PART 1
Technology and Theory

p0010

c0001 CHAPTER 1

The Science of Synthesis

p0010 **"Today's recording techniques would have been regarded as science fiction forty years ago." …**

p0020 Today's dance- and club-based music relies just as heavily on the technology as it does on the musicality; therefore, to be proficient at creating this genre of music it is first necessary to fully comprehend the technology behind its creation. Indeed, before we can even begin to look at how to produce the music, a thorough understanding of both the science and the technology behind the music is paramount. You wouldn't attempt to repair a car without some knowledge of what you were tweaking, and the same applies for dance- and club-based music.

p0030 Therefore, we should start at the very beginning and where better to start than the instrument that encapsulated the genre – the analogue synthesizer. Without a doubt, the analogue synthesizers were responsible for the evolution of the music, and whilst the early synthesizers are becoming increasingly difficult to source today, nearly all synthesizers in production, whether hardware or software, follow the same path first laid down by their predecessors. However, to make sense of the various knobs and buttons that adorn a typical synthesizer and observe the effects that each has on a sound, we need to start by examining some basic acoustic science.

s0010 ACOUSTIC SCIENCE

p0040 When any object vibrates, air molecules surrounding it begin to vibrate sympathetically in all directions creating a series of sound waves. These sound waves then create vibrations in the ear drum that the brain perceives as sound.

p0050 The movement of sound waves is analogous to the way that waves spread when a stone is thrown into a pool of water. The moment the stone hits the water, the reaction is immediately visible as a series of small waves spread outwards in every direction. This is almost identical to the way in which sound behaves, with each wave of water being similar to the vibrations of air particles.

SNOMAN 978-0-240-52107-7 00001
CH001.indd 3 8/25/2008 9:15:17 PM

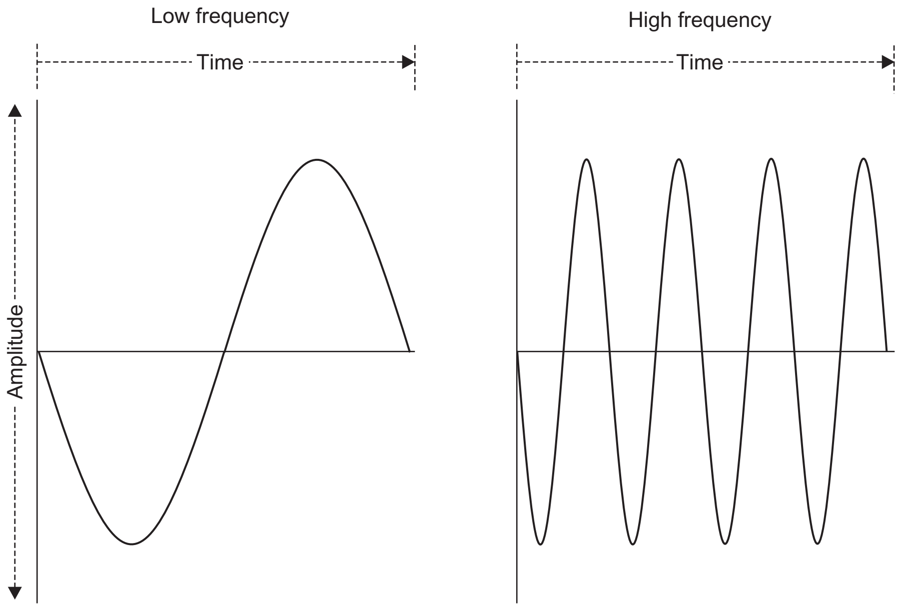

f0010 **FIGURE 1.1**
Difference between low and high frequencies

p0060 For instance, when a tuning fork is struck, the forks first move towards one another compressing the air molecules before moving in the opposite direction. In this movement from 'compression' to 'rarefaction' there is a moment where there are less air molecules filling the space between the forks. When this occurs, the surrounding air molecules crowd into this space and are then compressed when the forks return on their next cycle. As the fork continues to vibrate, the previously compressed air molecules are pushed further outwards by the next cycle of the fork and a series of alternating compressions and rarefactions pass through the air.

p0070 The numbers of rarefactions and compressions, or 'cycles', that are completed every second is referred to as the operating frequency and is measured in Hertz (Hz). Any vibrating object that completes, say, 300 cycles/s has a frequency of 300 Hz while an object that completes 3000 cycles/s has a frequency of 3 kHz.

p0080 The frequency of a vibrating object determines its perceived pitch, with faster frequencies producing sounds at a higher pitch than slower frequencies. From this we can determine that the faster an object vibrates, or 'oscillates', the shorter the cycle between compression and rarefaction. An example of this is shown in Figure 1.1.

p0090 Any object that vibrates must repeatedly pass through the same position as it moves back and forth through its cycle. Any particular point during this movement is referred to as the 'phase' of the cycle and is measured in degrees, similar to the measurement of a geometric circle. As shown in Figure 1.2, each cycle starts at position zero, passes back through this position, known as the 'zero crossing', and returns to zero.

f0020
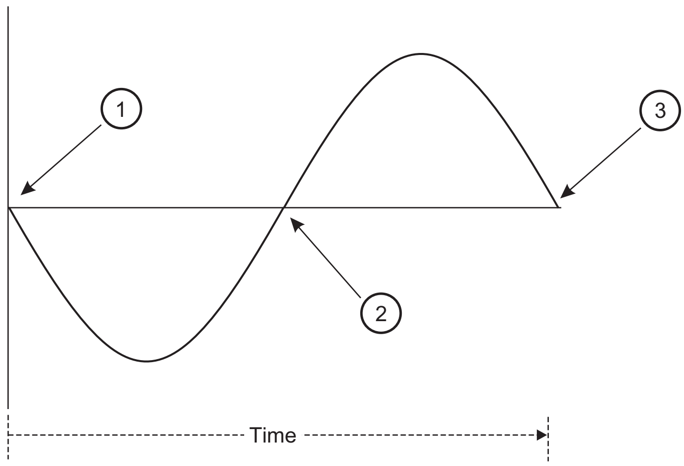

FIGURE 1.2
The zero crossing in a waveform

p0100 Consequently, if two objects vibrate at different speeds and the resulting waveforms are mixed together, both waveforms will start at the same zero point but the higher frequency waveform will overtake the phase of the lower frequency. Provided that these waveforms continue to oscillate, they will eventually catch up with one other and then repeat the process all over again. This produces an effect known as 'beating'.

p0110 The speed at which waveforms 'beat' together depends on the difference in frequency between them. It's important to note that if two waves have the same frequency and are 180° out of phase with one another there, one waveform reaches its peak while the second is at its trough, and no sound is produced. This effect, where two waves cancel one another out and no sound is produced, is known as 'phase cancellation' and is shown in Figure 1.3.

p0120 As long as waveforms are not 180° out of phase with one another, the interference between the two can be used to create more complex waveforms than the simple sine wave. In fact, every waveform is made up of a series of sine waves, each slightly out of phase with one another. The more complex the waveform this produces, the more complex the resulting sound. This is because as an increasing number of waves are combined a greater number of harmonics are introduced. This can be better understood by examining how an everyday piano produces its sound.

p0130 The strings in a piano are adjusted so that each oscillates at an exact frequency. When a key is struck, a hammer strikes the corresponding string forcing it to oscillate. This produces the fundamental pitch of the note and also, if the vibrations from this string are the same as any of the other strings' natural vibration rates, sets these into motion too. These are called 'sympathetic vibrations' and are important to understand because most musical instruments are based around this principle. The piano is tuned so that the strings that vibrate sympathetically with the originally struck string create a series of waves that are slightly out of phase with one another producing a complex sound.

p0140 Any frequencies that are an integer multiple of the lowest frequency (i.e. the fundamental) will be in harmony with one another, a phenomenon that was first realized by Pythagoras, from which he derived the following three rules:

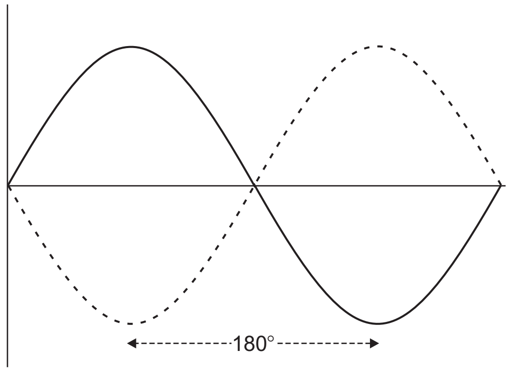

f0030 **FIGURE 1.3**
Two waves out of phase

- u0010 If a note's frequency is multiplied or divided by two, the same note is created but in a different octave.
- u0020 If a note's frequency is multiplied or divided by three, the strongest harmonic relation is created. This is the basis of the western musical scale. If we look at the first rule, the ratio 2:3 is known as a perfect fifth and is used as the basis of the scale.
- u0030 If a note's frequency is multiplied or divided by five, this also creates a strong harmonic relation. Again, if we look at the first rule, the ratio 5:4 gives the same harmonic relation but this interval is known as the major third.

p0180 A single sine wave produces a single tone known as the fundamental frequency, which in effect determines the pitch of the note. When further sine waves that are out of phase from the original are introduced, if they are integer multiples of the fundamental frequency they are known as 'harmonics' and make the sound appear more complex, otherwise if they are not integer multiples of the fundamental they are called 'partials', which also contribute to the complexity of the sound. Through the introduction and relationship of these harmonics and partials an infinite number of sounds can be created.

p0190 AUQ1 AUQ2 As Figure 1.4 shows, the harmonic content or 'timbre' of a sound determines the shape of the resulting waveform. Because these diagrams are very simple and the waveform produced by any instrument is very complex, it's difficult, if not impossible, to accurately reproduce it on paper.

p0200 In an attempt to overcome this, Joseph Fourier, a French scientist, discovered that no matter how complex any sound is it could be broken down into its frequency components and, using a given set of harmonics, it was possible to reproduce it in a simple form.

p0210 To use his words, "Every periodic wave can be seen as the sum of sine waves with certain lengths and amplitudes, the wave lengths of which have harmonic relations". This is based around the principle that the content of any sound is determined by the relationship between the level of the fundamental frequency and its harmonics and their evolution over a period of time. From this theory, known as the Fourier theorem, the waveforms that are common to most synthesizers are derived.

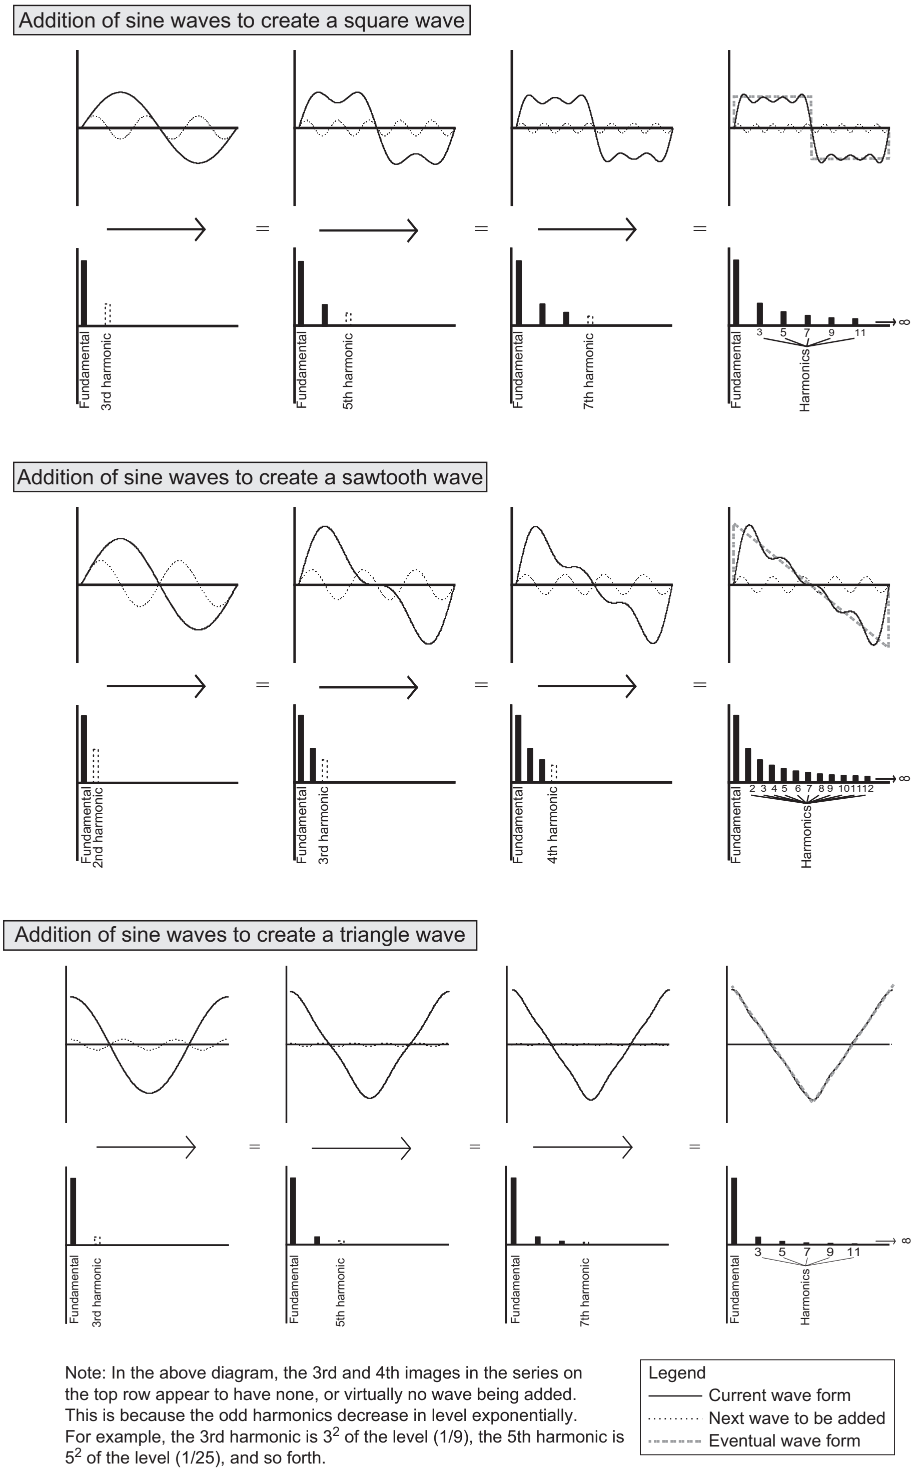

FIGURE 1.4
How multiple sound waves create harmonics

f0040

p0220 So far we've looked at how both the pitch and the timbre are determined. The final characteristic to consider is volume. Changes in volume are caused by the amount of air molecules an oscillating object displaces. The more air an object displaces, the louder the perceived sound. This volume, also called 'amplitude', is measured by the degree of motion of the air molecules within the sound waves, corresponding to the extent of rarefaction and compression that accompanies a wave. The problem, however, is that many simple vibrating objects produce a sound that is inaudible to the human ear because so little air is displaced; therefore, for the sound wave to be heard most musical instruments must amplify the sound that's created. To do this, acoustic instruments use the principle of forced vibration that utilizes either a sounding board, as in a piano or similar stringed instruments, or a hollow tube, as in the case of wind instruments.

p0230 When a piano string is struck, its vibrations not only set other strings in motion but also vibrate a board located underneath the strings. Because this sounding board does not share the same frequency as the vibrating wires, the reaction is not sympathetic and the board is forced to resonate. This resonance moves a larger number of air particles than the original sound alone, in effect amplifying the sound. Similarly, when a tuning fork is struck and placed on a tabletop, the table's frequency is forced to match that of the tuning fork and the sound is amplified.

p0240 Of course, neither of these methods of amplification offers any physical control over the amplitude. If the level of amplification can be adjusted, then the ratio between the original and the changed amplitude is called the 'gain'.

p0250 It should be noted, however, that loudness itself is difficult to quantify because it's entirely subjective to the listener. Generally speaking, the human ear can detect frequencies from as low as 20 Hz up to 20 kHz; however, this depends on a number of factors. Indeed, whilst most of us are capable of hearing (or more accurately feeling) frequencies as low as 20 Hz, the perception of higher frequencies changes with age. Most teenagers are capable of hearing frequencies as high as 18 kHz while the middle-aged tend not to hear frequencies above 14 kHz. A person's level of hearing may also have been damaged, for example, by overexposure to loud noise or music. Whether it is possible for us to perceive sounds higher than 18 kHz with the presence of other sounds is a subject of debate that has yet to be proven. However, it is important to remember that sounds that are between 3 and 5 kHz appear perceivably louder than frequencies that are out of this range.

s0020 SUBTRACTIVE SYNTHESIZER

p0260 Having looked into the theory of sound, we can look at how this relates to a synthesizer. Subtractive synthesizer is the basis of many forms of synthesizers f0050 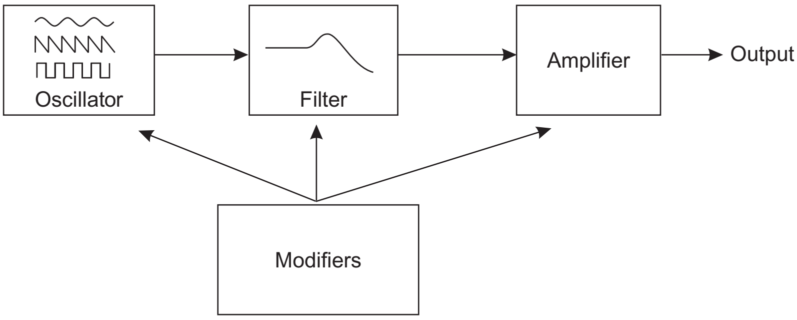

FIGURE 1.5
Layout of a basic synthesizer

and is commonly related to analogue synthesizer. It is achieved by combining a number of sounds or 'oscillators' together to create a timbre that is very rich in harmonics.

p0270 This rich sound can then be sculpted using a series of 'modifiers'. The number of modifiers available on a synthesizer is entirely dependent on the model, but all synthesizers offer a way of filtering out certain harmonics and of shaping the overall volume of the timbre.

p0280 The next part of this chapter looks at how a real analogue synthesizer operates, although any synthesizer that emulates analogue synthesizer (i.e. digital signal processing (DSP) analogue) will operate in essentially the same way, with the only difference being that the original analogue synthesizer voltages do not apply to their DSP equivalents.

p0290 An analogue synthesizer can be said to consist of three components (Figure 1.5):

u0040 ■ An oscillator to make the initial sound.
u0050 ■ A filter to remove frequencies within the sound.
u0060 ■ An amplifier to define the overall level of the sound.

Each of these components and their role in synthesizer are discussed in the sections below.

s0030 VOLTAGE-CONTROLLED OSCILLATOR (VCO)

p0330 When a key on a keyboard is pressed, a signal is sent to the oscillator to activate it, followed by a specific control voltage (CV) to determine the pitch. The AUQ3
CV that is sent is unique to the key that is pressed, allowing the oscillator to determine the pitch it should reproduce. For this approach to work correctly,

the circuitry in the keyboard and the oscillator must be incredibly precise in order to prevent the tuning from drifting, so the synthesizer must be serviced regularly. In addition, changes in external temperature and fluctuations in the power supply may also cause the oscillator's tuning to drift.

p0340 This instability gives analogue synthesizers their charm and is the reason why many purists will invest small fortunes in second-hand models rather than use the latest DSP-based analogue emulations. Although, that said, if too much detuning is present, it will be immediately evident and could become a major problem! There is still an ongoing argument over whether it's possible for DSP oscillators to faithfully reproduce analogue-based synthesizers, but the argument in favour of DSP synthesizers is that they offer more waveforms and do not drift too widely, and therefore prove more reliable in the long run.

p0350 In most early subtractive synthesizers the oscillator generated only three types of waveforms: square, sawtooth and triangle waveforms. Today this number has increased and many synthesizers now offer additional sine, noise, tri-saw, pulse and numerous variable wave shapes as well.

p0360 Although these additional waveforms produce different sounds, they are all based around the three basic wave shapes and are often introduced into synthesizers to prevent mixing of numerous basic waveforms together, a task that would reduce the number of oscillators.

p0370 For example, a tri-saw wave is commonly a sample of three sawtooth waves blended together to produce a sound that is rich in harmonics, with the advantage that the whole sound is contained in one oscillator. Without this waveform it would take three oscillators to recreate this sound, which could be beyond the capabilities of the synthesizer. Even if the synthesizer could utilize three oscillators to produce this one sound, the number of available oscillators would be reduced. Subsequently, while there are numerous oscillator waves available, knowledge of only the following six types is required.

s0040

The Sine Wave

p0380 A sine wave is the simplest wave shape and is based on the mathematical sine function (Figure 1.6). A sine wave consists of the fundamental frequency alone and does not contain harmonics. This means that they are not suitable for sole use in a subtractive sense, because if the fundamental is removed no sound is produced (and there are no harmonics upon which the modifiers could act). Consequently, the sine wave is used independently to create sub-basses

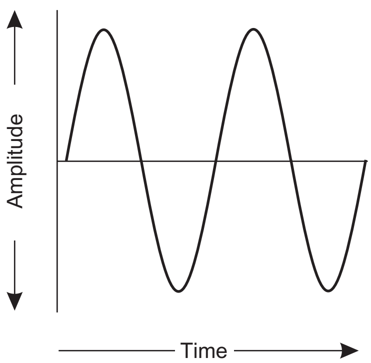

FIGURE 1.6
f0060 A sine wave

or whistling timbres or is mixed with other waveforms to add extra body or bottom end to a sound.

s0050
The Square Wave

p0390 A square wave is the simplest waveform for an electrical circuit to generate because it exists in only two states: high and low (Figure 1.7). This wave produces only odd harmonics resulting in a mellow, hollow sound. This makes it particularly suitable for emulating wind instruments, adding width to strings and pads, or for the creation of deep, wide bass sounds.

f0070
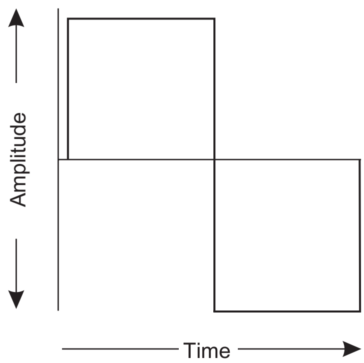

FIGURE 1.7
A square wave

s0060
The Pulse Wave

p0400 Although pulse waves are often confused with square waves, there is a significant difference between the two (Figure 1.8). Unlike a square wave, a pulse wave allows the width of the high and low states to be adjusted, thereby varying the harmonic content of the sound. Today it is unusual to see both square and pulse waves featured in a synthesizer. Rather the square wave offers an additional control allowing you to vary the width of the pulses. The benefit of this is that reductions in the width allow you to produce thin reed-like timbres along with the wide, hollow sounds created by a square wave.

f0080
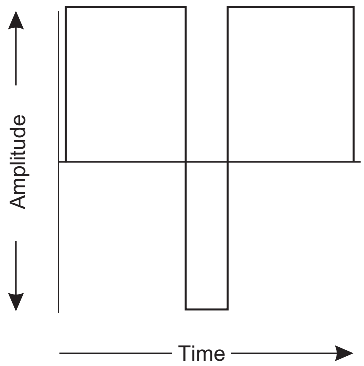

FIGURE 1.8
A pulse wave

s0070
The Sawtooth Wave

p0410 A sawtooth wave produces even and odd harmonics in series and therefore produces a bright sound that is an excellent starting point for brassy, raspy sounds (Figure 1.9). It's also suitable for creating the gritty, bright sounds needed for leads and raspy basses. Because of its harmonic richness, it is often employed in sounds that will be filter swept.

s0080
The Triangle Wave

p0420 The triangle wave shape features two linear slopes and is not as harmonically rich as a sawtooth wave since it only contains odd

f0090
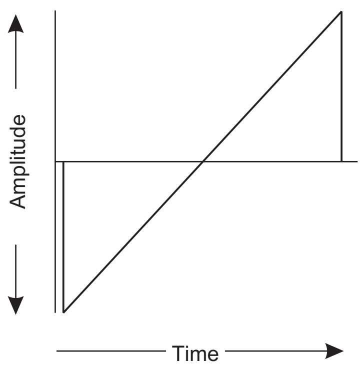

FIGURE 1.9
A sawtooth wave

SNOMAN 978-0-240-52107-7 00001
CH001.indd 11 8/25/2008 9:15:19 PM

harmonics (partials) (Figure 1.10). Ideally, this type of waveform is mixed with a sine, square or pulse wave to add a sparkling or bright effect to a sound and is often employed on pads to give them a glittery feel.

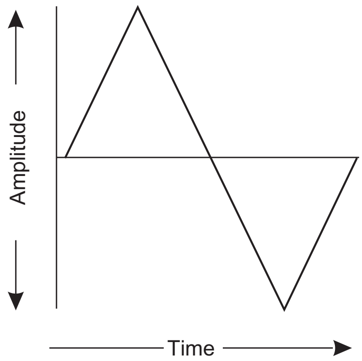

f0100 **FIGURE 1.10** A triangle wave

s0090
The Noise Wave

p0430 Noise waveforms are unlike the other five waveforms because they create a random mixture of all frequencies rather than actual tones (Figure 1.11). Noise waveforms can be 'pink' or 'white' depending on the energy of the mixed frequencies they contain. White noise contains equal amounts of energy at every frequency and is comparable to radio static, while pink noise contains equal amounts of energy in every musical octave and therefore we perceive it to produce a heavier, deeper hiss.

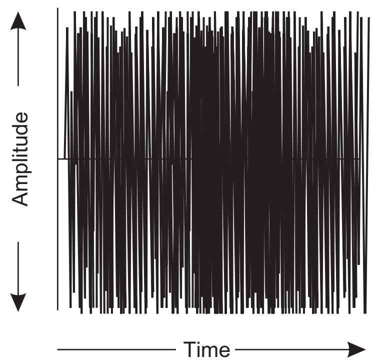

f0110 **FIGURE 1.11** A noise wave

p0440 Noise is useful for generating percussive sounds and was commonly used in early drum machines to create snares and handclaps. Although this remains its main use, it can also be used for simulating wind or sea effects, for producing breath effects in wind instrument timbres or for producing the typical trance leads.

s0100
CREATING MORE COMPLEX WAVEFORMS

p0450 Whether oscillators are created by analogue or DSP circuitry, listening to individual oscillators in isolation can be a mind-numbing experience. To create interesting sounds, a number of oscillators should be mixed together and used with the available modulation options.

p0460 This is achieved by first mixing different oscillator waveforms together and then detuning them all or just those that share the same waveforms so that they are out of phase from one another, resulting in a beating effect. Detuning is accomplished using the detune parameter on the synthesizer, usually by odd rather than even numbers. This is because detuning by an even number introduces further harmonic content that may mirror the harmonics already provided by the oscillators, causing the already present harmonics to be summed together.

p0470 It should be noted here that there is a limit to the level that oscillators can be detuned from one another. As previously discussed, oscillators should be detuned so that they beat, but if the speed of these beats is increased by any more than 20 Hz the oscillators separate, resulting in two noticeably different sounds. This can sometimes be used to good effect if the two oscillators are to be mixed with a timbre from another synthesizer because the additional timbre can help to fuse the two separate oscillators. As a general rule of thumb, it is unusual to detune an oscillator by more than an octave.

p0480 Additional frequencies can also be added into a signal using ring modulation and sync controls. Oscillator sync, usually found within the oscillator section of a synthesizer, allows a number of oscillators' cycles to be synced to one another. Usually all oscillators are synced to the first oscillator's cycle; hence, no matter where in the cycle any other oscillator is, when the first starts its cycle again the others are forced to begin again too.

p0490 For example, if two oscillators are used, with both set to a sawtooth wave and detuned by −5 cents (one-hundredth of a tone), every time the first oscillator restarts its cycle so too will the second, regardless of the position in its own cycle. This tends to produce a timbre with no harmonics and can be ideal for creating big, bold leads. Furthermore, if the first oscillator is unchanged and pitch bend is applied to the second to speed up or slow its cycle, screaming lead sounds typical of the Chemical Brothers are created as a consequence of the second oscillator fighting against the syncing with the first.

p0500 After the signals have left the oscillators, they enter the mixer section where the volume of each oscillator can be adjusted and features such as ring modulation can be applied to introduce further harmonics. (The ring modulation feature can sometimes be found within the oscillator section but is more commonly located in the mixer section, directly after the oscillators). Ring modulation works by providing a signal that is the sum and difference compound of two signals (while also removing the original tones). Essentially, this means that both signals from a two-oscillator synthesizer enter the ring modulator and come out from the other end as one combined signal with no evidence of the original timbre remaining.

p0510 As an example, if one oscillator produces a signal frequency of 440 Hz (A4 on a keyboard) and the second produces a frequency of 660 Hz (E5 on a keyboard), the frequency of the first oscillator is subtracted from the second.

$$660\,\text{Hz} - 440\,\text{Hz} = 220\,\text{Hz}(\text{A3})$$

Then the first oscillator's frequency is added to that of the second.

$$660\,\text{Hz} + 440\,\text{Hz} = 1100\,\text{Hz}(\text{C\#6})$$

Based on this example, the difference of 220 Hz provides the fundamental frequency while the sum of the two signals, 1100 Hz, results in a fifth harmonic overtone. When working with synthesizer, though, this calculation is rarely performed. This result is commonly achieved by ring modulating the oscillators together at any frequency and then tuning the oscillator. Ring modulation is typically used in the production of metallic-type effect (ring modulators were used to create the Dalek voice from *Dr Who*) and bell-like sounds. If ring modulation is used to create actual pitched sounds, a large number of in-harmonic overtones are introduced into the signal creating dissonant, unpitched results.

p0520 The option to add noise may also be included in the oscillator's mix section to introduce additional harmonics, making the signal leaving the oscillator/mix section full of frequencies that can then be shaped further using the options available.

s0110 VOLTAGE-CONTROLLED FILTERS

p0530 Following the oscillator's mixer section are the filters for sculpting the previously created signal. In the synthesizer world, if the oscillator's signal is thought of as a piece of wood that is yet to be carved, the filters are the hammer and chisels that are used to shape it. Filters are used to chip away pieces of the original signal until a rough image of the required sound remains.

p0540 This makes filters the most vital element of any subtractive synthesizer because if the available filters are of poor quality, few sound sculpting options will be available and it will be impossible to create the sound you require. Indeed, the choice of filters combined with the oscillator's waveforms is often the reason why specific synthesizers must be used to recreate certain 'classic' dance timbres.

p0550 The most common filter used in basic subtractive synthesizers is a low-pass filter. This is used to remove frequencies above a defined cut-off point. The effect is progressive, meaning that more frequencies are removed from a sound, the further the control is reduced, starting with the higher harmonics and gradually moving to the lowest. If this filter cut-off point is reduced far enough, all harmonics above the fundamental can be removed, leaving just the fundamental frequency. While it may appear senseless to create a bright sound with oscillators only to remove them later with a filter, there are several reasons why you may wish to do this.

u0070 ■ Using a variable filter on a bright sound allows you to determine the colour of the sound much more precisely than if you tried to create the same effect using oscillators alone.
u0080 ■ This method enables you to employ real-time movement of a sound.

This latter movement is an essential aspect of sound design because we naturally expect dynamic movement of sound throughout the length of the note. Using our previous example of a piano string being struck, the initial sound is very bright, becoming duller as it dies away. This effect can be simulated by opening the filter as the note starts and then gradually sweeping the cut-off frequency down to create the effect of the note dying away.

p0580 Notably, when using this effect, frequencies that lie above the cut-off point are not attenuated at right angles to the cut-off frequency; therefore, the rate at which they die away will depend on the transition period. This is why different f0120 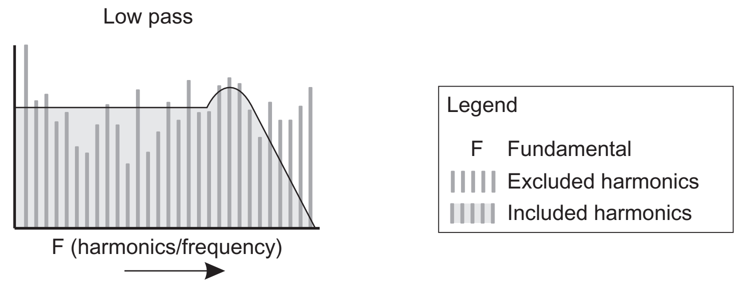

FIGURE 1.12
Action of the low-pass filter

filters that essentially perform the same function can make beautiful sweeps, whilst others can produce quite uneventful results (Figure 1.12).

p0590 When a cut-off point is designated, small quantities of the harmonics that lie above this point are not removed completely and are instead attenuated by a certain degree. The degree of attenuation is dependent on the transition band of the filter being used. The gradient of this transition is important because it defines the sound of any one particular filter. If the slope is steep, the filter is said to be 'sharp' and if the slope is more gradual the filter is said to be 'soft'. To fully understand the action of this transition, some prior knowledge of the electronics involved in analogue synthesizer is required.

p0600 When the first analogue synthesizers appeared in the 1960s, different voltages were used to control both the oscillators and the filters. Any harmonics produced by the oscillators could be removed gradually by physically manipulating the electrical current. This was achieved using a resistor (to reduce the voltage) and a capacitor (to store a voltage), a system that is often referred to as a resistor–capacitor (RC) circuit. Because a single RC circuit produces a 6 dB transition, the attenuation increases by 6 dB every time a frequency is doubled.

p0610 One RC element creates a 6 dB per octave 1-pole filter that is very similar to the gentle slope created by a mixing desks EQ. Consequently, manufacturers soon implemented additional RC elements into their designs to create 2-pole filters, which attenuated 12 dB per octave, and 4-pole filters, to provide 24 dB per octave attenuation. Because 4-pole filters attenuate 24 dB per octave, making substantial changes to the sound, they tend to sound more synthesized than sounds created by a 2-pole filter; so it's important to decide which transition period is best suited to the sound. For example, if a 24 dB filter is used to sweep a pad, it will result in strong attenuation throughout the sweep, while a 12 dB will create a more natural flowing movement (Figure 1.13).

p0620 If there is more than one of these available, some synthesizers allow them to be connected in series or parallel, which gives more control over the timbre from the oscillators. This means that two 12 dB filters could be summed together to produce a 24 dB transition, or one 24 dB filter could be used in isolation for aggressive tonal adjustments with the following 12 dB filter used to perform a real-time filter sweep.

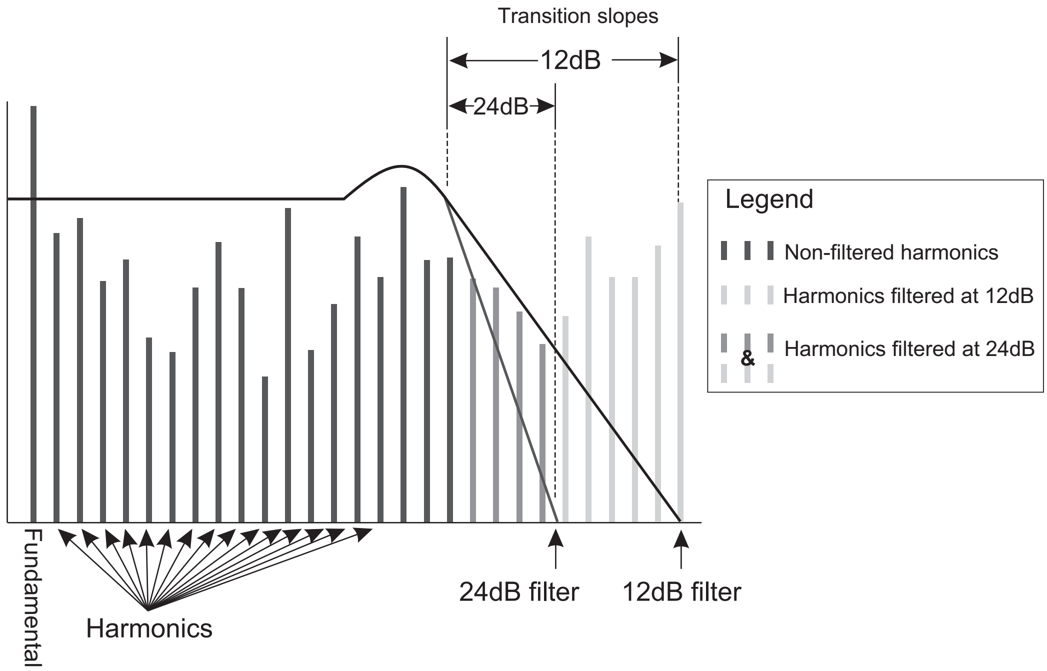

f0130 **Figure 1.13**
The difference between 12 dB and 24 dB slopes

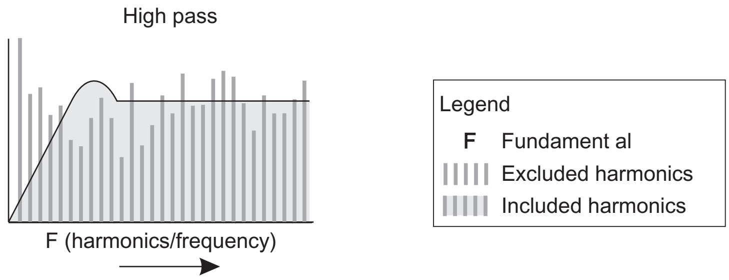

f0140 **FIGURE 1.14**
Action of a high-pass filter

p0630 Although low-pass filters are the most commonly used type, there are numerous variations including high pass, band pass, and notch and comb. These utilize the same transition periods as the low-pass filter but each has a widely different effect on the sound (Figure 1.14).

p0640 A high-pass filter has the opposite effect to a low-pass filter, first removing the low frequencies from the sound and gradually moving towards the highest. This is less useful than the low-pass filter because it effectively removes the fundamental frequency of the sound, leaving only the fizzy harmonic overtones. Because of this, high-pass filters are rarely used in the creation of instruments

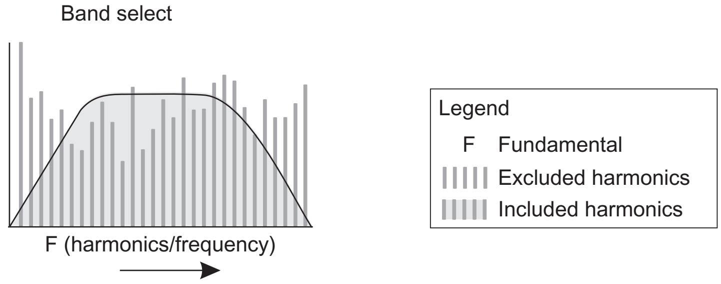

f0150 **FIGURE 1.15** Action of the band-pass filter

and are predominantly used to create effervescent sound effects or bright timbres that can be laid over the top of another low-pass sound to increase the harmonic content.

p0650 The typical euphoric trance leads are a good example of this, as they are often created from a tone with the fundamental overlaid with numerous other tones that have been created using a high-pass filter. This prevents the timbre from becoming too muddy as a consequence of stacking together fundamental frequencies. In both remixing and dance music, it's commonplace to run a high-pass filter over an entire mix to eliminate the lower frequencies, creating an effect similar to a transistor radio or a telephone. By reducing the cut-off control, gradually or immediately, the track morphs from a thin sound to a fatter one, which can produce a dramatic effect in the right context.

p0660 If high- and low-pass filters are connected in series, then it's possible to create a band-pass, or band-select, filter. These permit a set of frequencies to pass unaltered through the filter while the frequencies either side of the two filters are attenuated. The frequencies that pass through unaltered are known as the 'bandwidth' or the 'band pass' of the filter, and clearly, if the low pass is set to attenuate a range of frequencies that are above the current high-pass setting, no frequencies will pass through and no sound is produced.

p0670 Band-pass filters, like high-pass filters, are often used to create timbres consisting of fizzy harmonics (Figure 1.15). They can also be used to determine the frequency content of a waveform, as by sweeping through the frequencies each individual harmonic can be heard. Because this type of filter frequently removes the fundamental, it is often used as the basis of sound effects or lo-fi and trip-hop timbres or to create very thin sounds that will form the basis of sound effects.

p0680 Although band-pass filters can be used to thin a sound, they should not be confused with band-reject filters, which can be used for a similar purpose. Band-reject filters, often referred to as notch filters, attenuate a selected range of frequencies effectively creating a notch in the sound – hence the name – and usually leave the fundamental unaffected. This type of filter is handy for

SNOMAN 978-0-240-52107-7 00001
CH001.indd 17 8/25/2008 9:15:20 PM

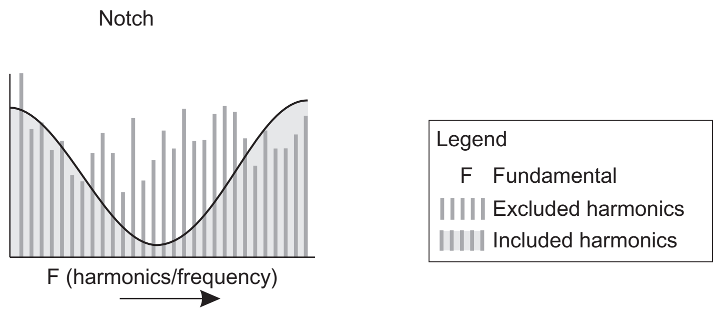

f0160 **FIGURE 1.16**
Action of the notch filter

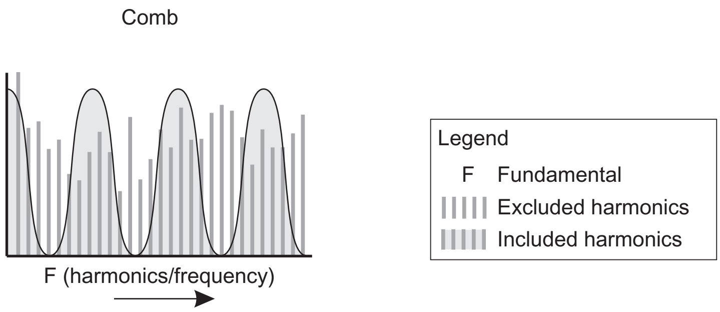

f0170 **FIGURE 1.17**
Action of the comb filter

scooping out frequencies, thinning out a sound while leaving the fundamental intact, making them useful for creating timbres that contain a discernable pitch but do not have a high level of harmonic content (Figure 1.16).

p0690 One final form of filter is the comb filter. With these, some of the samples entering the filter are delayed in time and the output is then fed back into the filter to be reprocessed to produce the results, effectively creating a comb appearance, hence the name. Using this method, sounds can be tuned to amplify or reduce specific harmonics based on the length of the delay and the sample rate, making it useful for creating complex sounding timbres that cannot be accomplished any other way. Because of the way they operate, however, it is rare to find these featured on a synthesizer and are usually available only as a third-party effect.

p0700 As an example, if a 1 kHz signal is put through the filter with a 1 ms delay, the signal will result in phase because 1 ms is coincident with the inputted signal, equalling one. However, if a 500 Hz signal with a 1 ms delay were used instead, it would be half of the period length and so would be shifted out of phase by 180°, resulting in a zero. It's this constructive and deconstructive period that creates the continual bump then dip in harmonics, resulting in a comb-like appearance when represented graphically, as in Figure 1.17. This method applies to all frequencies, with integer multiples of 1 kHz producing ones and odd multiples of 500 Hz (1.5, 2.5, 3.5 kHz etc.) producing zeros. The effect of

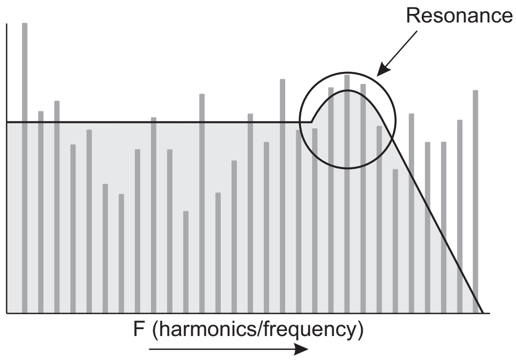

f0180 **FIGURE 1.18**
The effect of resonance

using this filter can at best be described as highly resonant, and forms the basis of flanger effects; therefore, its use is commonly limited to sound design rather than the more basic sound sculpting.

p0710 One final element of sound manipulation in a synthesizer's filter section is the resonance control. Also referred to as peak, this refers to the amount of the output of the filter that is fed back directly into the input, emphasizing any frequencies that are situated around the cut-off frequency. This has a similar effect to employing a band-pass filter at the cut-off point, effectively creating a peak. Although this also affects the filter's transition period, it is more noticeable at the actual cut-off frequency than anywhere else. Indeed, as you sweep through the cut-off range the resonance follows the curve, continually peaking at the cut-off point. In terms of the final sound, increasing the resonance makes the filter sound more dramatic and is particularly effective when used in conjunction with low-pass filter sweeps (Figure 1.18).

p0720 On many analogue and DSP-analogue-modelled synthesizers, if the resonance is turned up high enough it will feed back on itself. As more and more of the signal is fed back, the signal is exaggerated until the filter breaks into self-oscillation. This produces a sine wave with a frequency equal to that of the set cut-off point and is often a purer sine wave than that produced by the oscillators. Because of this, self-oscillating filters are commonly used to create deep, powerful sub-basses that are particularly suited to the drum 'n' bass and rap genres.

p0730 Notably, some filters may also feature a saturation parameter which essentially overdrives the filters. If applied heavily, this can be used to create distortion effects, but more often it's used to thicken out timbres and add even more harmonics and partials to the signal to create rich sounding leads or basses.

p0740 The keyboard's pitch can also be closely related to the action of the filters, using a method known as pitch tracking, keyboard scaling or more frequently 'key follow'. On many synthesizers the depth of this parameter is adjustable,

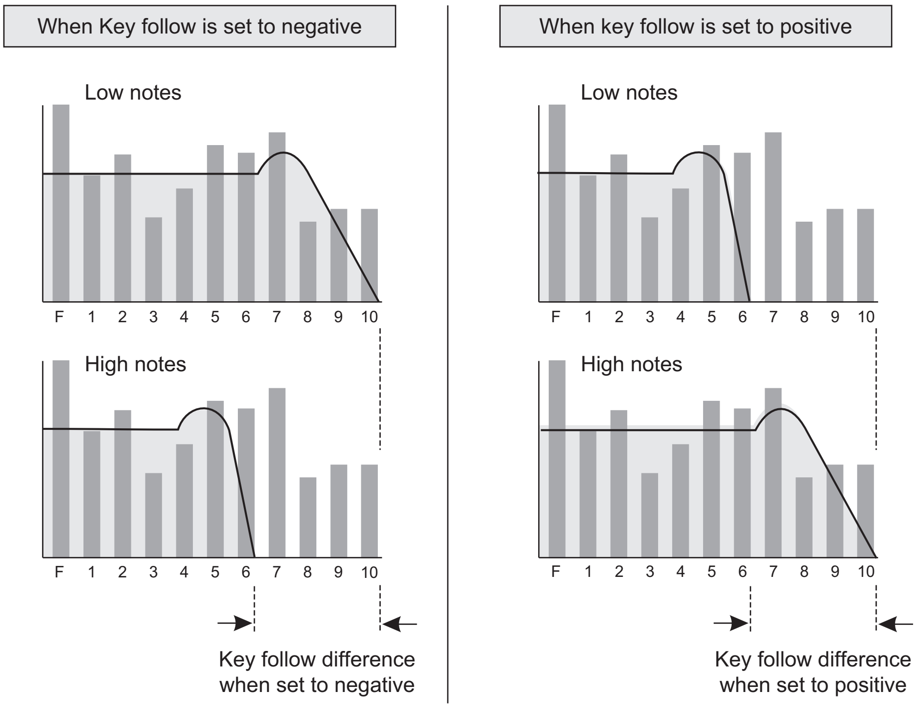

f0190 **FIGURE 1.19** The effect of filter key follow

allowing you to determine how much or how little the filter should follow the pitch.

p0750 When this parameter is set to its neutral state (neither negative nor positive), as a note is played on the keyboard the cut-off frequency tracks the pitch and each note is subjected to the same level of filtering. If this is used on a low-pass filter, for example, the filter setting remains fixed, so as progressively higher notes are played fewer and fewer harmonics will be present in the sound, making the timbre of the higher notes mellower than that of the lower notes. If the key follow parameter is set to positive, the higher notes will have a higher cut-off frequency and the high notes will remain bright (Figure 1.19). If, on the other hand, the key follow parameter is set to negative, the higher notes will lower the cut-off frequency, making the high notes even mellower than when key follow is set to its neutral state. Key follow is useful for recreating real instruments such as brass, where the higher notes are often mellower than the lower notes, and is also useful on complex bass lines that jump over an octave, adding further variation to a rhythm.

s0120 VOLTAGE-CONTROLLED AMPLIFIER (VCA)

p0760 Once the filters have sculpted a sound, the signal then moves into the final stage of synthesizer: the amplifier. When a key is pressed, rather than the volume f0200 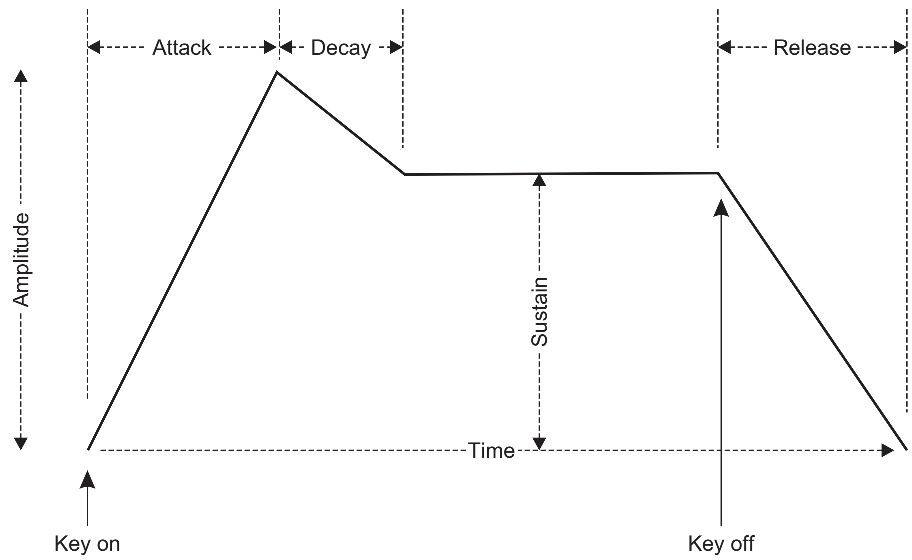

FIGURE 1.20
The ADSR envelope

rising immediately to its maximum and falling to zero when released, an 'envelope generator' is employed to emulate the nuances of real instruments.

p0770 Few, if any, acoustic instruments start and stop immediately. It takes a finite amount of time for the sound to reach its amplitude and then decay away to silence again; thus, the 'envelope generator' – a feature of all synthesizers – can be used to shape the volume with respect to time. This allows you to control whether a sound starts instantly the moment a key is pressed or builds up gradually and how the sound dies away (quickly or slowly) when the key is released. These controls usually comprise four sections called attack, decay, sustain, and release (ADSR), each of which determines the shaping that occurs at certain points during the length of a note. An example of this is shown in Figure 1.20.

u0090 ■ Attack: The attack control determines how the note starts from the point when the key is pressed and the period of time it takes for the sound to go from silence to full volume. If the period set is quite long, the sound will 'fade in', as if you are slowly turning up a volume knob. If the period set is short, the sound will start the instant a key is pressed. Most instruments utilize a very short attack time.

u0100 ■ Decay: Immediately after a note has begun it may initially decay in volume. For instance, a piano note starts with a very loud, percussive part but then drops quickly to a lower volume while the note sustains as the key is held down. The time the note takes to fade from the initial peak at the attack stage to the sustain level is known as the 'decay time'.

u0110 ■ Sustain: The sustain period occurs after the initial attack and decay periods and determines the volume of the note while the key is held down. This means that if the sustain level is set to maximum, any decay period

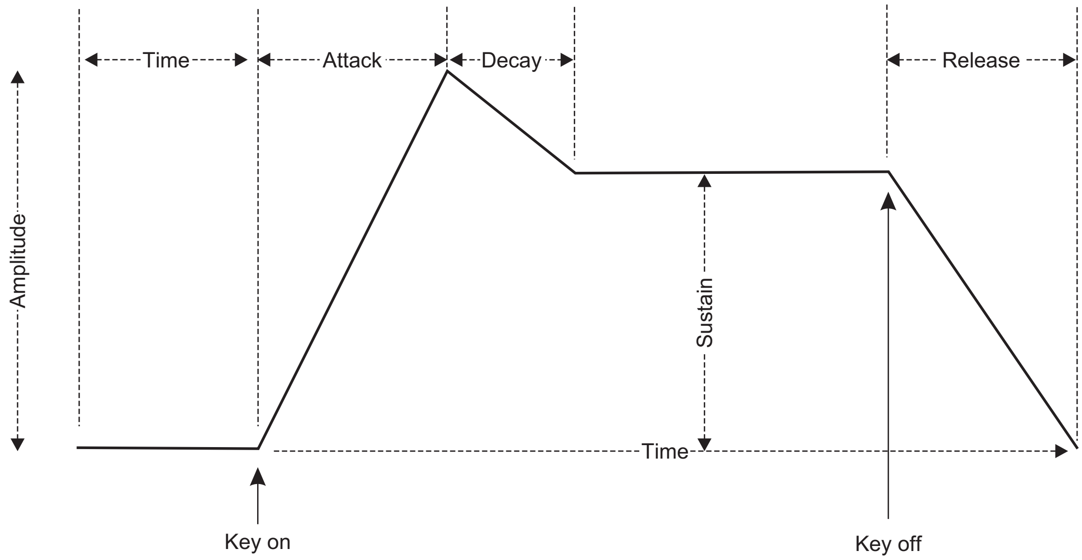

f0210 **FIGURE 1.21**
The TADSR envelope

will be ineffective, because at the attack stage the volume is at maximum and so there is no level to decay down to. Conversely, if the sustain level were set to zero, the sound peaks following the attack period and will fade to nothing even if you continue to hold down the key. In this instance, the decay time determines how quickly the sound decays down to silence.

u0120
- Release: The release period is the time it takes for the sound to fade from the sustain level to silence after the key has been released. If this is set to zero, the sound will stop the instant the key is released, while if a high value is set the note will continue to sound, fading away as the key is released.

p0820 AUQ4 Although ADSR envelopes are the most common, there are some subtle variations such as attack–release (AR), time–attack–delay–sustain–release (TADSR), and attack–delay–sustain–time–release (ADSTR). Because there are no decay or sustain elements contained in most drum timbres, AR envelopes are often used on drum synthesizers. They can also appear on more economical synthesizers simply because the AR parameters are regarded as having the most significant effect on a sound, making them a basic requirement. Both TADSR and ADSTR envelopes are usually found on more expensive synthesizers. With the additional period, T (time), in TADSR, for instance, it is possible to set the amount of time that passes before the attack stage is reached (Figure 1.21).

p0830 It's also important to note that not all envelopes offer linear transitions, meaning that the attack, decay and release stages will not necessarily consist

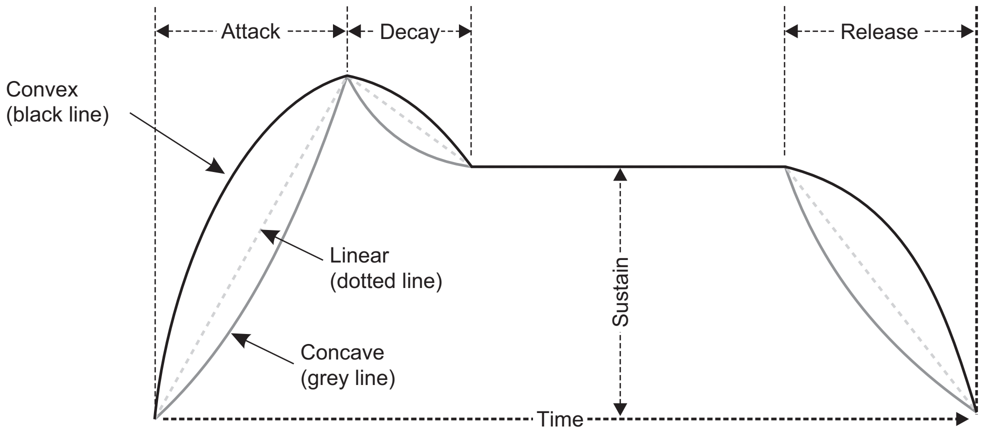

FIGURE 1.22
f0220 Linear and exponential envelopes

entirely of a straight line as it is shown in Figure 1.22. On some synthesizers these stages may be concave or convex, while other synthesizers may allow you to state whether the envelope stages should be linear, concave, or convex. The differences between the linear and the exponential envelopes are shown in Figure 1.22.

s0130 MODIFIERS

p0840 Most synthesizers also offer additional tools for manipulating sound in the form of modulation sources and destinations. Using these tools, the response or movement of one parameter can be used to modify another totally independent parameter, hence the name 'modifiers'.

p0850 The number of modifiers available, along with the destinations they can affect, is entirely dependent on the synthesizer. Many synthesizers feature a number of envelope generators that allow the action of other parameters alongside the amplifier to be controlled.

p0860 For example, in many synthesizers, an envelope may be used to modify the filter's action and by doing so you can make tonal changes to the note while it plays. A typical example of this is the squelchy bass sound used in most dance music. By having a zero attack, short decay and zero sustain level on the envelope generator, a sound that starts with the filter wide open before quickly sweeping down to fully closed is produced. This movement is archetypal to most forms of dance music but does not necessarily have to be produced by envelopes. Instead, some synthesizers offer one-shot low-frequency oscillators (LFOs) which can be used in the envelope's place. For instance, by using a triangle waveform LFO to modulate the amp, there is a slow rise in volume before a slowdrop again.

s0140 LOW-FREQUENCY OSCILLATOR

p0870 LFOs produce output frequencies in much the same way as VCOs. The difference is that a VCO produces an audible frequency (within the 20 Hz–20 kHz range) while an LFO produces a signal with a relatively low frequency that is inaudible to the human ear (in the range 1–10 Hz).

p0880 The waveforms an LFO can utilize depend entirely upon the synthesizer in question, but they commonly employ sine, saw, triangle, square, and sample and hold waveforms. The sample and hold waveform is usually constructed with a randomly generated noise waveform that momentarily freezes every few samples before beginning again.

p0890 LFOs should not be underestimated because they can be used to modulate other parameters, known as 'destination', to introduce additional movement into a sound. For instance, if an LFO is set to a relatively high frequency, say 5 Hz, to modulate the pitch of a VCO, the pitch of the oscillator will rise and fall according to the speed and shape of the LFO waveform and an effect similar to that of vibrato is generated. If a sine wave is used for the LFO, then it will essentially create an effect similar to that of a wailing police siren. Alternatively, if this same LFO is used to modulate the filter cut-off, then the filter will open and close at a speed determined by the LFO, while if it were used to modulate an oscillator's volume, it would rise and fall in volume recreating a tremolo effect.

p0900 This means that an LFO must have an amount control (sometimes known as depth) for varying how much the LFO's waveform augments the destination, a rate control to control the speed of the LFO's waveform cycles, and a fade-in control in some. The fade-in control adjusts how quickly the LFO begins to affect the waveform after a key has been pressed. An example of this is shown in Figure 1.23.

p0910 The LFO on more capable synthesizers may also have access to its own envelope. This gives control of the LFO's performance over a specified time period, allowing it not only to fade in after a key has been pressed but also to decay, sustain, and fade away gradually. It is worth noting, however, that the destinations an LFO can modulate are entirely dependent on the synthesizer being used. Some synthesizers may only allow LFOs to modulate the oscillator's pitch and the filter, while others may offer multiple destinations and more LFOs. Obviously, the more LFOs and destinations that are available, the more creative options you will have at your disposal.

p0920 If required, further modulation can be applied with an attached controller keyboard or the synthesizer itself in the form of two modulation wheels. The first, pitch bend, is hard-wired and provides a convenient method of applying a modulating CV to the oscillator(s). By pushing the wheel away from you, you can bend the pitch (i.e. frequency) of the oscillator up. Similarly, you can bend the pitch down by pulling the wheel towards you. This wheel is normally spring loaded to return to the centre position, where no bend is applied,

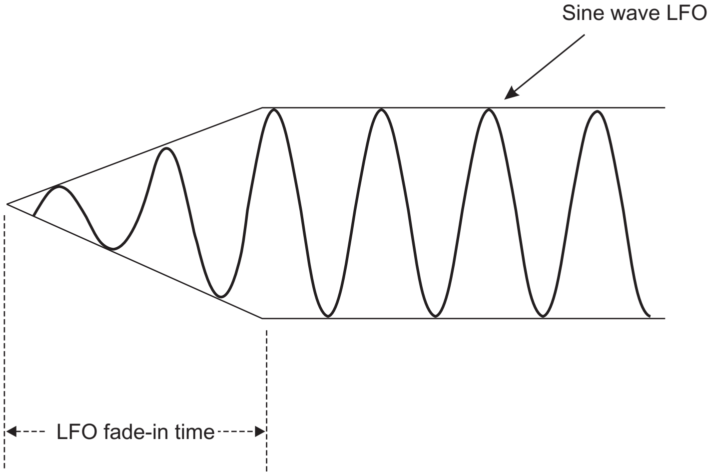

FIGURE 1.23
f0230 LFO fade-in

if you let go of it, and is commonly used in synthesizer solos to give additional expression. The second wheel, modulation, is freely assignable and offers a convenient method of controlling any on-board parameters, such as the level of the LFO signal sent to the oscillator, filter or VCA or to control the filter cut-off directly. Again, whether this wheel is assignable will depend on the manufacturer of the synthesizer.

p0930 On some synthesizers the wheels are hard coded to only allow oscillator modulation (for a vibrato effect), while some others do not have a separate modulation wheel and instead the pitch bend lever can be pushed forward to produce LFO modulation.

s0150 PRACTICAL APPLICATIONS

p0940 While there are other forms of synthesis – which will be discussed later in this chapter – most synthesizers used in the production of dance music are of an analogue/subtractive nature; therefore, it is vital that the user grasps the concepts behind all the elements of subtractive synthesis and how they can work together to produce a final timbre. With this in mind, it is sensible to experiment with a short example to aid in the understanding of the components.

p0950 Using the synthesizer of your choice, clear all the current settings so that you start from nothing. On many synthesizers, this is known as 'initializing a patch', so it may be a button labelled 'init', 'init patch' or similar.

p0960 Begin by pressing and holding C3 on your synthesizer, or alternatively con-
trolling the synthesizer via MIDI programme in a continual note. If not, place AUQ5
something heavy on C3. The whole purpose of this exercise is to hear how the

sound develops as you begin to modify the controls of the synthesizer, so the note needs to play continually.

p0970 Select sawtooth waves for two oscillators; if there is a third oscillator that you cannot turn off, choose a triangle for this third oscillator. Next, detune one sawtooth from the other until the timbre begins to thicken. This is a tutorial to grasp the concept of synthesis, so keep detuning until you hear the oscillators separate from one another and then move back until they become one again and the timbre is thickened out. Generally speaking, detuning of 3 cents should be ample but do not be afraid to experiment – this is a learning process. If you are using a triangle wave, detune this against the two saws and listen to the results. Once you have a timbre you feel you can work with, move onto the next step.

p0980 Find the VCA envelope and start experimenting. You will need to release C3 and then press it again so you can hear the effect that the envelope is having on the timbre. Experiment with these envelopes until you have a good grasp on how they can adjust the shape of a timbre; once you're happy you have an understanding, apply a fast attack with a short decay, medium sustain and a long release. As before, for this next step you will need to keep C3 depressed.

p0990 Find the filter section, and experiment with the filter settings. Start by using a high-pass filter with the resonance set around midway and slowly turn the filter cut-off control. Note how the filter sweeps through the sound, removing the lower frequencies first, slowly progressing to the higher frequencies. Also experiment with the resonance by rotating it to move upwards and downwards and note how this affects the timbre. Do the same with the notch and band pass etc. (if the synthesizer has these available) before finally moving to the low pass. Set the low-pass filter quite low, along with a low-resonance setting – you should now have a static buzzing timbre.

p1000 The timbre is quite monotonous, so use the filter envelope to inject some life into the sound. This envelope works on exactly the same principles as the VCA, with the exception that it will control the filter's movement. Set the filter's envelope to a long attack and decay, but use a short release and no sustain and set the filter envelope to maximum positive modulation. If the synthesizer has a filter key follow, use this as it will track the pitch of the note being played and adjust itself. Now try depressing C3 to hear how the filter envelope controls the filter, essentially sweeping through the frequencies as the note plays.

p1010 Finally, to add some more excitement to the timbre, find the LFO section. Generally, the LFO will have a rotary control to adjust the rate (speed), a selector switch to choose the LFO waveform, a depth control and a modulation destination. Choose a triangle wave for the LFO waveform, Hold down C3 on the synthesizer's keyboard, turn the LFO depth control up to maximum and set the LFO destination to pitch. As before, hold down the C3 key and slowly rotate the LFO rate (speed) to hear the results. If you have access to a second LFO, try modulating the filter cut-off with a square wave LFO, set the LFO depth to maximum and experiment with the LFO rate again.

p1020 If you would like to experiment more with synthesis to help get to grips with the principles, jump to Chapter 4 for further information on programming specific synthesizer timbres. Note, however, that different synthesizers will produce timbres differently and some are more suited to reproducing particular timbres than others.

s0160 OTHER SYNTHESIS METHODS

s0170 Frequency Modulation (FM)

p1030 FM is a form of synthesizer developed in the early 1970s by Dr John Chowning of Stanford University, then later developed further by Yamaha, leading to the release of the now-legendary DX7 synthesizer: a popular source of bass sounds for numerous dance musicians.

p1040 Unlike analogue, FM synthesizer produces sound by using operators, which are very similar to oscillators in an analogue synthesizer but can only produce simple sine waves. Sounds are generated by using the output of the first operator to modulate the pitch of the second, thereby introducing harmonics. Like an analogue synthesizer, each FM voice requires a minimum of two oscillators in order to create a basic sound, but because FM only produces sine waves the timbre produced from just one carrier and modulator isn't very rich in harmonics.

p1050 In order to remedy this, FM synthesizers provide many operators that can be configured and connected in any number of ways. Many will not produce musical results, so to simplify matters various algorithms are used. These algorithms are preset as combinations of modulator and carrier routings. For example, one algorithm may consist of a modulator modulating a carrier, which in turn modulates another carrier, before modulating a modulator that modulates a carrier to produce the overall timbre. The resulting sound can then be shaped and modulated further using LFOs, filters and envelopes using the same subtractive methods as in any analogue synthesizer.

p1060 This means that it should also be possible to emulate FM synthesizer in an analogue synthesizer with two oscillators, where the first oscillator acts as a modulator and the second acts as a carrier. When the keyboard is played, both oscillators produce their respective waveforms with the frequency dictated by the particular notes that were pressed. If the first oscillator's output is routed into the modulation input of the second oscillator and further notes are played on the keyboard, both oscillators play their respective notes but the pitch of the second oscillator will change over time with the frequency of the first, essentially creating a basic FM synthesizer. Although this is, in effect, FM, it is usually called 'cross modulation' in analogue synthesizers.

p1070 Due to the nature of FM, many of the timbres created are quite metallic and digital in character, particularly when compared to the warmth generated by the drifting of analogue oscillators. Also due to the digital nature of FM synthesizer, the facia generally contains few real-time controllers. Instead, numerous buttons adorn the front panel forcing you to navigate and adjust any parameters through a small LCD display.

p1080 Notably, although both FM and analogue synthesizers were originally used to reproduce realistic instruments, neither can fabricate truly realistic timbres. If the goal of the synthesizer system is to recreate the sound of an existing instrument, this can generally be accomplished more accurately using digital sample-based techniques.

s0180 SAMPLES AND SYNTHESIS

p1090 Unlike analogue or FM, sample synthesizer utilizes samples in place of the oscillators. These samples, rather than consisting of whole instrument sounds, also contain samples of the various stages of a real instrument along with the sounds produced by normal oscillators. For instance, a typical sample-based synthesizer may contain five different samples of the attack stage of a piano, along with a sample of the decay, sustain and release portions of the sound. This means that it is possible to mix the attack of one sound with the release of another to produce a complex timbre.

p1100 Commonly, up to four of these individual 'tones' can be mixed together to produce a timbre and each of these individual tones can have access to numerous modifiers including LFOs, filters and envelopes. This obviously opens up a whole host of possibilities not only for emulating real instruments, but also for creating complex sounds. This method of synthesizer has become the de
AUQ6 facto standard for any synthesizer producing realistic instruments. By combining both samples of real-world sounds with all the editing features and functionality of analogue synthesizers, they can offer a huge scope for creating both realistic and synthesized sounds.

s0190 GRANULAR SYNTHESIS

p1110 One final form of synthesizer that has started to make an appearance with the evolution of technology is granular synthesizer. It is rare to see a granular synthesizer employed in hardware synthesizers due to its complexity, but software synthesizers are being developed for the public market that utilize it. Essentially, it works by building up sounds from a series of short segments of sounds called 'grains'. This is best compared to the way that a film projector operates, where a series of still images, each slightly different from the last, are played sequentially at a rate of around 25 pictures per second, fooling the eyes and brain into believing there is a smooth continual movement.

p1120 A granular synthesizer operates in the same manner with tiny fragments of sound rather than still images. By joining a number of these grains together, an overall tone is produced that develops over a period of time. To do this, each grain must be less than 30 ms in length as, generally speaking, the human ear is unable to determine a single sound if they are less than 30–50 ms apart. This also means that a certain amount of control has to be offered over each grain. In any one sound there can be anything from 200 to 1000 grains, which is the main reason why this form of synthesizer appears mostly in the form of software. Typically, a granular synthesizer will offer most, but not necessarily all, of the following five parameters:

u0130 ■ Grain length: This can be used to alter the length of each individual grain. As previously mentioned, the human ear can differentiate between two grains if they are more than 30–50 ms apart, but many granular synthesizers usually go above this range, covering 20–100 ms. By setting this length to a higher value, it's possible to create a pulsing effect.

u0140 ■ Density: This is the percentage of grains that are created by the synthesizer. Generally, it can be said that the more the grains created, the more complex a sound will be, a factor that is also dependent on the grain shape.

u0150 ■ Grain shape: Commonly, this offers a number between 0 and 200 and represents the curve of the envelopes. Grains are normally enveloped so that they start and finish at zero amplitude, helping the individual grains mix together coherently to produce the overall sound. By setting a longer envelope (a higher number) two individual grains will mix together, which can create too many harmonics and often result in the sound exhibiting lots of clicks as it fades from one grain to the other.

u0160 ■ Grain pan: This is used to specify the location within the stereo image where each grain is created. This is particularly useful for creating timbres that inhabit both speakers.

u0170 ■ Spacing: This is used to alter the period of time between each grain. If the time is set to a negative value, the preceding grain will continue through the next created grain. This means that setting a positive value inserts space between each grain; however, if this space is less than 30 ms, the gap will be inaudible.

p1180 The sound produced with granular synthesizer depends on the synthesizer in question. Usually, the grains consist of single frequencies with specific waveforms or occasionally they are formed from segments of samples or noise that have been filtered with a band-pass filter. Thus, the constant change of grains can produce sounds that are both bright and incredibly complex, resulting in a timbre that's best described as glistening. After creating this sound by combing the grains, the whole sound can be shaped by using envelopes, filters and LFOs.

Author Queries

{AUQ1}	Figures 1.4 and 1.5 are set as a single figure, and consequently the following figures have been renumbered. Please check.
{AUQ2}	Please check and confirm the edits for correctness in the sentence 'Though these diagrams are very simple…'
{AUQ3}	Please check and confirm the edits in the following sentence for correctness: 'When a key on a keyboard is pressed...'
{AUQ4}	The four sections of ADSR (attack, decay, sustain, release) have been turned into a list. Please check.
{AUQ5}	Please check and confirm the edits in the following sentence for correctness: 'Begin by pressing and holding ...'
{AUQ6}	Please check the following sentence for clarity and rephrase if necessary: 'This method of synthesizer...'