation is called *movable waveshaping*, invented by Xin Chong at the Beijing Central Music Conservatory (Xin 1987). In this technique the shaping function itself varies with time. This can be accomplished by storing a longer shaping function and moving an index to scan various parts of it at different times. Starting from simple input signals and simple time-varying shaping functions, a multiplicity of results can be obtained.

Fractional Waveshaping

De Poli (1984) analyzed a configuration in which the shaping function is a fraction, specifically, a ratio between two polynomials. He calls this *fractional waveshaping*. Fractional waveshaping can generate such effects as exponential spectra and spectra whose shapes resemble a damped cosine wave. The multiple bumps of the damped cosine wave spectrum are heard as formants. Dynamically varying spectra are achieved as in regular waveshaping by varying the amplitude and offset of the input cosine signal.

Postprocessing and Parameter Estimation

The waveshaped signal can be passed through another signal processing device, what we could call *postprocessing* the waveshaped signal. This device

could be, for example, an AM oscillator, a FM oscillator, or a filter. AM and FM can enrich the waveshaping spectrum by adding, for example, in-harmonic partials to a harmonic spectrum (Arfib 1979; Le Brun 1979; De Poli 1984).

De Poli (1984) and Volonnino (1984) developed an experimental filtering method called *frequency-dependent waveshaping*. This was aimed at providing independent control of the phase and amplitude of each harmonic generated by the waveshaping process. See the cited literature for more details on these techniques.

Beauchamp (1979) inserted a highpass filter on the output of his waveshaping model of brass tones in order to mimic the damping effects of brass pipes. More recently, Beauchamp and Horner (1992) have simulated instrumental tones by a multiple waveshaper + filter model. They first perform a parameter estimation of an instrumental tone and approximate its spectrum with a single waveshaper + filter model. They subtract this approximation from the original sound to obtain a difference or *residual* signal. They then approximate the residual with another waveshaper + filter model. Using two or even three waveshaping models in this way results in much closer simulations than a single model.

General Modulations

Many synthesis techniques can be turned into modulation techniques by substituting a time-varying function for a constant term in the equation of the original technique. If the time-varying function is periodic, the technique is one of a family of synthesis techniques known under the rubric of *waveshape parameter modulation*. For example, amplitude modulation and frequency modulation can be classified as waveshape parameter modulation techniques. For more on this classification scheme see Mitsuhashi (1980).

James A. Moorer (1976) showed that the equation for single FM is one instance of a general class of equations called *discrete summation formulas* (DSFs). DSFs refers to a set of formulas that are the *closed form* solution of the sums of finite and infinite trigonometric series. By “closed form” is meant a more compact and efficient representation of a longer summation formula. These formulas are relevant to sound synthesis if we assume that they describe waveforms that are sums of sinusoidal waves. For example, the right-hand side of the following equation is the closed-form solution to the summation shown in the left-hand side: