Digital Sound Modelling
lecture notes for Com Sci 295

Michael J. O'Donnell

10 March 1994—revised 24 March 1996

Contents

List of Figures

For supplementary material, particularly better versions of graphs, see the associated Maple files.

Chapter 1

Physical and Mathematical Foundations of Sound Modelling

In order to have useful discussions about sound, we need a very simplistic, but practical, understanding of the physics and mathematics associated with sound.

1.1 Physics: What Is Sound?

For our purposes, sound is any kind of vibration that is detectible by the ear or devices analogous to the ear. Treatments of sound in physics books tend to focus attention on the transmission of sound vibrations through the air. We will focus instead on the vibrating systems that produce and detect sounds, and just assume that the air is capable of transmitting vibrations from sound producers to the detectors in the ear.

1.1.1 Vibrating Springs

The simplest sort of vibration to understand is that of a spring. To really simplify things, imagine an environment with no gravity, and with a mass (a solid chunk of something) moving along a frictionless track that is fixed so the track cannot move. The track constrains motion of the mass to a straight line, so we do not need to consider the three dimensions of space.

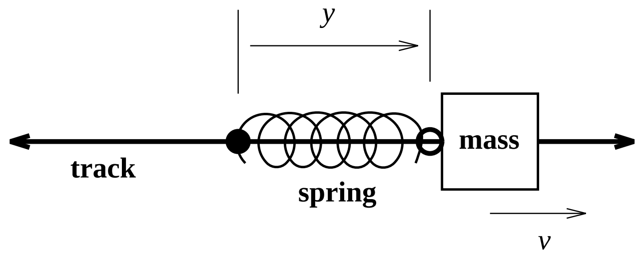

Figure 1.1: Ideal spring resonator

Finally, imagine a spring attached at one end to the mass, and at the other end to some fixed point on the track. Be a bit liberal-minded, and imagine that the spring has length 0 when it is not stretched, and that the mass can move freely past the point where the spring is attached. A picture of our imaginary system is given in Figure 1.1.

At any moment in time, the state of the spring system can be described by two real numbers: the *displacement* y of the mass to the right of the point on the track to which the spring is fixed, and the *velocity* v of the mass to the right. Displacement to the left is represented by negative values of y, and motion to the left is represented by negative values of v. Now, imagine that we displace the mass to the right and hold it in a fixed position, stretching the spring. That is, we establish an initial condition where $y > 0$ and $v = 0$. When we release the mass, the spring pulls it to the left, causing a state where $y > 0$ and $v < 0$. Eventually the mass reaches the center of the track at $y = 0$, but at this moment $v < 0$ and inertia carries the mass beyond the center, to the left where $y < 0$. Now, the spring pulls the mass to the right, cancelling out the motion $v < 0$ to the left. Eventually the mass stops with $v = 0$, but at this moment the spring is stretched to the left with $y < 0$, so the pull to the right continues and causes the mass to move right with $v > 0$. This motion to the right eventually moves the mass past the center, so $y > 0$. The leftward pull of the spring opposes the motion until $v = 0$. So we return to a condition that is similar to the initial one: $y > 0$ and $v = 0$, and the cycle repeats. Figure 1.2 shows a schematic qualitative view of the vibration of the spring.

To complete the simplistic physics of a vibrating spring, we need to convert the qualitative observations above into quantitative information that we can use in a mathematical analysis. For this purpose, let t be a real number representing the *time* that has passed since some arbitrary starting moment

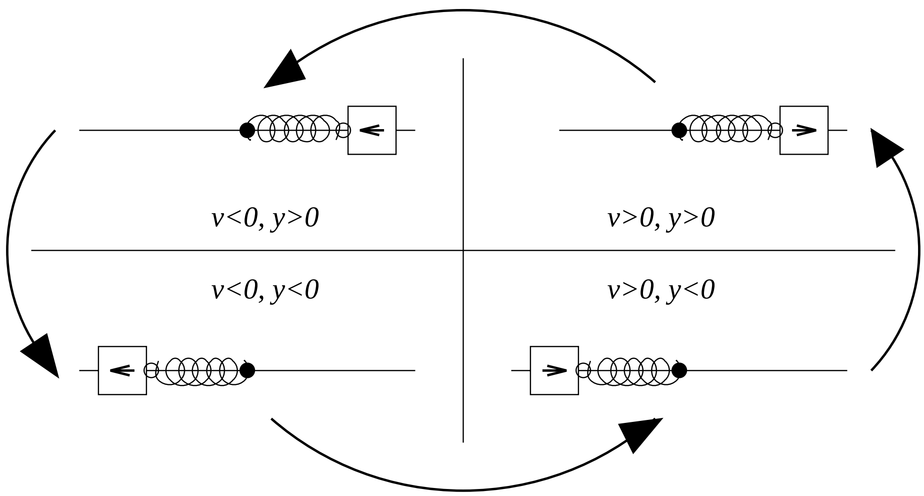

Figure 1.2: Qualitative states of spring

when $t = 0$. For any quantity q that depends on time, dq/dt means the instantaneous rate of change of q with respect to t—when the independent variable t is understood from context, dq/dt is often abbreviated q'. When no outside force acts on the spring and mass, its behavior is described by the following two equations:

$$y' = Av \quad (1.1)$$
$$v' = -By \quad (1.2)$$

A and B are positive real number constants (independent of time)—their actual values do not matter to us. Equation 1.1 holds because velocity is defined to be the change of location over time, and displacement is just a measure of location from a particular origin—the constant A takes care of any conversion of units between y' and v (normally the units are the same and $A = 1$). Equation 1.2 represents the fact that the force exerted by a spring increases in magnitude proportionally to the distance that the spring is stretched, and the force acts to pull the ends of the spring together. The value of B is determined by the stiffness of the spring. Equation 1.2 is an approximation, because no real spring exerts a force precisely proportional to the stretching distance—in particular when a spring is stretched too far it changes radically, becoming stiffer, or becoming softer, or breaking, depending on its construction. The right practical approach to understanding vibration is to do as much analysis as possible based on the simple approximate equations above, and then do the potentially complicated corrections only when greater accuracy is required.

Vibrating objects that produce sound, and others (such as the hairs in the cochlea of the ear) that detect sound, can be modelled fairly well by systems of vibrating springs connected together in various ways. Other vibrating systems have other physical parameters that measure the vibrating behavior, but in most cases there are two real numbers—for example *pressure* and *flow* of vibrating air, *potential* and *current* of vibrating electrical charge—that behave analogously to displacement and velocity in a vibrating spring.

1.2 Mathematics: How Do We Model Sound?

The key to understanding the mathematical analysis of sound is to *visualize* the mathematics using graphs and geometric diagrams. The right way to visualize the mathematics does *not* look like the physical system of vibrating springs or other objects that it is describing. The value of the mathematics is to give us a *different* way of visualizing sound, that is much more convenient for analytic reasoning than the actual physical configuration of vibrating objects. Mathematically, the important properties of a vibrating spring are just Equations 1.1 and 1.2. We can forget that they arose from the physical properties of a spring, and just consider the numerical behavior of two real numbers x and y as functions of t, when they satisfy the equations.

$$y' = Ax \tag{1.3}$$

$$x' = -By \tag{1.4}$$

From now on, lower case Roman variables, such as x and y, stand for real numbers that are functions of a time parameter t. Upper case Roman variables, such as A and B, stand for real number constants, which are the same as unvarying functions of time (but widely used notations, such as $e = 2.71828\ldots$, are left alone). Occasionally, we will use the form $x(t)$ to denote the value of a function x at a particular time t, but usually we will refer to entire functions rather than individual values. When an expression $\alpha(t)$ containing an independent variable, such as t, should refer to an entire function, rather than a single value of the function, we write $[t]\alpha(t)$.

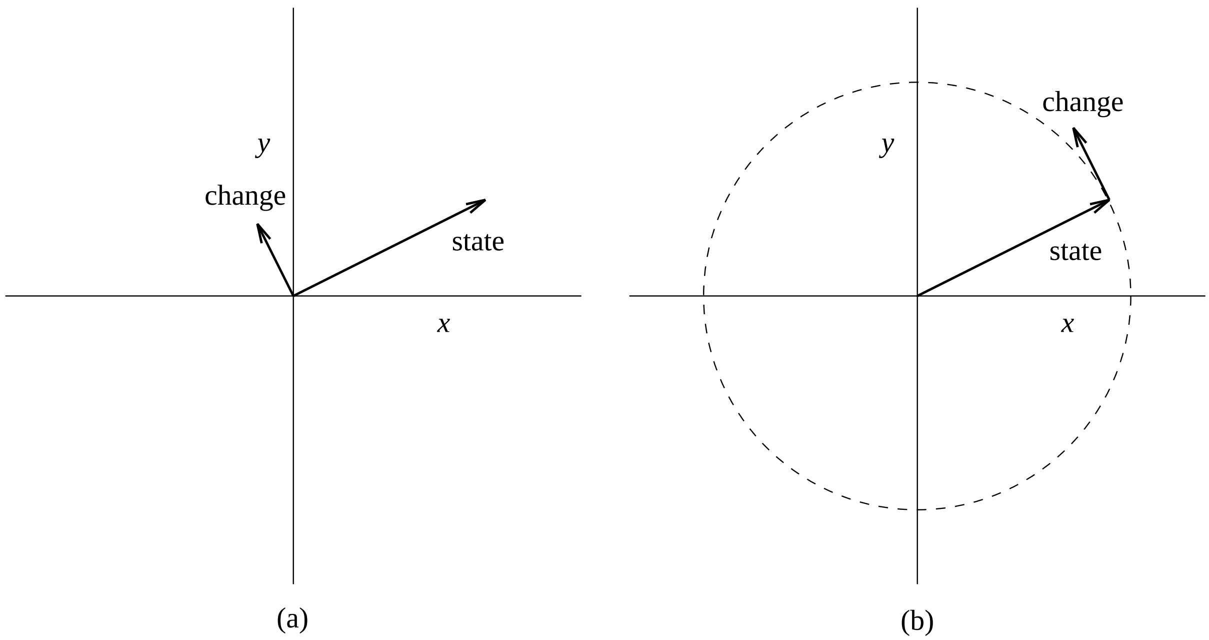

Figure 1.3: State and change vectors for vibrating spring

1.2.1 Vibration as the Circular Movement of a Rotor

To visualize all the possible states of a vibrating system, consider a plane in which the horizontal axis gives the value of x and the vertical axis gives the value of y—in this way each possible state of the system is a point in the plane.

First, consider the simple case where $B = A$, so Equations 1.3 and 1.4 specialize to

$$y' = Ax \tag{1.5}$$
$$x' = -Ay \tag{1.6}$$

Figure 1.3(a) shows an example point $\langle x, y\rangle$ and the corresponding point $\langle x', y'\rangle = \langle -Ay, Ax\rangle$ as vectors in the plane, when $A = 3/8$. Notice that the angle between these two vectors is always a right angle. Since $\langle x', y'\rangle$ represents a *change* in $\langle x, y\rangle$ it is useful to displace the origin of the vector representing $\langle x', y'\rangle$ to the end of the vector representing $\langle x, y\rangle$, as shown in Figure 1.3(b). Now, it is easy to see that the state of the system must trace out a circle in the plane centered about the origin, because the direction of

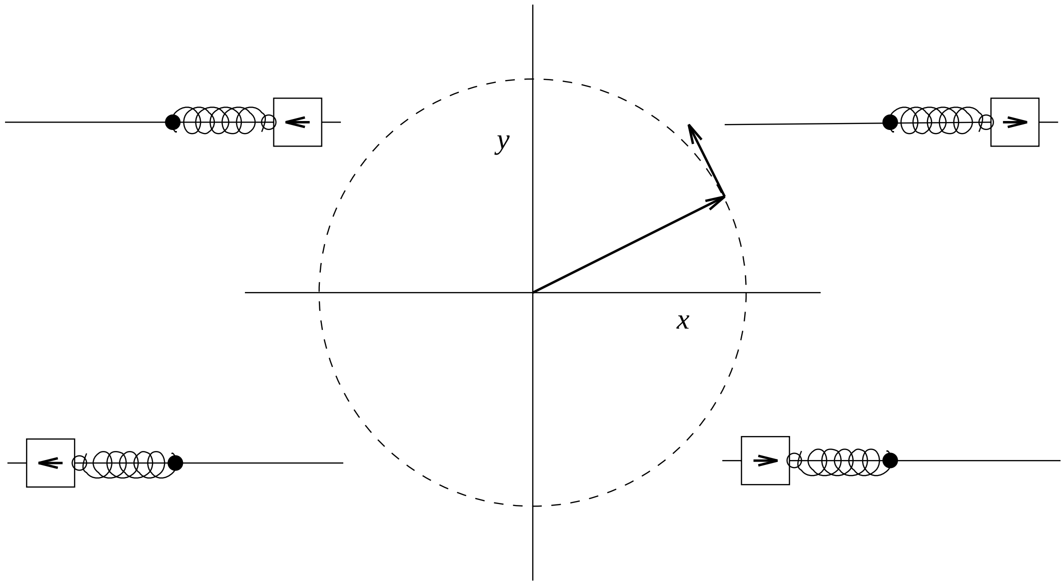

Figure 1.4: Spring system vs. rotor

change is always at right angles to the state vector. The size of the circle can be any nonnegative real number—setting the size corresponds to providing an initial displacement to the mass on the spring. Furthermore, since the magnitude of the state vector always stays the same, the magnitude of the change vector is also always the same (A times the magnitude of the state vector), so the state moves around the circle at a constant speed. I call such a system with a point moving around a circle at a constant speed a *rotor.* The time required for one full rotation is the *period* P of the rotor. The number of full rotations in a unit of time is the *frequency* of the rotor: its value is $1/P$. The magnitude of the state vector is the *amplitude* of the rotor. The angle of the state vector with respect to the x axis ($\langle 1, 0\rangle$) at time $t = 0$ is the *phase* of the rotor.

Take 5 minutes to visualize the relationship between the rotor and the vibrating spring system, as suggested in Figure 1.4. Notice that we have no interest in actual physical devices that look like rotors—the rotor is purely a mathematical concept that allows us to analyze the behavior of a vibrating system. Now forget about springs, and always visualize vibration in terms of rotors and similar mathematical systems that we investigate later.

While the speed of a rotor state around its circular path is constant, the

x and y components of the rotor state oscillate sinusoidally. Consider a rotor with amplitude R (that is, the circular path has radius R) and frequency F (F full rotations per unit time), starting at time $t = 0$ in state $\langle x, y\rangle = \langle R, 0\rangle$. The values of x and y at any time are given by the trignometric cos and sin functions.

$$\begin{aligned} x &= R\cos(2\pi Ft) && (1.7)\\ y &= R\sin(2\pi Ft) && (1.8) \end{aligned}$$

The multiplication by 2π is required because we measure angles in *radians,* and one full rotation is 2π radians. Notice that the maximum (minimum) values for x and y are both R ($-R$), and each reaches its maximum and minimum when the other is 0. Figure 1.5 shows x (solid line) and y (dashed line) as functions of time t for a rotor with a frequency of 1/4 rotation per unit time. Figure 1.6 shows a three-dimensional plot of x, y, and t. The path of the state is a helix, circling about the t axis. Think of the helix as the trace of a point running around the circle from Figure 1.3.

When $A \neq B$ in Equations 1.3 and 1.4, the state vector traces out an ellipse, whose aspect ratio is $\sqrt{A/B}$. The speed of the state vector around the ellipse is not constant (but the period and frequency are still well defined). Instead of figuring out a detailed description of an elliptical rotor, notice that we can always normalize a rotor to have circular motion, by changing the units in which x and y are measured. In an elliptical rotor with frequency F, starting at time $t = 0$ in state $\langle x, y\rangle = \langle R_x, 0\rangle$ and crossing the y axis in state $\langle x, y\rangle = \langle 0, R_y\rangle$, the values of x and y at any time are still given by the cos and sin functions, but with different scaling factors for each.

$$\begin{aligned} x &= R_x\cos(2\pi Ft) && (1.9)\\ y &= R_y\sin(2\pi Ft) && (1.10) \end{aligned}$$

In this case, the maximum (minimum) value for x is R_x ($-R_x$), and for y it is R_y ($-R_y$). As before, each parameter reaches its maximum and minimum when the other is 0.

1.2.2 Rotor State as a Complex Number

It is mathematically convenient to think of the two-dimensional rotor state vector $\langle x, y\rangle$ as a single complex number $x + \boldsymbol{i}y$, where $\boldsymbol{i}$ is the "imaginary"

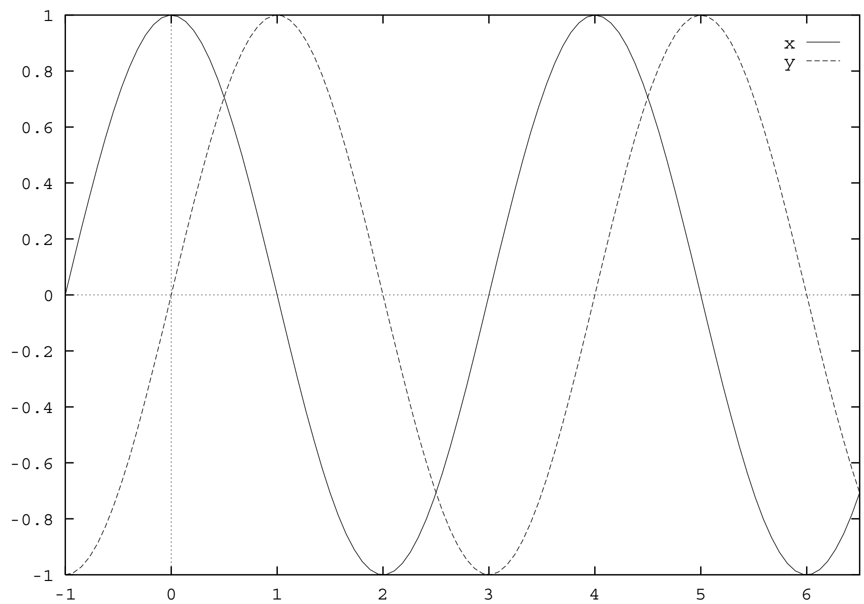

$$x = \cos(2\pi t/4) \text{ solid line}$$

$$y = \sin(2\pi t/4) \text{ dashed line}$$

Figure 1.5: Rotor state parameters x and y as functions of time t

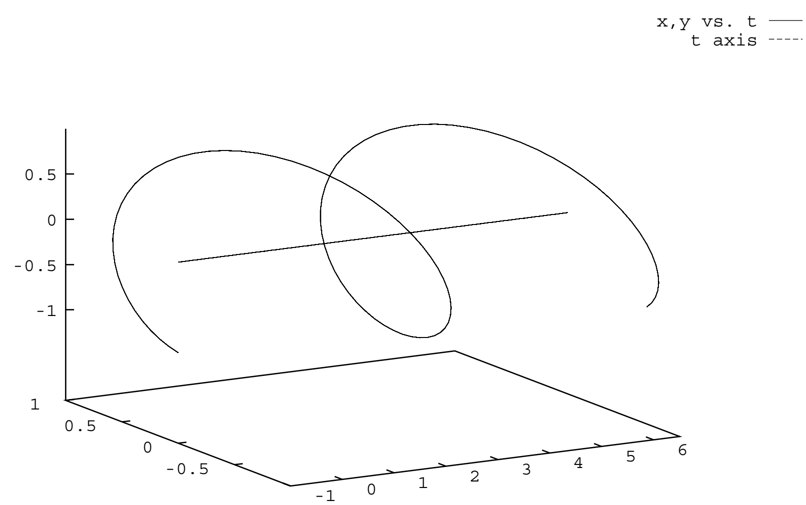

t	breadth axis
$x = \cos(2\pi t/4)$	depth axis
$y = \sin(2\pi t/4)$	vertical axis

Figure 1.6: Rotor parameters x and y vs. t as a helix in three dimensions

number defined to be the principal square root of -1 (if you read engineering books and articles, you may see this number written as $\boldsymbol{j}$ instead of $\boldsymbol{i}$). Do *not* look for deep significance in the names "real number," "imaginary number," "complex number." These names are just tags made up by mathematicians—"real" numbers are no more *real* than other numbers, "imaginary" numbers are no more *imaginary,* and "complex" numbers are used to *simplify* a lot of the analysis that we need to do. For our purposes, the complex number $x + \boldsymbol{i}y$ is just a particular notation for the vector $\langle x, y \rangle$, which is particularly convenient because the familiar operations of addition, multiplication, and exponentiation on the real numbers extend very naturally to operations on complex numbers that are just right for analyzing vibration.

Review of Complex Arithmetic

From now on, we use Greek letters α, β, γ, etc. as variables ranging over complex number functions depending on the time variable t. Complex number constants independent of time are denoted by bold face Greek letters $\boldsymbol{\alpha}$, $\boldsymbol{\beta}$, $\boldsymbol{\gamma}$, etc. (but widely used notations, such as $\pi = 3.14159\ldots$ are left alone). It is important to be fluent in the following facts about complex numbers, and to be able to do complex arithmetic and algebra just as easily as you learned to do real arithmetic and algebra in calculus class. Make sure that you *visualize* each of the facts below in terms of vectors in the plane.

Cartesian form of complex numbers

$$x_1 + \boldsymbol{i}y_1 = x_2 + \boldsymbol{i}y_2 \text{ if and only if } x_1 = x_2 \text{ and } y_1 = y_2 \tag{1.11}$$

Addition and multiplication extend to complex numbers by using the commutative, associative, and distributive laws, and the fact that $\boldsymbol{i}\boldsymbol{i} = \boldsymbol{i}^2 = -1$. Addition of complex numbers may be visualized in terms of the vectors represented by the two numbers: shift the origin of one vector to the head of the other vector as shown in Figure 1.7. The conjugate of a complex number, written $\overline{\alpha}$, is the reflection of α through the real axis, as shown in Figure 1.8.

$$(x_1 + \boldsymbol{i}y_1) + (x_2 + \boldsymbol{i}y_2) = (x_1 + x_2) + \boldsymbol{i}(y_1 + y_2) \tag{1.12}$$

$$(x_1 + \boldsymbol{i}y_1)(x_2 + \boldsymbol{i}y_2) = (x_1x_2 - y_1y_2) + \boldsymbol{i}(x_1y_2 + x_2y_1) \tag{1.13}$$

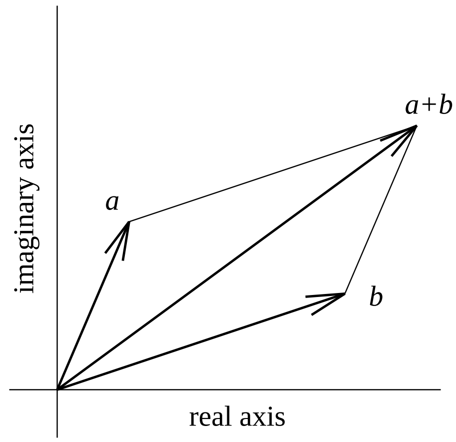

Figure 1.7: Adding two complex numbers

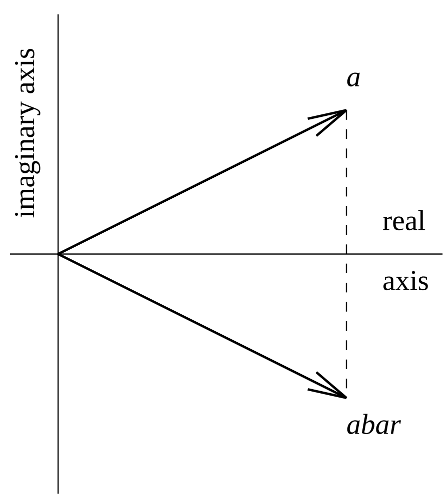

Figure 1.8: The conjugate of a complex number

$$
\begin{aligned}
\overline{x_1 + \boldsymbol{i}y_1} &= x_1 - \boldsymbol{i}y_1 && (1.14)\\
\alpha + \beta &= \beta + \alpha && (1.15)\\
(\alpha + \beta) + \gamma &= \alpha + (\beta + \gamma) && (1.16)\\
0 + \alpha &= \alpha && (1.17)\\
\alpha\beta &= \beta\alpha && (1.18)\\
(\alpha\beta)\gamma &= \alpha(\beta\gamma) && (1.19)\\
1\alpha &= \alpha && (1.20)\\
0\alpha &= 0 && (1.21)\\
\alpha(\beta + \gamma) &= \alpha\beta + \alpha\gamma && (1.22)
\end{aligned}
$$

The *real* and *imaginary* parts of a complex number are defined to select out the two components of the vector.

$$
\begin{aligned}
\Re(x + \boldsymbol{i}y) &= x && (1.23)\\
\Im(x + \boldsymbol{i}y) &= y && (1.24)\\
\Re(\alpha + \beta) &= \Re(\alpha) + \Re(\beta) && (1.25)\\
\Im(\alpha) + \Im(\beta) &= \Im(\alpha) + \Im(\beta) && (1.26)\\
\Re(\alpha\beta) &= \Re(\alpha)\Re(\beta) - \Im(\alpha)\Im(\beta) && (1.27)\\
\Im(\alpha\beta) &= \Re(\alpha)\Im(\beta) + \Im(\alpha)\Re(\beta) && (1.28)\\
\alpha &= \Re(\alpha) + \boldsymbol{i}\Im(\alpha) && (1.29)
\end{aligned}
$$

$$\alpha = \beta \text{ if and only if } \Re(\alpha) = \Re(\beta) \text{ and } \Im(\alpha) = \Im(\beta) \qquad (1.30)$$

$\Re(\alpha)$ and $\Im(\alpha)$ are called the *Cartesian* coordinates of the complex number α.

Polar form of complex numbers. The reason why complex numbers are particularly convenient for analyzing vibration is that they may be manipulated according to the *magnitude* (length of the vector) and *argument* (angle of the vector with respect to 1) as well. A magnitude is just a real number ≥ 0, representing the length of a vector. Angles are a bit trickier.

Rotational and directional angles. There are really two connected but different concepts that are both called "angles." First, there are *rotational angles* that measure an amount of rotation. A rotational angle may be any real number—positive numbers represent counterclockwise rotation, and negative numbers represent clockwise rotation. A rotational angle of 2π represents a full rotation counterclockwise. Even though the direction that an object points after a full rotation is the same as before the rotation, 2π represents a different rotation than 0 or -2π or 4π—suppose for example that we are measuring rotation of a wheel that winds up a spring.

The other sorts of angles are *directional angles* that measure the direction that a vector is pointing with reference to some conventional 0 direction (for complex numbers, 0 is the directional angle of the vector represented by 1). Directional angles must be in the half-open interval $[0, 2\pi)$. Many books and articles prefer to describe directional angles in the interval $(-\pi, \pi]$ (so, for example, the angle $3\pi/2$ in our notation becomes $-\pi/2$). It makes no essential difference which interval is used, since all arithmetic on directional angles is done on a circle of circumference 2π, rather than the usual real line. We may convert rotational angles to directional angles with the function $\bmod 2\pi$.

$$x \bmod 2\pi \;=\; x - 2\pi\lfloor x/(2\pi)\rfloor \tag{1.31}$$

$$0 \leq x \bmod 2\pi < 2\pi \tag{1.32}$$

When $x \bmod z = y \bmod z$ we often write $x = y \pmod z$ instead. This form suggests an alternate view of modular arithmetic: $x = y \pmod z$ means that x and y are two names for the same thing in the $\pmod z$ universe, even though they may be different numbers in the usual real number universe. If we apply a rotational angle x to rotate a vector from a starting position with directional angle 0, we get a new vector with directional angle $x \bmod 2\pi$. Notice that all rotational angles $x + 2k\pi$ for integers k correspond to the same directional angle. Given a directional angle x resulting from a rotation, there is no way to tell which of the infinitely many possible rotational angles generated x. To avoid becoming confused by the ambiguity in the word "angle," visualize each angle as either an amount of rotation or a static direction, instead of a pure abstract real number.

The angle of a complex number is a directional angle, so it is restricted to the interval $[0, 2\pi)$.

$$|x + \boldsymbol{i}y| \;=\; \sqrt{x^2 + y^2} \tag{1.33}$$

$$\arg(x + \boldsymbol{i}y) \;=\; \arctan(y/x) \bmod 2\pi \tag{1.34}$$

$$|\alpha| \;\geq\; 0 \tag{1.35}$$

$$\arg(\alpha) \geq 0 \quad (1.36)$$
$$\arg(\alpha) < 2\pi \quad (1.37)$$
$$(1.38)$$

$$\alpha = \beta \text{ if and only if } |\alpha| = |\beta| \text{ and } \arg(\alpha)\arg(\beta) \quad (1.39)$$

$\arg(0)$ is undefined, since it makes no sense to take the angle of a vector with magnitude 0. But, we let $\arg(\boldsymbol{i}y) = \pi/2$ for $y > 0$ and $\arg(\boldsymbol{i}y) = 3\pi/2$ for $y < 0$ in spite of the division by 0 in Equation 1.34, since $\boldsymbol{i}$ and $-\boldsymbol{i}$ are clearly at a right angles to the real axis (notice that $\lim_{z\to\infty} \arctan(z) = \pi/2$, $\lim_{z\to-\infty} \arctan(z) + 2\pi = -\pi/2 + 2\pi = 3\pi/2$). Figure 1.9 shows the relation between $\Re(\alpha)$, $\Im(\alpha)$, $|\alpha|$, and $\arg(\alpha)$ when α is drawn as a vector in a two-dimensional space. $|\alpha|$ and $\arg(\alpha)$ are called the *polar* coordinates of the complex number α. Addition of complex numbers is easiest to do by manipulating the real and imaginary parts, but multiplication and division may be defined very nicely on the magnitude and angle.

$$|\alpha\beta| = |\alpha||\beta| \quad (1.40)$$
$$\arg(\alpha\beta) = \arg(\alpha) + \arg(\beta) \bmod 2\pi \quad (1.41)$$
$$|\alpha/\beta| = |\alpha|/|\beta| \quad (1.42)$$
$$\arg(\alpha/\beta) = \arg(\alpha) - \arg(\beta) \bmod 2\pi \quad (1.43)$$
$$|-\alpha| = |\alpha| \quad (1.44)$$
$$\arg(-\alpha) = \arg(\alpha) + \pi \bmod 2\pi \quad (1.45)$$
$$|x\alpha| = |-x\alpha| = x|\alpha| \text{ for } x \geq 0 \quad (1.46)$$
$$\arg(x\alpha) = \arg(\alpha) \text{ for } x > 0 \quad (1.47)$$
$$|\overline{\alpha}| = |\alpha| \quad (1.48)$$
$$\arg(\overline{\alpha}) = -\arg(\alpha) \bmod 2\pi \quad (1.49)$$
$$\alpha = |\alpha|(\cos(\arg(\alpha)) + \boldsymbol{i}\sin(\arg(\alpha))) \quad (1.50)$$

Using Equations 1.40 and 1.41, we see that a complex number α of magnitude 1 acts as a rotator: the multiplication $\alpha\beta$ rotates β by the angle $\arg(\alpha)$. In particular, multiplication by $\boldsymbol{i}$ rotates a vector counterclockwise by $\pi/2$ (right angle). So, letting the single complex number $\rho = x + \boldsymbol{i}y$ represent the rotor state $\langle x, y\rangle$, we may express Equations 1.5 and 1.6 as a single equation.

$$\rho' = \boldsymbol{i}A\rho \quad (1.51)$$

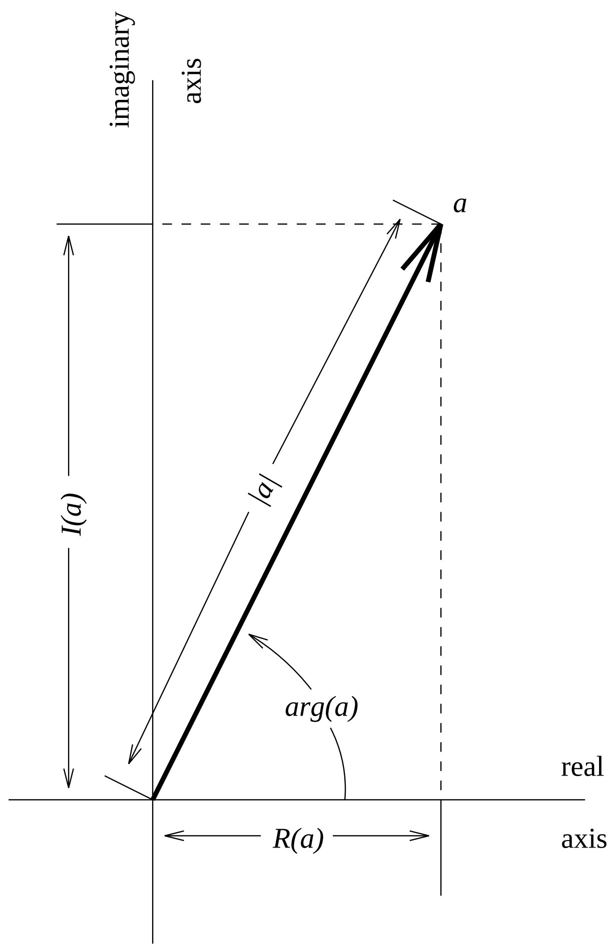

Figure 1.9: Complex number: Cartesian and polar coordinates

The elliptical rotor system of Equations 1.3 and 1.4 may also be expressed as a single equation using the conjugate operation.

$$\rho' = \boldsymbol{i}((A+B)\rho + (A-B)\overline{\rho})/2 \qquad (1.52)$$

Complex number to real power. Equation 1.41 is also the key to understanding exponentiation of complex numbers. Notice that the angle behaves *logarithmically* with respect to addition and multiplication (think of the analogous equation $\ln(xy) = \ln(x) + \ln(y)$). Since exponentiation is essentially iterated multiplication, and complex multiplication is additive on angles, complex exponentiation has a multiplicative effect on angles. Consider first a complex number α raised to the power of a real number x.

$$\begin{aligned}
|\alpha^x| &= |\alpha|^x && (1.53)\\
\arg(\alpha^x) &= x\arg(\alpha) && (1.54)\\
\alpha^{x+y} &= \alpha^x\alpha^y && (1.55)\\
\alpha^{xy} &= (\alpha^x)^y && (1.56)\\
(\alpha\beta)^x &= \alpha^x\beta^x && (1.57)\\
\overline{\alpha}^x &= \overline{\alpha^x} && (1.58)\\
\alpha^1 &= \alpha && (1.59)\\
\alpha^0 &= 1 && (1.60)
\end{aligned}$$

Equation 1.54 begins to reveal the power of complex numbers for analyzing vibration. Think of a rotor with amplitude R, frequency F, and phase 0 (starting value $R + \boldsymbol{i}0$ at time $t = 0$). The state of the rotor at any time t is $\rho = R\boldsymbol{i}^{4Ft}$ (we multiply F by 4 because $\arg(\boldsymbol{i}) = \pi/2$ is $1/4$ of 2π radians, which is a full rotation). For an elliptical rotor with starting value R_1 that crosses the imaginary axis at $\boldsymbol{i}R_2$, the state at time t is $(R_1 + R_2)\boldsymbol{i}^{4Ft}/2 + (R_1 - R_2)\overline{\boldsymbol{i}^{4Ft}}/2$. Equivalent expressions include $(R_1+R_2)\boldsymbol{i}^{4Ft}/2+(R_1-R_2)(-\boldsymbol{i})^{4Ft}/2$ and $(R_1+R_2)\boldsymbol{i}^{4Ft}/2+(R_1-R_2)\boldsymbol{i}^{-4Ft}/2$. The last of these is the most popular, and leads to the notion of a $-F$ frequency component in an elliptical rotor.

Complex exponents. The most important fact about complex numbers for the study of vibration is the rule for raising the real number $e = 2.71828\ldots$ to a complex power.

$$e^{\boldsymbol{i}y} = \cos(y) + \boldsymbol{i}\sin(y) \qquad (1.61)$$

$$
\begin{aligned}
e^{x+\boldsymbol{i}y} &= e^x(\cos(y)+\boldsymbol{i}\sin(y)) && (1.62)\\
|e^\alpha| &= e^{\Re(\alpha)} && (1.63)\\
\arg(e^\alpha) &= \Im(\alpha) \bmod 2\pi && (1.64)\\
e^{\overline{\alpha}} &= \overline{e^\alpha} && (1.65)\\
\alpha &= |\alpha|e^{\boldsymbol{i}\arg(\alpha)} && (1.66)\\
\alpha &= e^{\ln(|\alpha|)+\boldsymbol{i}\arg(\alpha)} \text{ for } \alpha \neq 0 && (1.67)
\end{aligned}
$$

Equation 1.61, known as *Euler's formula* in honor of the famous mathematician who discovered it, is the most important single equation for the study of vibration. It allows us to reason about trigonometric functions by using the relatively easy-to-remember properties of exponentiation. Computer algebra systems typically convert trigonometric formulae into exponential form in order to simplify them more efficiently.

For our purposes, the derivation of Euler's formula is not as important as the formula itself. To see why the formula is sensible, consider the ordinary differential equation defining the exponential function for real numbers:

$$x' = Ax \qquad (1.68)$$

The most interesting solution to equation 1.68 is the one with initial condition $x(0) = 1$, and this leads to

$$x(t) = e^{At} \qquad (1.69)$$

That is, the (scaled) exponential function e^{At} is characterized by its inital value and the fact that its slope at each time t is A times its value at time t—the larger it gets, the faster it grows. Notice that equation 1.51 has the same form as equation 1.68, but it describes a complex-valued function, and the multiplier is an imaginary number $\boldsymbol{i}A$ rather than a real number A. So, it is sensible to regard the natural solution to equation 1.51, which is a rotor, as the function $e^{\boldsymbol{i}At}$.

Euler's formula also gives us another way to represent each complex number α—instead of the usual form $\Re(\alpha) + \boldsymbol{i}\Im(\alpha)$ we may write $|\alpha|e^{\boldsymbol{i}\arg(\alpha)}$. For $\alpha \neq 0$ we may also write $e^{\ln(|\alpha|)+\boldsymbol{i}\arg(\alpha)}$. Unlike the additive Cartesian form $x + \boldsymbol{i}y$, the exponential polar $re^{\boldsymbol{i}w}$ form for a complex number is not unique, since $re^{\boldsymbol{i}w} = re^{\boldsymbol{i}(w+2k\pi)}$ and $0e^{\boldsymbol{i}w_1} = 0e^{\boldsymbol{i}w_2}$.

$$r_1 e^{\boldsymbol{i}w_1} = r_2 e^{\boldsymbol{i}w_2} \text{ if and only if } \quad \begin{array}{c} r_1 = r_2 = 0 \\ \text{or} \\ r_1 = r_2 \text{ and } w_1 = w_2 \pmod{2\pi} \\ \text{or} \\ r_1 = -r_2 \text{ and } w_1 = (w_2 + \pi) \pmod{2\pi} \end{array} \tag{1.70}$$

In essence, exponentiation is a kind of conversion between Cartesian and polar coordinates: the polar coordinates of e^α are $|e^\alpha| = e^{\Re(\alpha)}$ and $\arg(e^\alpha) = (\Im(\alpha) \bmod 2\pi)$. So, the Cartesian coordinates of α turn into the polar coordinates of e^α. Notice that $\Im(\alpha)$ is naturally understood as a rotational angle, while $\arg(e^\alpha)$ is a directional angle.

Euler's formula (Equation 1.61) allows an even nicer way to analyze a rotor with amplitude R, frequency F, and phase 0 (starting value $R+\boldsymbol{i}0 = Re^0$ at time $t = 0$): the state at any time t is just $\rho = Re^{\boldsymbol{i}2\pi Ft}$ (such exponential expressions are sometimes called *phasors* in the engineering literature). For the elliptical rotor starting at R_1 and crossing the imaginary axis at $\boldsymbol{i}R_2$, the state at time t is $((R_1 + R_2)e^{\boldsymbol{i}2\pi Ft} + (R_1 - R_2)\overline{e^{\boldsymbol{i}2\pi Ft}})/2$, or $((R_1 + R_2)e^{\boldsymbol{i}2\pi Ft} + (R_1 - R_2)e^{-\boldsymbol{i}2\pi Ft})/2$. And, it is particularly easy to construct a complex number with magnitude 1 to rotate other numbers by a given angle w: use $e^{\boldsymbol{i}w}$. Look back at Figures 1.3 and 1.6 again, and interpret them in terms of complex numbers.

Now, the way to understand exponentiation β^α with an arbitrary complex base β is to first write $\beta = |\beta|e^{\boldsymbol{i}\arg(\beta)}$, and then use the rules for exponentiation with base e.

$$(re^{\boldsymbol{i}w})^{x+\boldsymbol{i}y} = r^v e^{-wy} e^{\boldsymbol{i}(wx+y\ln(r))} \tag{1.71}$$

$$|\beta^\alpha| = |\beta|^{\Re(\alpha)} e^{-\arg(\beta)\Im(\alpha)} \tag{1.72}$$

$$\arg(\beta^\alpha) = \arg(\beta)\Re(\alpha) + \ln(|\beta|)\Im(\alpha) \tag{1.73}$$

These equations are rather complicated, and fortunately we will not be using them much. Work them through for exercise with complex numbers, and convince yourself that they follow from the earlier rules. Notice how exponentiation mixes together the Cartesian coordinates of the exponent with the polar coordinates of the base.

Complex logarithms. Euler's formula makes it easy to define the natural (base e) logarithm of a complex number.

$$\begin{aligned}
\ln(\alpha) &= \ln|\alpha| + \boldsymbol{i}\arg(\alpha) \text{ for } \alpha \neq 0 && (1.74)\\
e^{\ln(\alpha)} &= \alpha \text{ for } \alpha \neq 0 && (1.75)\\
\Re(\ln(\alpha)) &= \ln(|\alpha|) && (1.76)\\
\Im(\ln(\alpha)) &= \arg(\alpha) && (1.77)
\end{aligned}$$

$\ln(0)$ is undefined. Notice that, for positive real numbers $x > 0$ there is a unique real number y such that $e^y = x$. But, even for positive real numbers x there are infinitely many complex numbers β such that $e^\beta = x$. That is, while Equation 1.75 defines the natural logarithm uniquely as a real value, it has infinitely many complex solutions, since $e^\alpha = e^{\alpha+2k\pi}$ for all integers k. The particular choice above for the imaginary part of $\ln(\alpha)$ is arbitrary, just as the particular interval $[0, 2\pi)$ for directional angles is arbitrary. Notice that this choice restricts all complex logarithms $\ln(\alpha)$ to the horizontal stripe in the complex plane where $0 \leq \Im(\ln(\alpha)) < 2\pi$.

$$\log_\beta(\alpha) = \ln(\alpha)/\ln(\beta) \text{ for } \alpha, \beta \neq 0 \qquad (1.78)$$

1.2.3 Sound Signals in the Time Domain

In general, the sounds that we would like to create and analyze are much more complicated than the sounds produced by simple rotors. But, we will continue to model sounds by complex-valued functions σ depending on a real number parameter t standing for time. Such functions are called *sound signals in the time domain.* Later, in Chapters 4 and 6, we will see other mathematical representations of sound, but signals in the time domain are the easiest models to relate intuitively to the physical signals that enter the ear. Widely used digital input and output devices for sound are also most easily understood in terms of signals in the time domain. Most books and papers on sound consider real-valued time signals, and most electronic devices, both digital and analog, for analyzing or creating sound deal only with real values. Many analysis and synthesis techniques, however, are best understood in terms of a complex signal σ. We may always project the complex signal σ to a real signal by taking $\Re(\sigma)$. Just as many systems for manipulating graphic images deal with three dimensional models, and project them to two dimensions at the last stage before displaying them on video screens, we will think of sound signals as two dimensional, and project to one dimension at the last stage before rendering them through loudspeakers.

Given a sound signal σ in the time domain, and a particular time t, the *instantaneous amplitude* of σ at time t is $|\sigma(t)|$, the *instantaneous phase* is $\arg(\sigma(t))$, and the *instantaneous frequency* is $(d\arg(\sigma)/dt)(t)$. These are interesting quantities to discuss, and may be useful in analyzing sound, but they do not necessarily have the perceptual impact of the corresponding constant quantities associated with a simple rotor.

Not every complex-valued function σ of t makes sense as a sound signal in the time domain. Some reasonable relation must hold between the real and complex components of σ. But, it is not clear precisely what relation to require in general. Particular physical interpretations of σ impose certain constraints—for example if $\Re(\sigma)$ is the velocity of a physical object, and $\Im(\sigma)$ is the displacement of the same object, then $\Re(\sigma) = d\Im(\sigma)/dt$. When $\Im(\sigma) = R\sin(Ft)$, $d\Im(\sigma)/dt = FR\cos(Ft)$, so this derivative constraint forces rotors to be elliptical, with aspect ratio proportional to the frequency. Circular rotors are much more convenient mathematically. Roughly speaking, we would like to restrict sound signals σ so that $\Im(\sigma)$ is essentially the same as $\Re(\sigma)$ with a phase difference of π (90°)—such signals are said to be *in quadrature,* since the angle π is one quarter of the full circle. The problem is that many different frequencies may be present in σ. In Chapter 4 we see a precise definition of this quadrature constraint.

Figures 1.10, 1.11, and 1.12 show examples of sound signals in the time domain that are slightly more complicated than the basic helix. In each case, part (a) shows a three-dimensional plot of the complex-valued function, and part (b) shows a two-dimensional plot of the real and imaginary components.

1.3 Exercises

1. Take the spring system of Figure 1.1, rotate the track to a vertical orientation, and let a constant gravitational force act on the mass. The stable position about which the mass oscillates is no longer at the point where the spring attaches to the track, but some distance below that point where the force exerted by the spring exactly cancels gravity. Does the frequency of the vibrating spring increase or decrease as a result of the influence of gravity? Explain briefly.

(a)

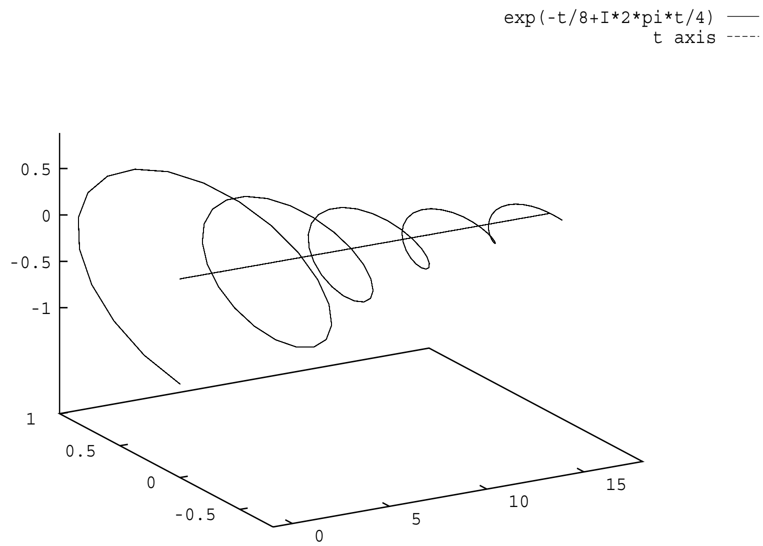

(b)

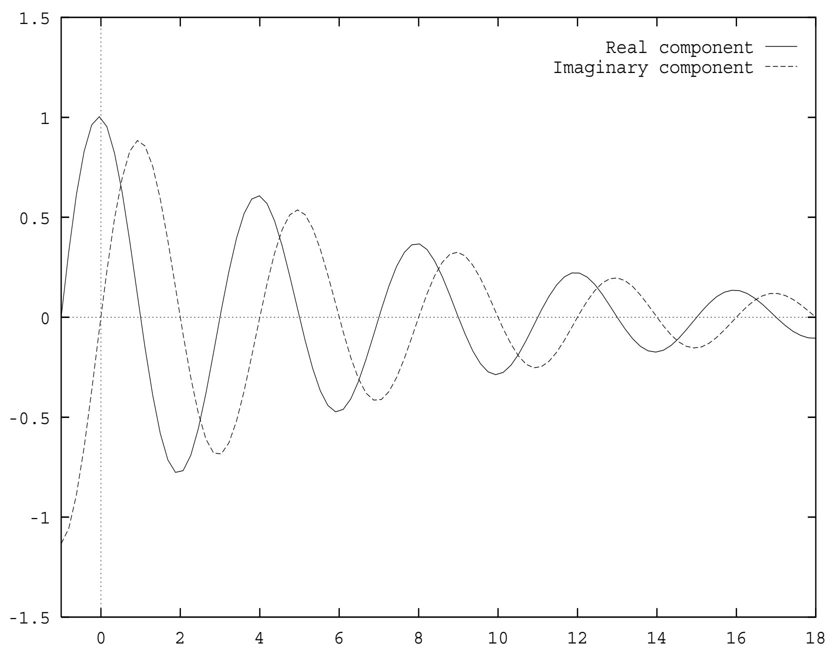

Figure 1.10: The signal $e^{-t/8+i2\pi t/4}$: decaying helix

(a)

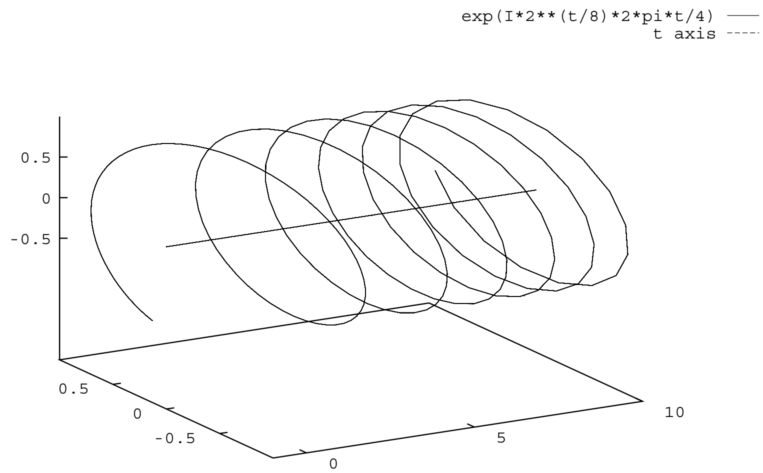

(b)

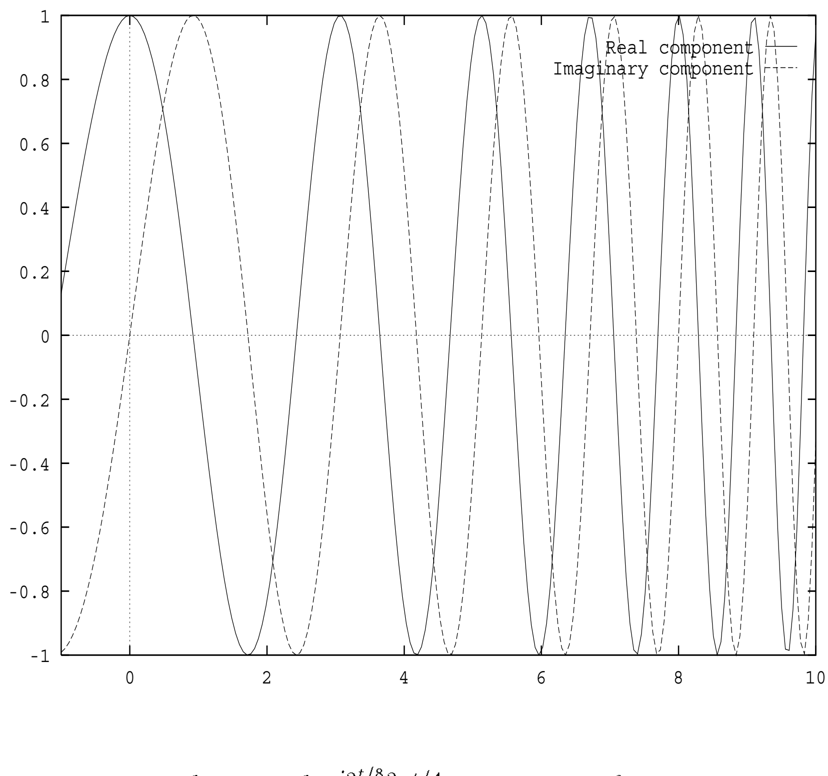

Figure 1.11: The signal $e^{i2^{t/8}2\pi t/4}$: increasing frequency

(a)

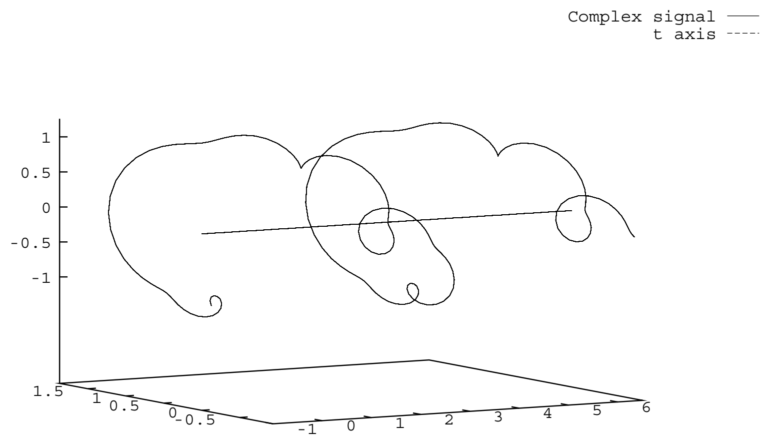

(b)

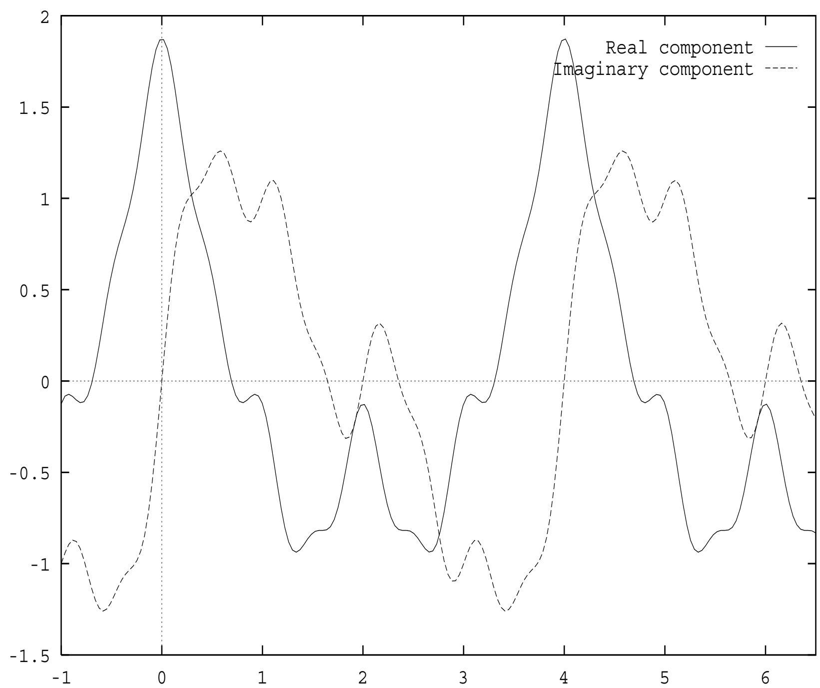

Figure 1.12: The signal $e^{i2\pi t/4} + e^{i4\pi t/4}/2 + e^{i8\pi t/4}/4 + e^{i16\pi t/4}/8$: sum of 4 helixes

2. Consider a vibrating spring system in which the motion of the mass is opposed by a certain amount of friction. In order to analyze such a system, do we change Equation 1.1, Equation 1.2, both, or neither? Explain briefly.

3. Consider a vibrating system with a mass that is attracted to the center of vibration by a gravitational force instead of a spring. In such a system, the period of vibration depends on the amplitude. As the amplitude of the vibration increases, does the period increase or decrease? Explain briefly.

4. Notice how Equation 1.5 has a positive multiplier A, while Equation 1.6 has a negative multiplier $-A$. There are three other possibilities: (a) both multipliers negative, (b) both multipliers positive, (c) the first multiplier negative and the second positive. Describe briefly and qualitatively the behavior of a system described by each of the variants (a–c). Draw pictures analogous to Figure 1.3(b) to help explain.

5. When $A > B$ in Equations 1.3 and 1.4, the path of the state vector $\langle x, y \rangle$ is an ellipse. Which axis of the ellipse is longer, the x axis or the y axis? Explain briefly, using precise mathematical information derived from Equations 1.3 and 1.4. Hint: Derive slightly different equations relating $C_y y'$ to $C_x x$ and $C_x x'$ to $C_y y$ for cleverly chosen constant multipliers C_x and C_y.

6. Derive simple formulae representing the frequency and period of the vibrating system of Equations 1.3 and 1.4 in terms of the constants A and B. Hint: Look at Equations 1.7 and 1.8. Differentiate both sides of both equations. Solve the special case where $A = B$. Then, apply the scaling of x and y by constants C_x and C_y that you used in Exercise 5.

7. In an elliptical rotor system obeying Equations 1.3 and 1.4 the speed with which the state point travels around the ellipse is not constant.

   (a) Where is this speed the least, and where is it greatest? Explain briefly.

   (b) Answer the same question for the *angular* speed of the state vector—the speed at which its angle with the x axis changes.

8. For each of the following operations on complex numbers α and β, state whether it is more convenient to represent each number in Cartesian or polar coordinates, or whether both are equally convenient. Sometimes the answer is different for β than for α.

   (a) $\alpha + \beta$
   (b) $\alpha - \beta$
   (c) $\overline{\alpha}$
   (d) $\alpha\beta$
   (e) α/β
   (f) β^α
   (g) $\log_\beta(\alpha)$

9. Derive formulae for the Cartesian coordinates $\Re(\beta^\alpha)$ and $\Im(\beta^\alpha)$ of β^α in terms of Cartesian and/or polar coordinates of α and β.

10. Derive the following trigonometric identities, using Euler's formula (Equation 1.61) and easy algebraic manipulations of additions, subtractions, multiplications, and divisions of complex numbers. Note that $\cos^2(x)$ and $\sin^2(x)$ are conventional ways of writing $(\cos(x))^2$ and $(\sin(x))^2$, respectively.

    (a) $$\cos(2x) = \cos^2(x) - \sin^2(x)$$
    (b) $$\sin(2x) = 2\cos(x)\sin(x)$$
    (c) $$\cos^2(x) = (1 + \cos(2x))/2$$
    (d) $$\sin^2(x) = (1 - \cos(2x))/2$$
    (e) $$\cos(x) + \cos(y) = 2\cos((x+y)/2)\cos((x-y)/2)$$
    (f) $$\sin(x) + \sin(y) = 2\sin((x+y)/2)\cos((x-y)/2)$$
    (g) $$\cos(x)\cos(y) = (\cos(x-y) + \cos(x+y))/2$$
    (h) $$\sin(x)\sin(y) = (\cos(x-y) - \cos(x+y))/2$$
    (i) $$\sin(x)\cos(y) = (\sin(x-y) + \sin(x+y))/2$$
    (j) $$\cos(x+y) = \cos(x)\cos(y) - \sin(x)\sin(y)$$
    (k) $$\sin(x+y) = \sin(x)\cos(y) + \cos(x)\sin(y)$$

11. We saw how to express the state of an elliptical rotor at time t in the form $ae^{i2\pi ft} + b\overline{e^{i2\pi ft}}$, where a is the average of the real and imaginary intercepts of the ellipse, and b is half their difference. Derive a nice formula for the state of an elliptical rotor whose major and minor axes are different from the real and imaginary axes.

Chapter 2

Perceptual Foundations of Sound

This chapter is a particularly rough draft, with a lot of missing information.

In every section: limits of human sound perception.

This chapter sketches the structure of human sound perception in a deliberately simplistic and superficial way. I believe that the best digital models for sound production will be informed by audiology, but they will sacrifice a lot of perceptual precision for mathematical simplicity. By rough analogy, there are lots of qualities of human visual perception that are ignored by the pixel model of graphics. Also, general-purpose models for sound production must work for almost all listeners, so they cannot be designed around details of perception that vary from person to person. Later chapters will investigate more precisely the mathematical qualities of sound that affect perception, but we will stick with a very approximate and intuitive notion of perception itself.

We are interested in sound as a medium that may be used for communication. Particular forms of audible communication, such as music and speech, may be highly specialized to their purposes and to the acoustic resources available to them for generating sound. There must be some very general structural qualities of sound that are present in essentially all uses of sound for communication. Each particular form of audible communication may exploit these general structural qualities in very different ways.

By rough analogy to visual communication, notice that almost all visual scenes may be described in terms of structural concepts such as region, edge, texture, color, brightness. Written communication in English exploits the shapes of regions with contrasting brightness, and the edges of those regions, to provide recognizable alphabetic characters of the Roman alphabet. Architectural drawings exploit edges in a radically different way. Perspective pictures draw on texture and color in yet other ways to communicate layouts of physical objects. In this chapter we seek an intuitive understanding of structural qualities of sound roughly at the level of region, edge, texture, color, brightness in video.

2.1 The Ear as a Frequency Analyzer

The key receptive structure in the ear is the *cochlea,* a spiral-shaped tube containing lots of little hairs that vibrate with the surrounding fluid. Is it air or some body liquid? From our point of view, each hair is a physical realization of a rotor. Somehow (the *how* is still the topic of some debate) each hair is tuned to a narrow range of frequencies, and stimulates an assigned nerve ending proportionally to the amount of excitation it receives within its frequency range. So, the human ear is roughly a frequency analyzer, passing on a spectral presentation of sound at each instant to the brain for further analysis.

2.2 Sound Imaging—What is "a Sound"?

I call a complex of sound that is presented to a listener an "audible scene." Many audible scenes deompose naturally into the sum of several components that are perceived as units, vaguely analogous to contiguous regions in a visual scene. The decomposition is often ambiguous, and sometimes there is no sensible decomposition, but the notion of a perceived contiguous piece of sound is likely to be useful whenever it applies. I call such an intuitive unit in an audible scene "a sound." In well-articulated musical pieces, a single note by a single instrument is a sound. In speech, the notion is more ambiguous, but perhaps a phoneme or segment of a phoneme may be understood as a sound.

Automated analysis of audible scenes into individual sounds is extremely difficult, because it must resolve all of the ambiguities that arise. Synthesis by adding up individual sounds to create audible scenes is much more tractable, since the instructions for synthesizing a given scene can specify an interpretation explicitly. A synthesis method based on adding individual sounds together might be very useful even if it doesn't guarantee that every object described as "a sound" by the system is perceived as a single sound—as long as there is a good heuristic correlation between description and perception the method can succeed.

The precise way in which the ear and brain decompose an audible scene into individual sounds is not understood. The spatial location of sound sources, as detected by the stereo effects of pairs of ears and by the asymmetric distortion induced by the funny shapes of our heads and external ears, certainly plays an important part. We will ignore spatial location, not because it isn't important, but because *for the purposes of synthesis,* it can probably be separated from monaural qualities. To synthesize an audible scene, we may describe sounds, then describe where each sound is placed, and these two parts of our description may be essentially independent. For analysis, they are probably tangled together inextricably.

Ignoring location, the qualities that make a particular complex vibration sensible to regard as an individual sound probably have to do with the frequency components of that vibration. We prefer to group frequency components together perceptually when their beginnings, and to a lesser extent their endings, are nearly simultaneous. Also, we prefer to group frequencies that are very close to being integer multiples of some audible frequency, which may or may not be present itself—stated another way we prefer to associate frequencies whose ratios are very close to rational numbers with small integer numerators and denominators. These qualitative observations are *very* far from providing a useful basis for analysis, but they may serve as heuristic guides in considering synthesis techniques.

2.3 Perceptual Parameters of a Sound

2.3.1 Pitch

Pitch is the quality of a sound that leads us to consider it "higher" or "lower" than another sound. Some sounds, such as engine noises and drum beats, yield only a vague sense of high or low pitch. Other sounds, such as notes of bird songs and of melodic musical instruments, yield a fairly precise sense of pitch that can be measured numerically, with most listeners agreeing that the measurement is correct.

At first approximation, the pitch of a sound is its frequency. The human ear detects frequencies from about 20 Hertz (cycles per second) to about 20,000 Hertz. Perceived pitch is essentially the *logarithm* of frequency. Multiplying a frequency is perceived as adding to the pitch. For example, on a piano keyboard, the interval (difference in pitch) called an *octave* is heard as the result of moving up 12 *half steps* (the interval from one key to the next higher—usually one is white and the other black), but it is essentially a multiplication of the frequency by 2. The interval called a *perfect fifth,* heard as moving up 7 half steps, multiplies frequency by approximately 3/2.

We need to know the pitch resolution of the human ear.

But, it's not that simple. Perception of pitch is affected by loudness (loud sounds tend to sound higher in pitch than soft sounds of the same frequency), and there may be many other small but significant influences on perceived pitch. My hunch is that *most* of these should not affect the structure of a general-purpose model of sound, but rather should be viewed as fine points to be applied outside of the model, when polishing a sound definition to its final form, only when great precision is truly required. For most purposes, lots of perceptual subtleties are best ignored.

One major complication in pitch perception probably *will* affect the structure of good digital models of sound. Although pitch is *essentially* the logarithm of frequency, perception of pitch is tied more closely to the *relation* between a number of component frequencies in a sound, rather than to the frequency of one particular component. Specifically, when a sound is nearly *harmonic*—when most of the frequency components of a sound are nearly integer multiples of another audible frequency F, called the *fundamental* pitch of that sound—we tend to hear a pitch given by $\ln F$. *The frequency F itself need not be present!* This seems spooky at first, but it is probably a very sensible adaptation of aural perception to the fact that some components of a sound may be masked by noise. Perception of the "missing fundamental" is roughly analogous to the visual perception of an entire object, even though parts of it are hidden behind other objects.

The perception of pitch intervals is also a bit more complicated than merely subtracting one pitch from another. When we perceive the pitch interval between two nearly harmonic sounds s_1 with fundamental frequency F_1 and s_2 with fundamental F_2, we seem to overlay their component frequencies. If the component of s_1 with *approximate* frequency MF_1 is close enough to the component of s_2 with *approximate* frequency NF_2 (M and N are integers), this influences us toward perceiving an interval determined by

$\ln N/M = \ln N - \ln M$, rather than $\ln F_1/F_2 = \ln F_1 - \ln F_2$. Each pair of components that overlays closely enough influences the perceived interval, and it is hard to characterize the way in which these influences add up. But, there are plenty of sounds that are nearly enough harmonic to have a musical effect, but far enough from perfect integer ratios to confuse the perception of intervals. The piano and the bagpipes are examples of instruments with substantial deviations from harmonic sound, and the comparison of pitch between them and nearly perfect harmonic sounds, such as the sounds of most orchestral instruments, is quite tricky.

The precision of pitch perception is roughly constant within audible frequency limits. Since pitch is the logarithm of frequency, this means that frequency precision is much better for lower frequencies and poorer for higher frequencies. Section 2.3.5 discusses the perception of time for sound, which has a variation of precision inverse to the variation of frequency precision.

2.3.2 Loudness

Loudness of a simple helical signal is roughly the logarithm of its power (the rate at which it delivers energy). Notice that when two helical signals have the same amplitude, the one with higher frequency also has higher power, because it moves faster. The exact power in a signal depends on the precise physical interpretation of the signal, but in general the power in a helical signal $Re^{\boldsymbol{i}2\pi Ft}$ is proportional to some polynomial in R and F, and at least as big as RF, so perceived loudness is roughly proportional to $\ln(RF) = \ln(R) + \ln(F)$. But, perceived loudness varies according to the sensitivity of the ear at the given frequency, so signals at frequencies near the limits of audible frequencies seem softer than signals of equal power near the center.

We need information on the units used to measure loudness and the limits of normal human perception.

When a number of frequencies are present in a sound, it seems sensible that the perceived loudness will be roughly proportional to the logarithm of the sum of all power within audible frequency limits. This seems sensible, but it's wrong. The perception of loudness in complex sounds is influenced by the *critical bands* of human sound perception—frequency bands containing a spread of frequencies roughly spanning a musical minor third, so the highest frequency in a band is roughly 6/5 times the lowest. These bands are not discrete, rather they overlap continuously across the range of audible frequencies, varying slightly in width depending on the center frequency.

Two helical signals within a critical band tend to add their powers, so that the perceived loudness within a critical band is close to the logarithm of the total power in the band. But, two helical signals whose frequencies differ by more than a critical band tend to add their perceived loudness *after* the individual loudnesses are taken as the logarithm of power. Since $\ln(x) + \ln(y) > \ln(x+y)$, the same amount of sonic power sounds louder when spread over a larger frequency range. The precise computation of loudness from power spectrum is quite complicated because of the overlapping of critical bands.

I doubt that the critical band concept will have an impact on the lowest levels of sound modelling, but it clearly has a profound effect on perception, and therefore on the construction of highly polished sounds.

2.3.3 Timbre

Two sounds of the same pitch and loudness may have recognizably different qualities: for instance the sounds of string instruments vs. reed instruments in the orchestra. These distinguishing qualities of sound are called *timbre,* and are sometimes compared to visible color. Compared to pitch and loudness, timbre is not at all well defined. It clearly has a lot to do with the relative strengths of different frequency components of a sound, called the *partials.* But, it is also affected seriously by some aspects of the time development of partials—particularly but not exclusively by the increase in amplitude of partials at the beginning of a sound, called the *attack* in music. Different partials of a musical sound typically increase at very different rates, and these differences are crucial to the identification of a sound with a particular instrument. For example, the sounds of brass instruments are recognized partly by the quicker development of lower frequencies than higher frequencies.

At first approximation it seems that two sounds of different pitch will have the same perceived timbre when the spectral content of one looks just like the other, but shifted in frequency. For example a sound with a component of amplitude 1 at 100 Hertz, amplitude 0.5 at 200 Hertz, and amplitude 0.25 at 300 Hertz might be expected to be qualitatively similar to one with amplitude 1 at 250 Hertz, 0.5 at 500 Hertz, and 0.25 at 750 Hertz. In this sort of case, the second sound might be produced by recording the first one on tape, then playing it back with the tape moving faster. The famous singing chipmunks demonstrate the fallacy in this expectation—they do not sound at all like their creator singing higher.

A more accurate notion of timbre must take into account the fact that sound perception has adapted to the way that many sound producers, including the human voice and most musical instruments, create their sounds by a two-stage process. First, there is some sort of vibrating structure, such as the vocal chord, violin string, oboe reeds, which may follow the shifted partials model fairly well. But, the sound coming from this first vibrating structure filters through another resonating structure, such as the human head, the body of the violin, the body of the oboe, which scales the amplitudes of partials according to its responsiveness at different frequencies. The responiveness of the second structure does *not* take a frequency shift when the incoming pitch changes, so it changes the relative strengths of partials depending on their absolute frequencies, and not just their ratios to the given pitch. This filtering structure is sometimes called a *formant filter,* because it may often be characterized by a small number of highly resonant frequency bands, called *formants.* Human sound perception seems to have adapted to recognizing the constancy of formant filters when they are stimulated by a variety of incoming sounds at different pitches. This is vaguely analogous to the tendency of human visual perception to perceive the reflective properties of a given pigment as its *color,* even under radically different illuminations that may change the actual spectrum reaching the eye quite severely.

2.3.4 Transient Effects

2.3.5 Sound Events

Although abstract physics recognizes *time* as a single one-dimensional continuum (at least for any single observer), different intervals of time may be perceived as if they are in completely different dimensions, depending on the lengths of the intervals and the sorts of perceptible changes that occur during them. For example, in visual perception, changes in electromagnetic flux on a scale of millionths of a second are not perceived as time at all, but rather determine the frequency of light, and thereby contribute to the perception of color. Changes on a scale of tenths of a second or longer are generally perceived as temporal events involving changes in visual qualities, including color. The huge gap between the electromagnetic time scale and the event-sequence time scale make it easy to classify particular changes unambiguously into one class or the other.

Sound perception seems to have at least three time scales that are perceived quite differently, and they all overlap to make things more complicated.

1. Changes in air pressure on a scale of about $1/20,000$th of a second to $1/20$th of a second are on the *sonic* time scale. They are not perceived as time developments at all, but determine the frequency of components of a sound, and thereby contribute to pitch, timbre, and loudness perception. Sonic time for sound is analogous to the electromagnetic time scale for visual perception.

2. Slightly slower what range? changes in the amplitudes of various frequency components are on the *transitional* time scale. They are perceived as time developments, such as the attack initiating a musical note, but the exact time sequence is hard or perhaps impossible to trace perceptually. The special quality of the transitional time scale is demonstrated by playing sounds backwards. While a sequence of events played backwards may be recognized accurately, even if it is physically ridiculous (for example, a reversed movie of someone walking), the time-reversal of sound transitions makes us perceive them completely differently. I suspect that a person who heard a time-reversed sound for the first time, with no clue such as a view of the record being spun backwards, might not even recognize it as the time reversal of something. Even knowing that a sound is time reversed, it is difficult to tell intuitively what the forward version sounds like. I am not aware of any visual phenomenon analogous to the transitional time scale for sound.

3. Changes in the frequency components of sound on a scale of more accurate lower bound? perhaps $1/30$th of a second and longer are often perceived as sequences of events. The *event-sequence* time scale for sound is analogous to the one for visual perception, and they operate in a similar range. For example, the sequence of notes in a scale, or the sequence of clicks in a rhythmic form, are perceived on the event-sequence time scale. The reversal of a sequence of events may be musically or physically peculiar, but it is relatively easy to recognize.

Even the sonic and event-sequence time scales overlap for sound, with the transitional scale in between and overlapping both. This makes the understanding of time developments in sound quite subtle in some cases. In particular, the boundaries are sensitive to frequency. For low-frequency components of sound, the boundaries of the scales move toward longer time intervals, and for high-frequency components they move toward shorter intervals. It takes about one full period (rotation) of a helix to recognize the frequency. So, changes in a helical component of sound can only be detected when they are not too short compared to the time of a complete period.

The inverse relation between frequency precision, which is best for low frequencies, and time precision, which is best for high frequencies, is striking. It is not an accident, but comes from fundamental physical limitations, which limit the *product* of time and frequency precision, so that when one improves, the other gets proportionately worse. The same mathematical form produces the Heisenberg uncertainty principle in quantum mechanics.

Chapter 3

Digital Sampled Sound

In Chapter 1 we modelled sound as a function from a real value t representing time to a complex value σ representing a two-dimensional state of some vibrating system. In order to deal with sound digitally, we must somehow reduce such a signal to a finite number of symbols from a discrete set of possibilities, such as the bits 0 and 1 that current digital computers use to represent all information. The digitization of sound is normally achieved by two independent steps, each of which has consequences for the fidelity with which sound is produced digitally.

3.1 Discrete time

Sound signals in the discrete time domain. The usual first logical step in digitizing sound is to approximate the continuous domain of real values representing time by a discrete set of equally spaced values. Let S be a positive real number, called the *sampling rate* (the number of samples to take in a unit of time). $\mathcal{N}$ represents the set of all positive and negative integers. $\mathcal{T}_S$ represents the domain of *discrete time* with sampling rate S.

$$\mathcal{N} = \{\ldots, -2, -1, 0, 1, 2, 3, \ldots\} \quad (3.1)$$

$$\mathcal{T}_S = \{k/S : k \in \mathcal{N}\} \quad (3.2)$$

Every finitely represented sound spans some finite interval $\{t_{min}, \ldots, t_{max}\}$ rather than the infinite domain $\mathcal{T}_S$, but we may only listen to a finite time-span of sound in a lifetime anyway, so the *discretization* of time is much more important than the limitation to a finite interval. A *sound signal in the* *discrete time domain* with sampling rate S is a complex-valued function σ on $\mathcal{T}_S$. When discussing the domain $\mathcal{T}_S$, we write the members of the domain as $\ldots, t_{-2}, t_{-1}, t_0, t_1, t_2, t_3, \ldots$, where

$$t_k = k/S \tag{3.3}$$

Converting from continuous to discrete. Given a continuous sound signal σ, the obvious and natural choice for a discrete sound signal to represent it is $\mathcal{D}_S(\sigma)$ where

$$\mathcal{D}_S(\sigma)(t) = \sigma(t) \text{ for } t \in \mathcal{T}_S \tag{3.4}$$

$\mathcal{D}_S(\sigma)$ is just σ restricted to the domain $\mathcal{T}_S$. Such a representation is inherently ambiguous—there are an infinite number of different continous sound signals represented by the same discrete sound signal. The confusion resulting from this ambiguity is called *aliasing.*

Aliasing. **—Note: serious problems with functional notation—** Whenever two continuous sound signals σ_1 and σ_2 agree on every point in $\mathcal{T}_S$ ($\sigma_1(t) = \sigma_2(t)$ for all $t \in \mathcal{T}_S$), then they have the same representation $\sigma_d = \mathcal{D}_S(\sigma_1) = \mathcal{D}_S(\sigma_2)$ as a discrete sound signal, and we say that σ_1 and σ_2 are *aliases.* The famous problem of wagon wheels appearing to roll backwards in old movies is an example of aliasing in a sampled video signal. But, when we generate a real physical sound from σ_d, a listener can only hear one of the infinitely many continous sound signals that it might represent. This is undesirable in the case where we discretize one continous sound, and the listener hears a different one. In particular, two continuous helical signals are aliases if and only if their amplitudes and phases are exactly the same, and their frequencies are the same $\pmod{S}$.

$$\begin{gathered}\text{For } R_1, R_2 > 0,\ P_1, P_2 \in [0, 2\pi): \\ \mathcal{D}_S(R_1 e^{\boldsymbol{i}(P_1 + 2\pi F_1 t)}) = \mathcal{D}_S(R_2 e^{\boldsymbol{i}(P_2 + 2\pi F_2 t)}) \text{ for all } t \in \mathcal{T}_S \\ \text{if and only if} \\ R_1 = R_2 \text{ and } P_1 = P_2 \text{ and } F_1 = F_2 \pmod{S}\end{gathered} \tag{3.5}$$

For real-valued signals, there is even more aliasing. Frequencies F_1 and F_2 may be aliased when $F_1 = -F_2 \pmod{S}$ and the signals are out of phase by exactly π (half a period).

$$\text{For } R_1, R_2 > 0,\ P_1, P_2 \in [0, 2\pi),\ F_1, F_2 \neq 0 \pmod{S/2}:$$

$$\begin{gathered} \mathcal{D}_S(R_1 \sin(P_1 + 2\pi F_1 t)) = \mathcal{D}_S(R_2 \sin(P_2 + 2\pi F_2 t)) \text{ for all } t \in \mathcal{T}_S \\ \text{if and only if} \\ R_1 = R_2 \text{ and } P_1 = P_2 \text{ and } F_1 = F_2 \pmod{S} \\ \text{or} \\ R_1 = R_2 \text{ and } P_1 = P_2 + \pi \pmod{2\pi} \text{ and } F_1 = -F_2 \pmod{S} \end{gathered} \tag{3.6}$$

For frequencies that are exact multiples of half the sampling rate, there is even confusion about the amplitude and phase. In the case of odd multiples of half the sampling rate, the samples are all equal in magnitude, and alternating in sign. The amplitude of the samples depends on the phase at which the samples are taken, which is the same for each half wave.

$$\text{For } R_1, R_2 > 0,\ P_1, P_2 \in [0, 2\pi),\ F_1 = F_2 = S/2 \pmod{S}:$$

$$\begin{gathered} \mathcal{D}_S(R_1 \sin(P_1 + 2\pi F_1 t)) = \mathcal{D}_S(R_2 \sin(P_2 + 2\pi F_2 t)) \text{ for all } t \in \mathcal{T}_S \\ \text{if and only if} \end{gathered} \tag{3.7}$$

$$R_1 \sin(P_1) = R_2 \sin(P_2)$$

In the special case where $P_1 = P_2 = 0$, the amplitudes could have any values. For multiples of the sampling rate, all samples have the same value, and again there is a tradeoff between amplitude and phase.

$$\text{For } R_1, R_2 > 0,\ P_1, P_2 \in [0, 2\pi),\ F_1 = F_2 = 0 \pmod{S}:$$

$$\begin{gathered} \mathcal{D}_S(R_1 \sin(P_1 + 2\pi F_1 t)) = \mathcal{D}_S(R_2 \sin(P_2 + 2\pi F_2 t)) \text{ for all } t \in \mathcal{T}_S \\ \text{if and only if} \end{gathered} \tag{3.8}$$

$$R_1 \sin(P_1) = R_2 \sin(P_2)$$

Finally, for completeness, notice that an odd multiple of half the sampling rate aliases with a multiple of the sampling rate precisely when the phases are both 0, so that all samples have the value 0.

For $R_1, R_2 > 0$, $P_1, P_2 \in [0, 2\pi)$, $F_1 = 0 \pmod S$, $F_2 = S/2 \pmod S$:

$$\mathcal{D}_S(R_1 \sin(P_1 + 2\pi F_1 t)) = \mathcal{D}_S(R_2 \sin(P_2 + 2\pi F_2 t)) \text{ for all } t \in \mathcal{T}_S$$

$$\text{if and only if}$$

$$P_1 = P_2 = 0 \tag{3.9}$$

Rendering a discrete sound signal as continuous sound. Having computed a discrete sound signal σ_d, we need to render it through a loudspeaker or similar controllable vibrating device in order to hear the sound. At some point in the rendering process, σ_d is converted to a continuous signal σ_c. It is natural to choose one of the infinitely many continous sound signals that σ_d might represent. In particular, it is natural to create σ_c by *interpolating* values between the ones given by σ_d in such a way as to make the resulting sound signal as smooth as possible, according to some appropriate definition of smoothness. The interpolating is normally done, not by a digital computation, but by the analog machinery, usually electronic, controlled by the computation. In Chapter 4 we find that when the sampling rate is high enough, the precise nature of the interpolation is relatively unimportant. But, the sorts of analog devices commonly used for sound production typically interpolate so that the final result is close to a sum of sinusoidal signals of the lowest possible frequency. So, if $\sigma_d(t) = Re^{i2\pi Ft}$, is a discrete helical signal, it is normally rendered as something very close to the continuous signal $\sigma_c(t) = Re^{i2\Pi(F \bmod S)t}$.

Not all aliases of a helical or sinusoidal signal are helical or sinusoidal themselves. For real-valued signals the frequencies, the frequencies near half the sampling rate ($S/2$) alias to signals that are amplitude modulations of a carrier with frequency S (see Figure ???). In many cases, these amplitude-modulated signals represent the way that a rendered signal is likely to be heard. For complex-valued helical signals, the problem of non-helical aliases seems to be less important, but I know very little about the rendering of complex-valued signals. It is interesting to note that we seem to need two real numbers per period to represent a given frequency, whether those two reals are separate samples or whether they are bundled into a single complex sample. But, we really need only the sign of the imaginary component, along with the entire value of the real component of a complex sample, to resolve the ambiguity between helical frequencies $F_1 = -F_2 \pmod S$, although the full value of both real and complex components is required to get full information about amplitude and phase, and about multiple frequency components.

Causes of and cures for aliasing. Aliasing occurs any time a continuous signal is converted to a discrete signal by sampling. It is natural to think of the case where a physical continuous signal is read by an electronic sampler, but this is not really the most important cause of aliasing. The engineers who design and build samplers are pretty smart, and they have had plenty of time to worry about aliasing and find ways to prevent its harmful consequences. The most troublesome cases of aliasing arise when a continuous mathematical model of a sound signal is converted to a calculation of samples. The original continuous model may only exist in the mind of a person who is designing sound—it need not be present as a data structure in a computer, or in any other realization in an artificial medium. Even when there is an explicit representation of the continuous sound signal available as a data structure, the problem of avoiding aliasing in software is far more complex, due to the variety of conceptual sources for continuous signals. Flexible sound-processing software has largely failed to prevent the introduction of harmful aliasing in sampled signals.

The only cure for the harmful consequences of aliasing is prevention. Once a continuous sound signal has been replaced by a sampled discrete representation, and the continuous signal is no longer available for inspection, there is no way to determine which of the infinitely many possible continuous signals was truly intended. In order to prevent one continuous sound signal σ_1, converted to the discrete signal $\sigma_d = \mathcal{D}_S(\sigma_1)$, from being rendered continuously as some alias σ_2 that sounds quite different, we must sample only signals that will be rendered accurately. With the usual "smooth" rendering techniques, a sampled complex-valued signal produces frequencies in the range $[0, S)$, and a sampled real-valued signal produces frequencies in the range $[0, S/2)$. To avoid harmful aliasing, all higher frequencies must be filtered out from the continuous signal before sampling. For this reason, sampling converters have analog filters that eliminate high frequencies before sampling. Even though digital filters have many advantages, they cannot be applied to the aliasing problem, because they cannot distinguish frequencies that differ by multiples of the sampling rate. To avoid the aliasing of frequencies near $S/2$ with an amplitude-modulated signal (and presumably there are simlar problems for complex-valued signals at frequencies near S), continuous signals should in fact be filtered to an even smaller frequency interval, but it is not clear precisely how much smaller it needs to be.

Discrete signal values other than samples.

- interval values
- average values—equivalent to filtering

3.2 Quantized Vibration State

Even a finite time segment from a discrete sound signal is an infinite object if the sample values are complex or real numbers. In order to get a completely digital representation of a sound signal, we must also approximate the continuous range of real or complex numbers by a discrete subset. Since the consequences of this quantization of the domain of values are largely independent of the consequences of discretizing the time domain, we consider signals from the continuous time domain to a discrete subset of the complex or real numbers.

Discrete sets of complex or real values. A subset $\mathcal{V}$ of the complex numbers is *discrete* if we may draw a circle around each point in $\mathcal{V}$, so that each circle contains only one point in $\mathcal{V}$. If $\mathcal{V}$ contains only real numbers, then it is also a discrete subset of the reals. While the discretization of time seems to make sense only with a constant interval between points, there are a number of different popular ways to quantize the real or complex values $\sigma(t)$. For the domain of real numbers, the two basic ideas are *linear* and *logarithmic* quantization.

Given a real number $Q > 0$, $\mathcal{V}_Q$ represents the linear quantization of the real domain with quantum interval Q.

$$\mathcal{V}_Q \quad = \quad \{k/S : \ k \in \mathcal{N}\} \tag{3.10}$$

Yes, this is mathematically the same thing as the discrete time domain $\mathcal{T}_Q$, but we think of it as having a different physical dimension. Just as in the case of the discrete time domain, a digital representation requires that we limit ourselves to a *finite* interval within $\mathcal{V}_Q$, but it is the use of a discrete subset of values, rather than the limitation to a finite interval, that has the most interesting consequences for sound modelling.

Given two real numbers $B > 1$ and $M > 0$, $\mathcal{L}_{B,M}$ represents the logarithmic quantization of the real domain with base B and minimum nonzero value M.

$$\mathcal{L}_{B,M} = \{0\} \cup \{B^k, -B^k : k \in \mathcal{N} \text{ and } B^k \geq M\} \tag{3.11}$$

Notice that the minimum value M is required to make the domain discrete, even if the 0 value is omitted. $\mathcal{L}_{B,M}$ is an idealized abstraction of several different essentially logarithmic quantizations, such as "mu-law" encoding, but it does not represent them precisely. The crucial quality of logarithmic domains is that the interval between points goes up exponentially with the magnitude of the points: the $\mathcal{L}_{B,M}$s are the mathematically simplest sort of domains with that crucial quality. Floating-point domains are a funny hybrid of linear and logarithmic: they consist of finite segments of different linear domains pieced together so that the progression over larger segments is essentially logarithmic.

The usual way to quantize the complex domain is to pick a quantization of the real domain, and then apply it to the real and imaginary components of complex numbers. So, we can define

$$\mathcal{V}_Q^2 = \{x + \boldsymbol{i}y : x, y \in \mathcal{V}_Q\} \tag{3.12}$$

$$\mathcal{L}_{B,M}^2 = \{x + \boldsymbol{i}y : x, y \in \mathcal{L}_{B,M}\} \tag{3.13}$$

Polar versions, others.

- Quantization (roundoff) noise.
- Nonlinear encodings.
- Recovering complex data from real.
- Why discretization and quantization are studied in such different ways.

3.3 Other Ways to Digitize a Sound Signal

- Delay conversion to discrete sampled representation as long as possible (analogy to bitmapping in graphics).

3.4 Direct Manipulation of Digital Sampled Sound

- Time shifting.
- Amplitude scaling.
- Amplitude clipping.
- Frequency/speed shifting.
- Time stretching by repetition of a "period".
- Changing sampling rate.
- AM enveloping.
- Adding sounds.
- Nonlinear waveshaping.
- Frequency modulation.

Chapter 4

The Frequency Spectrum

In this chapter we investigate the analysis of a sound signal into its components at different frequencies. For the moment, we are concerned only with *steady-state* sound signals. Imagine that you walk into a room in which some sound is in the air. You stay for some length of time much longer than the period of any frequency that you can hear, and then leave the room. Suppose that, while you are in the room, the sound is essentially stable: you do not hear any change in its quality. Roughly speaking, a steady-state sound signal is the infinite extension of such a stable sound into the past and the future—a sound that has always been and will always be qualitatively the same. It is natural and sensible in a mathematical analysis of a stable signal to ignore the fact that it has a beginning and an end.

4.1 Pure Helical, Periodic, and Quasiperiodic Signals

In order to analyze a signal into its components at different frequencies, we need to know what each component is like. We choose the standard helixes $Re^{\boldsymbol{i}(P+2\pi Ft)}$, characterized by amplitude $R > 0$, phase $P \in [0, 2\pi)$, and frequency $F > 0$, as the components for analysis. In principle, there are infinitely many other choices for the basic components: square waves, triangular waves, pulses, etc. We choose the helixes because they are mathematically very simple and suitable for analysis, and they agree quite well (but not perfectly) with the physical qualities of vibrating sound producers, as well as the sound detectors in the ear. In principle, the elliptical helixes with aspect ratios proportional to frequency match our physical analysis of a vibrating spring better than the circular helixes. But, we choose the mathematical simplicity of the circular helixes, and trust an intuitive hunch that they will be satisfactory for our practical needs.

A sound signal σ is *periodic* if there is a positive real number $P > 0$, called the *period* of σ, such that for all times t, $\sigma(t + P) = \sigma(t)$ (actually, the *smallest* such P is the period of σ, and all multiples of P satisfy the same equation). The helical signal $Re^{\mathbf{i}(P+2\pi Ft)}$ is periodic, with period $1/F$. If $\sigma_1, \sigma_2, \ldots$ are all periodic signals, and the ratios of their periods are rational numbers, then $\sum_i \sigma_i$ is also periodic. If the ratios of their periods are not all rational, then the sum is not periodic, but it may still be analyzed into its components at different frequencies—such signals are called *quasiperiodic.* Since physical measurements can never distinguish absolutely between rational and irrational values, it makes sense that we need to study quasiperiodic signals in essentially the same way as the periodic ones.

4.2 The Ear as a Spectral Analyzer

This section repeats Chapter 2, and should be merged in with that chapter.

The part of the mammalian ear that detects sound is called the *cochlea.* The cochlea is a tube (wound into a spiral, but that is not particularly important to us) containing a long sequence of tiny hairs. For our current purposes, each of those hairs is essentially a vibrating spring, tuned to a different frequency. The number of hairs is finite, but it is large enough that we will ignore the discreteness of the set of hairs, and suppose that every possible value in the continuous spectrum of real positive frequencies has a hair tuned to it. In Chapter 6 we find that there are other limitations on the accuracy with which frequency is measured in the ear, besides the finite number of detecting hairs. We also ignore the physiological limits on the range of frequencies. The most natural way to analyze signals is to find their components at all frequencies, and accept the fact that some of those frequencies are indetectible by a given ear.

Intuitively, the *spectrum* of a steady-state sound signal σ is the function ψ mapping each frequency $f > 0$ to a complex number $\psi(f)$ such that $|\psi(f)|$ is the amplitude of the stimulation delivered by σ to an idealized ear, and $\arg(\psi(f))$ is the phase of the stimulation. The spectrum ψ is often called a *signal in the frequency domain,* to complement the description of σ as a signal in the time domain. The perception of steady-state sound signals is explained very well (but not perfectly) in terms of the spectrum of the signal. In fact, most of the perception of steady-state sound appears to depend only on the magnitude $|\psi|$, but in Chapter 6 we find that the perception of changes in the spectrum depends on the relative phases of different components, so it is best to define the spectrum to include phase information.

4.2.1 Perceptual Parameters of Sound

It is tempting to think of the perceptual qualities of a steady-state sound signal that are derived from its spectrum as being analogous to the *color* of an optical signal. Such an analogy probably leads to more misunderstanding than useful insight. Human color perception depends only on three dimensions of the frequency spectrum of light, while the ear can distinguish at least hundreds of frequencies. Considering the number of independent parameters involved, a better analogy would relate each frequency component of a sound signal to a single point in a visual scene. This analogy also breaks down, because the ear relates individual frequency components in ways that are fundamentally and structurally different from the way the eye relates different points in a scene.

First, different frequency components of a steady-state sound signal are often grouped together and perceived as a unit, which I call *a sound.* A steady-state sound signal may contain many individual sounds going on simultaneously. A particular frequency, typically the lowest of those in a sound, may dominate the perceptual identification of frequency in a sound: such a dominant frequency is called the *fundamental* frequency of a sound; all of the component frequencies are called *partials* of the sound. In particular, components with frequencies that are very close to reasonably small integer multiples of a single frequency F $(1F, 2F, 3F, \ldots)$, are often heard as a single sound (a lot of information other than the spectrum may affect this grouping, particularly stereo effects that seem to locate components in space). In this case, the sound is a *(nearly) harmonic* sound, F is the fundamental, and the partial kF is called the *kth harmonic.* In many cases not all of the harmonics are present (for example, the presence of only odd harmonics is very common). Even the fundamental frequency F may not be present in a harmonic sound, but it still dominates perception of the identifying frequency of the sound.

For a harmonic sound, and for many others as well, the ear perceives a quality of highness or lowness called *pitch.* Pitch is essentially determined by the fundamental frequency (even when that frequency is not actually present), but the less harmonic a sound is the more subtle is the determination of a fundamental frequency. Doubling the fundamental frequency typically produces a perception of an additive increment in the pitch—the pitch increment associated with doubled fundamental frequency is called an *octave* in conventional European music. So, the perceived pitch of a sound is roughly the logarithm of the fundamental frequency. For perfectly harmonic sounds this definition works very well; for nearly harmonic sounds the perception of pitch intervals is affected by the inaccuracies in the nearly integer ratios of partials to fundamental frequencies. Notice that the essentially logarithmic relationship of pitch to frequency means that when a sound signal σ containing several sounds is transformed by time-scaling to the signal $[t]\sigma(St)$, the pitch intervals between the sounds stays the same. The tendency of frequency components to cluster into harmonic sounds is consistent with the logarithmic relationship of pitch to frequency, since the partials of different harmonic sounds will interleave in the same way as long as those sounds are separated by the same pitch interval.

While the fundamental frequency of a steady-state sound typically determines its pitch, the relative amplitudes of the frequency components produce a perceived quality of sound that is called *timbre.* Even with all the subtleties in determining perceived pitch, the perception of timbre is orders of magnitude more subtle, and has never been characterized with precision. Some acoustical scholars believe that the word "timbre" is simply a convenient label for those qualities of sound that we cannot describe or analyze satisfactorily, much in the way that "intelligence" sometimes seems to be used as a pleasant label for those aspects of human behavior that we want to admire, but cannot explain. My hunch is that timbre is susceptible to a much better analysis than has been achieved so far, but not necessarily to a complete analysis. Timbre perception certainly has a lot to do with the relative amplitudes of partials, but is also affected crucially by the initiation of a sound, the relation of amplitudes of partials to their absolute frequencies (rather than just the ratios with fundamental frequencies), and probably to a lot of other things that nobody has thought of yet.

The other perceptual quality of sound that has been analyzed fairly well is *loudness.* While pitch is naturally associated with an individual sound, loudness is perceived both for individual sounds and for entire sound signals.

Loudness is determined by the amplitude of sound. Although the mathematically simplest measure of the amplitude of a helical signal $Re^{i2\pi Ft}$ is the multiplier R, perceived loudness is better related to the *power* of the signal, which is proportional to R^2F. In a system based on elliptical helixes, if the aspect ratio of the ellipses grows proportionally to frequency, then product of the lengths of the axes is proportional to power, but this advantage of elliptical helixes does not seem to justify their extra complexity. Just as perceived pitch is logarithmically related to frequency, perceived loudness is logarithmically related to power, so the *decibel* system for measuring loudness associates an additive increase of 6 decibels with a doubling of power. Perceived loudness is also affected by frequency in at least two ways. First, loudness naturally tends to drop off as the frequency of a signal approaches the limits of frequency perception in the ear. Second, when several frequency components combine, the percieved loudness of the combination is different depending on whether the frequencies are close together (lying in what are called *critical bands*) or farther apart. My hunch is that the dependence of perceived loudness on frequencies should not affect the structure of a basic, general-purpose system for sound modelling, but it certainly will become very important to a sound designer for the finer polishing of a sound signal.

4.3 Mathematical Spectral Analysis with the Fourier Transform

The *frequency spectrum* of a sound signal σ is another complex-valued function ψ of one real parameter, but in this case the parameter represents *frequency*, rather than time, and the value of ψ at a frequency f describes the component of σ at frequency f. In particular, $|\psi(f)|$ is the magnitude of the component at frequency f, and $\arg(\psi(f))$ is the phase of that component. If σ is described by a formula in the form

$$\sigma(t) = R_1 e^{i(P_1+2\pi F_1 t)} + R_2 e^{i(P_2+2\pi F_2 t)} + \cdots$$

then it is easy to see that its spectrum ψ should have value $\psi(f) = 0$ for all frequencies f not in the list $F_1, F_2, \ldots$, and it appears that its value at F_i should be $\psi(F_i) = R_i e^{iP_i}$. This is the right basic idea, but for mathematical consistency with more complex spectra, the value $\psi(F_i)$ is interpreted in a slightly more peculiar way, described in Section 4.3.1. But, what if we are given a completely unknown signal σ, and need to characterize its frequency spectrum? Just as, in the physical world, a *prism* is used to analyze a light signal into its different frequency components, in the mathematical world the *Fourier transform* analyzes a mathematically given signal in a similar way.

The essential idea behind the Fourier transform is to perform a kind of *pattern matching* between the given signal σ and each of the standard helical components $e^{i2\pi Ft}$. The strength of the match will determine the value $\psi(F)$ of the spectrum at frequency F. Through a fortunate stroke of mathematics, we need not test all the different amplitudes R and phases P in the form $Re^{i(P+2\pi Ft)}$, because our definition of pattern-matching produces information about all amplitudes and phases as a result of matching against any helix of the right frequency. To understand how the peculiar integral formula that we will introduce as the definition of the Fourier transform represents a kind of pattern matching, we take a detour through simpler sorts of signals.

Suppose that we are given two signals σ_1 and σ_2 over a *discrete and finite* time domain $\{1, 2, \ldots, n\}$, and suppose in addition that both signals take only the values 1 and -1. At any time t, the product $m = \sigma_1(t)\sigma_2(t)$ is $m = 1$ if the signals agree, and $m = -1$ if the signals disagree. So, the sum

$$M(\sigma_1, \sigma_2) = \sum_{t=1}^{n} \sigma_1(t)\sigma_2(t)/n$$

represents the accuracy with which σ_1 matches σ_2—a perfect match gives the value $M\sigma_1, \sigma_2 = 1$; a complete failure to match gives the value $M(\sigma_1, \sigma_2) = -1$, and partial matches give values between -1 and 1. From another point of view, $(M(\sigma_1, \sigma_2)+1)/2$ is the probability that σ_1 and σ_2 agree at a randomly chosen time $t \in \{1, 2, \ldots, n\}$.

Now, generalize σ_1 and σ_2 to real-valued functions on the same discrete and finite time domain. The product $\sigma_1(t)\sigma_2(t)$ is still a very reasonable measure of the extent of agreement between σ_1 and σ_2 at time t. $m > 0$ when σ_1 and σ_2 have the same sign, $m < 0$ when they have opposite sign, and the magnitude $|m|$ indicates the strength of their agreement or opposition in a way that credits agreement with large numbers more than agreement with small numbers. Again, the extent of agreement or disagreement between the entire signals σ_1 and σ_2 may be taken as the average $M(\sigma_1, \sigma_2)$ of the products. Now, however, the average is no longer restricted to the interval $[-1, 1]$, but might be any real number. If σ_1 is a function that we are analyzing, and σ_2 is a pattern that we are comparing it to, then $M(\sigma_1, \sigma_2)$ gives information about the matching of M against all multiples of σ_2 as well, since $M(\sigma_1, S\sigma_2) = SM(\sigma_1, \sigma_2)$.

What if σ_1 and σ_2 are complex-valued functions on the discrete and finite time domain? The product $\sigma_1(t)\sigma_2(t)$ is no longer a good measure of agreement at time t, because of the way that complex multiplication adds the angles of multiplicands. For example, if $\arg(\sigma_1(t)) = \arg(\sigma_2(t)) = \pi/2$, then $\arg(\sigma_1(t)\sigma_2(t)) = \pi$, and π is the angle of the negative real numbers. Similarly, if $\arg(\sigma_1(t)) = \pi/2$ and $\arg(\sigma_2(t)) = 3\pi/2$, then $\arg(\sigma_1(t)\sigma_2(t) = 0$, so the product is a positive real number. These results are the opposite of what we want, since in the first case the signals agree in direction, but the product is negative, and in the second case the signals are opposite in direction but the product is positive. Instead of the product $\sigma_1(t)\sigma_2(t)$, we need the product of one signal with the *conjugate* of the other: $m = \sigma_1(t)\overline{\sigma_2(t)}$. Notice that

$$\arg(\sigma_1(t)\overline{\sigma_2(t)}) = \arg(\sigma_1(t)) + \arg(\overline{\sigma_2(t)}) = \arg(\sigma_1(t)) - \arg(\sigma_2(t)) \pmod{2\pi}$$

So, when $\sigma_1(t)$ and $\sigma_2(t)$ are in the same direction, m is a positive real number (angle 0), when they are in opposite directions m is a negative real number (angle π), and in other cases the angle of m gives the amount of disagreement between the angles of $\sigma_1(t)$ and $\sigma_2(t)$. In the special case where σ_1 and σ_2 are real-valued, $m = \sigma_1(t)\sigma_2(t)$, since each real number is its own conjugate. But, the mathematical symmetry between σ_1 and σ_2 in the real-valued case is lost in the complex-valued case: pattern matching the same two signals in opposite order yields results that are conjugates of one another. Let

$$M(\sigma_1, \sigma_2) = \sum_{t=1}^{n} \sigma_1(t)\overline{\sigma_2(t)}/n$$

Because $M(\sigma_1, \boldsymbol{\alpha}\sigma_2) = \overline{\boldsymbol{\alpha}}M(\sigma_1, \sigma_2)$ for all complex constants $\boldsymbol{\alpha}$, the result of pattern matching σ_1 against a single pattern σ_2 gives information about the matching of σ_1 against all signals that are the same as σ_2 except for amplitude and phase. In particular, matching against $e^{\boldsymbol{i}2\pi Ft}$ gives information about the match with all helixes $Re^{\boldsymbol{i}(P+2\pi FT)}$ of the same frequency.

Now, to define the Fourier transform $\mathcal{F}(\sigma)$ of a signal σ, we need only generalize the ideas above to the infinite continuous time domain, using the integral as the natural continuous analog of the sum. $\mathcal{F}(\sigma)$ is itself a function of a real variable f representing frequency, and $\mathcal{F}(\sigma)(f)$ is the result of pattern matching σ against a standard helix at frequency f ($e^{\boldsymbol{i}2\pi ft}$). Notice that $\overline{e^{i2\pi ft}} = e^{-i2\pi ft}$, and it is conventional to use the $-i2\pi$ form in the formula for the transform.

$$\mathcal{F}(\sigma)(f) = \int_{-\infty}^{\infty} \sigma(t)\overline{e^{i2\pi ft}}dt = \int_{-\infty}^{\infty} \sigma(t)e^{-i2\pi ft}dt \tag{4.1}$$

Although the first integral formula above displays the character of the Fourier transform as the results of pattern-matching a signal against the standard helixes, the second form is the one commonly seen in books and papers about the transform. There are a number of variations in the definition of the Fourier transform—some authors give it as $\int_{-\infty}^{\infty} \sigma(t)e^{-ift}dt$, others as $(2\pi)^{-1/2}\int_{-\infty}^{\infty} \sigma(t)e^{-ift}dt$—but these variations are merely the result of scaling to different units of measurement. The form chosen in Equation 4.1, which measures frequency in cycles per unit time, is the mathematically most convenient one for our purposes.

Other variations on the Fourier transform normalize the magnitude of points in the spectrum in sensible ways. Since the power in the helical signal $Re^{i2\pi ft}$ is proportional to $R^2 f$, rather than to R, we might argue for one of the following forms

$$\mathcal{F}_{freq}(\sigma)(f) = f^{-1}\mathcal{F}(f) \tag{4.2}$$

$$\mathcal{F}_{power}(\sigma)(f) = f^{-1}|\mathcal{F}(f)|^{(1/2)}e^{i\arg(\mathcal{F}(f))} \tag{4.3}$$

Or, if we interpret the real and imaginary components of the signal in such a way that the real component is the derivative of the imaginary component, then it could make sense to match the signal against elliptical helixes:

$$\begin{aligned}\mathcal{F}_{ellipse}(\sigma)(f) &= \int_{-\infty}^{\infty} \sigma(t)((f+1)e^{-i2\pi ft} + (f-1)e^{i2\pi ft})/2dt \\ &= ((f+1)\mathcal{F}(f) + (f-1)\mathcal{F}(-f))/2 \qquad (4.4)\end{aligned}$$

Finally, since perceived pitch is roughly the logarithm of frequency, a sensible variation is

$$\mathcal{F}_{logfreq}(\sigma)(p) = \int_{-\infty}^{\infty} \sigma(t)e^{-i2\pi 2^p t}dt \tag{4.5}$$

$$= \mathcal{F}(\sigma)(2^p) \tag{4.6}$$

The logarithmic frequency scale rules out negative frequencies, which may be an advantage or a disadvantage in different contexts. Since each of these variations is easy to calculate from the conventional transform in Equation 4.1, we stick with that simpler formula.

When using complex-valued signals in the time domain, it seems most sensible to use only positive frequencies—that is, signals σ whose spectra ψ have the property that $\psi(f) = 0$ for $f < 0$. But, the mathematics of the Fourier transform allows negative-frequency components, and it is best to understand the mathematics in its full generality, and then make whatever restrictions seem appropriate in a given application. Real-valued signals always have negative-frequency components. Also, elliptical helixes have negative-frequency components when the spectrum is defined in terms of circular helixes; similarly, circular helixes have negative-frequency components in a spectrum defined in terms of elliptical helixes.

We see in the remainder of this chapter how the Fourier transform $\mathcal{F}(\sigma)$ gives a sensible representation of the spectrum of the steady-state signal σ. It is also important to calculate a sound signal in the time domain from a given frequency spectrum. For this purpose, we need the *inverse Fourier transform.* The Fourier transform is almost self-inverting.

$$\mathcal{F}(\mathcal{F}(\sigma))(t) = \sigma(-t) \tag{4.7}$$

The Fourier transform of the Fourier transform of a signal σ is the same signal played backwards in time. The inverse of the Fourier transform is defined just like the forward transform, except running the helical patterns backwards in time.

$$\mathcal{F}^{-1}(\psi)(t) = \int_{-\infty}^{\infty} \psi(f)\overline{e^{-i2\pi ft}}df = \int_{-\infty}^{\infty} \psi(f)e^{i2\pi ft}df \tag{4.8}$$

$$\mathcal{F}^{-1}(\mathcal{F}(\sigma)) = \sigma \text{ for well-behaved } \sigma \tag{4.9}$$

Equation 4.9 fails for certain weird functions, but all of the sound signals that interest us are well behaved (see [?] for a characterization of the well-behaved functions).

4.3.1 Discrete Spectra

Ironically, for precisely the simple case of a signal that is the sum of helixes at discrete frequencies, the Fourier transform is ill-defined over conventional functions from reals to complex values. Consider the simplest case of a pure helical signal:

$$\mathcal{F}([t]e^{-i2\pi Ft})(f) = \int_{-\infty}^{\infty} e^{i2\pi Ft}e^{-i2\pi ft}dt = \int_{\infty}^{\infty} e^{i2\pi(F-f)t}$$

For $f \neq F$, the product $e^{i2\pi Ft}e^{-i2\pi ft} = e^{i2\pi(F-f)t}$ oscillates, and its integral is 0. But, for $f = F$, we integrate $e^{i2\pi Ft}e^{-i2\pi Ft} = e^{i2\pi(F-F)t} = e^0 = 1$, so the integral is infinite. Similarly, the Fourier transform of a sum of helixes is 0 except at the frequencies of the helixes, where it is infinite. The fact that the integral is infinite tells us that there exists a component at the given frequency, but we lose all information about the magnitude of the component.

Generalized Functions

In order to give more informative values for the Fourier transform of a signal with components at discrete frequencies, we leave the defining formula of Equation 4.1 alone, but we reinterpret the nature of the functions that it may evaluate to, and the rules of calculus for evaluating it. These changes do not affect the well-defined values given by the usual integral calculus, but they provide specific values in some cases where the usual integral calculus is ill defined (consider the analogy to complex arithmetic, which extends real arithmetic to provide a value for $\sqrt{-1}$, which is undefined in the reals). The basic idea is that we want to let $\mathcal{F}([t]e^{i2\pi Ft})$ represent a function whose integral is 1, but which has the value 0 everywhere except at F. There is no such function in the conventional calculus, but just as the real number system may be extended with the new value $\boldsymbol{i}$ with the property $\boldsymbol{i}^2 = -1$, the system of functions from a real number to a complex value may be extended with a new function δ, called the *Dirac* function, or the *impulse* function, with the properties

$$\delta(f) = 0 \text{ for } f \neq 0 \tag{4.10}$$

$$\int_{-\infty}^{\infty} \delta(f)df = 1 \tag{4.11}$$

Intuitively, δ is a function that has value 0 except for a spike at input 0. The spike is infinitesimally narrow, and infinitely high, so that the area inside it is 1. A more careful development defines *generalized functions* to be certain infinite sequences of conventional functions, and considers only the properties of generalized functions in the limit. δ is formally an infinite sequence of narrower and higher spikes, all with area 1. Notice that to place the spike at F instead of 0 we merely shift the input to δ, in the generalized function $[f]\delta(f - F)$. To increase the area under the spike to R, we multiply $R\delta$.

Now, we have a sensible value for the Fourier transform of a helix

$$\mathcal{F}([t]e^{i2\pi Ft})(f) = \delta(f - F) \tag{4.12}$$

For a helix with arbitrary amplitude and phase, the amplitude and phase pass through to the Fourier transform at the helix frequency:

$$\mathcal{F}([t]Re^{\boldsymbol{i}(P+2\pi Ft)})(f) \quad = \quad R\delta(f-F)e^{\boldsymbol{i}P} \tag{4.13}$$

It is often tempting to think that $\delta(0)=1$, but this is not quite right. In order to yield the integral value 1 in Equation 4.11, $\delta(0)$ must be bigger than every conventional real number. But, it is not merely ∞, since it is twice as big as $\delta(0)/2$, half as big as $2\delta(0)$, etc., and it has a well defined angle as a complex number. Just as the new number $\boldsymbol{i}$ was introduced to denote $\sqrt{-1}$, we may use $\Delta=\delta(0)$ when it is convenient to denote the value of the impulse function at its spike.

Also, unlike the discontinuous real-valued function that has value 1 at 0, and value 0 everywhere els, δ should be understood as continuous, and continuously differentiable, even at 0. Intuitively, there is a continuous connection from the 0 values to the infinite value, infinitesimally close to the 0 input. The derivative of δ is an even more peculiar function that has the value 0 at every real input, but goes infinite infinitesimally before input 0, and negatively infinite infinitesimally after input 0.

Now, consider the sum of helixes from the beginning of this section:

$$\sigma(t)=R_1e^{\boldsymbol{i}(P_1+2\pi F_1t)}+R_2e^{\boldsymbol{i}(P_2+2\pi F_2t)}+\cdots$$

The Fourier transform of σ is just a sum of shifted impulse functions:

$$\mathcal{F}(\sigma)(f)=R_1\delta(f-F_1)e^{\boldsymbol{i}P_1}+R_2\delta(f-F_2)e^{\boldsymbol{i}P_2}+\cdots$$

That is, $\mathcal{F}(\sigma)(f)=0$ for $f\neq F_1,F_2,\ldots$, and $\mathcal{F}(\sigma)(F_i)=\Delta R_ie^{\boldsymbol{i}P_i}$. The location of the nonzero values in $\mathcal{F}(\sigma)$ shows the frequencies in the spectrum of σ, the magnitude $|\mathcal{F}(\sigma)(f)/\Delta|$ gives the amplitude of the component at frequency f, and the angle $\arg(\mathcal{F}(\sigma)(f)/\Delta)$ gives its phase.

There is a *lot* more to the theory of generalized functions than I have even hinted here. There are generalized functions that cannot be defined from conventional functions plus δ—their values may not therefore be expressed in terms of Δ. In addition to introducing a particular sort of infinite values, generalized functions allow much of the useful qualities of continuous functions to be associated with mappings to values that are discontinuous in the conventional sense. That is, generalized functions may connect distant values in an infinitesimal range of inputs. So, the *unit step,* or *Heaviside* function H that has value 0 on negative inputs, value $\frac{1}{2}$ at 0, and value 1 on positive inputs, may be regarded as connecting from 0 to 1 in the infinitesimal region around 0, as the intuitive picture of H shows in Figure **??**. As long as the infinitesimally narrow connections work in the most obvious way, I omit them from the definition of a function.

$$H(x) = \begin{cases} 0 & \text{if } x < 0 \\ \frac{1}{2} & \text{if } x = 0 \\ 1 & \text{if } x > 0 \end{cases}$$

The impulse function δ also contains an instantaneous connection from 0 to Δ and back again to 0 in the infinitesimal region around 0.

$$\delta(x) = \begin{cases} 0 & \text{if } x < 0 \\ \Delta & \text{if } x = 0 \\ 0 & \text{if } x > 0 \end{cases}$$

4.3.2 Continuous Spectra and Noise

The putative advantage of the Fourier transform is to connect arbitrary sound signals to their spectra. From Section 4.3.1, we see that the Fourier transform discovers the spectrum of a sum of helixes at discrete frequencies, but we should also be able to use the Fourier transform to analyze signals with components spread continuously across some range of frequencies. Since we do not know in advance what such signals are like in the time domain, it makes sense to define a continuous spectrum, and then apply the inverse Fourier transform to get a sound signal. For example, the spectrum given by the Heaviside impulse function H seems like a sensible representation for the spectrum of a signal with components at all frequencies, all with equal amplitude and phase 0. But, the inverse Fourier transform yields

$$\begin{aligned} \mathcal{F}^{-1}(H)(t) &= \int_{-\infty}^{\infty} H(f) e^{\boldsymbol{i} 2\pi f t} df \\ &= \int_{0}^{\infty} e^{\boldsymbol{i} 2\pi f t} df \\ &= \text{missing steps} \\ &= \delta(t)/2 + \boldsymbol{i}/(2\pi t) \end{aligned}$$

This signal is not perceived as a steady-state sound: rather it is an infinitely sharp sort of click. The problem is that the different frequency components cancel out in a systematic way, instead of presenting a steady sound.

One might object to the spectrum H because it has equal sized components spread over an infinite range of frequencies, and therefore appears to represent an infinite total amount of sound. But, the finite continuous spectrum with frequencies over the range $[1,2]$ has similar problems.

$$\Pi_{1,2}(f) = \begin{cases} 0 & \text{if } f < 1 \\ \frac{1}{2} & \text{if } f = 1 \\ 1 & \text{if } 1 < f < 2 \\ \frac{1}{2} & \text{if } f = 2 \\ 0 & \text{if } f > 2 \end{cases}$$

Intuitively, a finite amount of sound at a discrete frequency F is represented by an infinite value Δx in the spectrum at F. So, $\Pi_{1,2}$ allots only an infinitesimal amount of sound at each frequency in the range $[1,2]$, and the total amount of sound is finite. But, the inverse Fourier transform yields:

$$\begin{aligned} \mathcal{F}^{-1}(\Pi_{1,2})(t) &= \int_{-\infty}^{\infty} \Pi_{1,2}(f) e^{\boldsymbol{i}2\pi ft} df \\ &= \int_{1}^{2} e^{\boldsymbol{i}2\pi ft} df \\ &= (e^{\boldsymbol{i}2\pi 2t} - e^{\boldsymbol{i}2\pi t})/(\boldsymbol{i}2\pi t) \\ &= (e^{\boldsymbol{i}(-\pi/2+2\pi 2t)} - e^{\boldsymbol{i}(-\pi/2+2\pi t)})/(2\pi t) \\ &= \text{missing steps} \\ &= \sin(\pi t) e^{-\boldsymbol{i}2\pi(3/2)t}/(\pi t) \end{aligned}$$

That is, $\mathcal{F}^{-1}(\Pi_{1,2})$ looks like a helix at frequency $3/2$, whose amplitude oscillates and dies out by multiplication with $\sin(\pi t)/(\pi t)$. Such a signal is not perceived as a steady-state sound, but rather as a sound that rises from silence to a maximum loudness at 0 and then dies out again. Again, the continuous spread of frequency components with synchronized phases produces a strange sort of systematic cancellation.

Random Functions

In order to give the perception of steady-state sound with components spread over a continuous range of frequencies, we need to use sound signals whose values vary randomly at each point in the time domain. Let

$$Z(x) = e^{-x^2/2}/\sqrt{2\pi} \tag{4.14}$$

and notice that

$$\int_{-\infty}^{\infty} Z(x)dx \quad = \quad 1 \tag{4.15}$$

Z is called the *Gaussian function,* or the *normal probability density function,* and it has the shape of a bell (see Figure **??**). Z is particularly convenient mathematically because many probability calculations starting with Gaussian functions produce scaled versions of Gaussian functions at the end. In particular, the sum of an infinite sequence of independent random variables with a Gaussian density is itself a random variable with a Gaussian density.

Let r be a real-valued function, such that for each input t, $r(t)$ is a random value with Gaussian density, independent of all other values of r. And, let p be another real-valued function such that for each t, $p(t)$ is a uniformly-distributed value in the range $[0, 2\pi)$, independent of all other values of p and all values of r. Now, define the complex-valued sound signal in the time domain ρ by

$$\rho \quad = \quad re^{ip} \tag{4.16}$$

Each angle $\arg(\rho(t))$ is a random value uniformly distributed over the range $[0, 2\pi)$, each $|\rho(t)|$ is a random positive real value distributed with density $[x]2Z(x)$, and all of these random values are independent of one another. Unlike the generalized functions, random functions such as ρ are legitimate functions according to the conventional mathematical definition, but they are unusual in that they are discontinuous everywhere. Imagine the graph of ρ as an infinite bristle brush, of the sort where bristles stick out in a full circle about the handle, which is the t axis of the graph. **(As generalized functions, do the ρs connect each value back to 0, or to "adjacent" values? I'm not sure at the moment. Later: certainly not to 0, but the exact sense in which adjacent values connect is tricky. The key is to get a sensible derivative.)** Although random functions are mathematical functions in the conventional sense, our definitions of random functions, and our notation for them, are somewhat odd. When we write ρ, we are referring to any one of the infinitely many functions with the random properties described above, but within one discussion, all instances of the symbol ρ refer to the same random function. If we attach different subscripts to ρ, such as $\rho_1, \rho_2, \ldots$, then we are referring to different functions with the same random properties (independent?).

The marvellous thing about the signal ρ is that its Fourier transform is another function with the same random properties:

$$\mathcal{F}(\rho_1)(f) \quad = \quad \int_{-\infty}^{\infty} \rho_1(t)e^{-i2\pi ft}dt = \rho_2(f) \text{ with probability 1} \tag{4.17}$$

The reason for this similarity is that, for each f, $[t]\rho_1(t)e^{-i2\pi ft} = \rho_{1,f}(t)$ is another Gaussian-random function, and then the integral gives yet another random function with Gaussian density. The independence of each pair of values $\mathcal{F}(\rho_1)(F_1)$, $\mathcal{F}(\rho_1)(F_2)$ for $F_1 \neq F_2$ holds because the functions $e^{-i2\pi F_1 t}$ and $e^{-i2\pi F_2 t}$ are linearly independent. So, the spectrum of a Gaussian-random sound signal ρ_1 is given by another Gaussian-random function ρ_2. ρ_1 is perceived as a steady-state sound, and ρ_2 represents essentially a uniform spread of sound over all frequencies. The randomization of the precise values of the spectrum avoids the systematic cancellation in the signals with spectra H and $\Pi_{1,2}$.

The spectrum ρ_2 above has the same qualitative behavior over negative frequencies as over positive frequencies. In order to have a steady-state sound signal with only positive-frequency components, define the spectrum by

$$\rho^+(f) \quad = \quad \begin{cases} 0 & \text{if } f < 0 \\ \rho(f) & \text{if } f \geq 0 \end{cases} \tag{4.18}$$

The inverse Fourier transform of ρ^+ is a random function with the same density as ρ.

$$\mathcal{F}^{-1}(\rho_1^+)(t) \quad = \quad \int_{-\infty}^{\infty} \rho_1^+(t)e^{-i2\pi ft}dt = \rho^{\oplus}(t) = \rho_2(t) \text{ with probability 1} \tag{4.19}$$

It is strange that ρ and ρ^+ seem to have the same inverse Fourier transform. In fact, they do not. They both have inverse transforms that are random functions of the same density, but $\mathcal{F}^{-1}(\rho^+)$ inhabits a small, probability-0, subspace of the space of all such functions. $\rho^{\oplus}$ denotes a function in this subspace.

$$\rho^{\oplus} \quad = \quad \mathcal{F}^{-1}(\rho^+) \tag{4.20}$$

Every $\rho^{\oplus}$ is also a ρ, but a randomly chosen ρ has probability 0 of being a $\rho^{\oplus}$.

Classical probability theory based on measure is not really the right foundation for this discussion. Each particular signal either does or does not sound like white noise. What determines the sound is not the process that produced the signal, but specific statistics that must agree with the statistics of a random process.

Steady-state signals with spectra, such as ρ and ρ^+, that distribute sound evenly across all frequencies or all positive frequencies are called *white noise,* by analogy to the perception of light with an even mixture of all fequencies as the color white. Lots of other randomized functions behave as white noise. In particular, either the phase or the magnitude of the signal may be constant or deterministic as long as the other one is appropriately randomized. For example, real-valued functions with random values in uniform or Gaussian density, and complex-valued functions with constand magnitude and random angle, both generate forms of white noise.

4.3.3 A Calculus of Fourier Transforms

In order to apply the Fourier transform to generate useful insight, we need some rules for deriving Fourier transforms of interesting functions. In principle, we may use the definition of Fourier transform in Equation 4.1, and apply the rules of the integral calculus. In fact, that is far too cumbersome, and we need to manipulate functions and their transforms at a much higher level. The best thing is to forget the meaningful application of the transform for a while, and just systematically absorb some useful functions and operators on functions and the rules for their interactions with the transform.

Basic Functions and Their Transforms

Many of the basic functions turn out to be useful in both the time and frequency domains, so I define them in terms of a generic variable x. For a systematic presentation, I repeat some functions that we are already familiar with.

$$1(x) = 1 \tag{4.21}$$

$$\xi(x) = e^{\boldsymbol{i}2\pi x} \tag{4.22}$$

$$\delta(x) = \begin{cases} 0 & \text{if } x < 0 \\ \Delta & \text{if } x = 0 \\ 0 & \text{if } x > 0 \end{cases} \tag{4.23}$$

$$
\begin{aligned}
\text{III}(x) &= \sum_{i=-\infty}^{\infty} \delta(x-i) & (4.24)\\
\kappa(x) &= \delta(x)/2 + \boldsymbol{i}/(2\pi x) & (4.25)\\
H(x) &= \begin{cases} 0 & \text{if } x < 0 \\ \frac{1}{2} & \text{if } x = 0 \\ 1 & \text{if } x > 0 \end{cases} & (4.26)\\
Z(x) &= e^{-x^2/2}/\sqrt{2\pi} & (4.27)\\
\Pi(x) &= \begin{cases} 0 & \text{if } x < -1/2 \\ \frac{1}{2} & \text{if } x = -\frac{1}{2} \\ 1 & \text{if } -1/2 < x < 1/2 \\ \frac{1}{2} & \text{if } x = -\frac{1}{2} \\ 0 & \text{if } x > 1/2 \end{cases} & (4.28)\\
\text{sinc}(x) &= \sin(\pi x)/(\pi x) & (4.29)\\
\rho &= re^{\boldsymbol{i}p} \text{ for Gaussian random } r\text{, uniform random } p & (4.30)\\
\rho^{+} &= H\rho & (4.31)\\
\rho^{\oplus} &= \mathcal{F}^{-1}(\rho^{+}) & (4.32)
\end{aligned}
$$

Now, the transforms of the basic functions:

$$
\begin{aligned}
\mathcal{F}(1) &= \delta & (4.33)\\
\mathcal{F}(\xi) &= [s]\delta(s-1) & (4.34)\\
\mathcal{F}(\delta) &= 1 & (4.35)\\
\mathcal{F}(\text{III}) &= \text{III} & (4.36)\\
\mathcal{F}(\kappa) &= H & (4.37)\\
\mathcal{F}(H) &= \overline{\kappa} & (4.38)\\
\mathcal{F}(Z) &= Z & (4.39)\\
\mathcal{F}(\Pi) &= \text{sinc} & (4.40)\\
\mathcal{F}(\text{sinc}) &= \Pi & (4.41)\\
\mathcal{F}(\rho_1) &= \rho_2 & (4.42)\\
\mathcal{F}(\rho^{+}) &= \rho^{\oplus} & (4.43)\\
\mathcal{F}(\rho^{\oplus}) &= \rho^{+} & (4.44)
\end{aligned}
$$

Functional Operators

Functions may be combined by performing a given arithmetic operation on their corresponding values. Cross correllation ($\star$) and convolution ($*$) operate on the entirety of two functions with an integral formula to produce a new function.

$$(\alpha+\beta)(x) = \alpha(x)+\beta(x) \tag{4.45}$$

$$(\alpha-\beta)(x) = \alpha(x)-\beta(x) \tag{4.46}$$

$$(\alpha\beta)(x) = \alpha(x)\beta(x) \tag{4.47}$$

$$(\alpha/\beta)(x) = \alpha(x)/\beta(x) \tag{4.48}$$

$$(\beta^{\alpha})(x) = \beta(x)^{\alpha(x)} \tag{4.49}$$

$$\overline{\alpha}(x) = \overline{\alpha(x)} \tag{4.50}$$

$$(\alpha\star\beta)(x) = \int_{-\infty}^{\infty}\overline{\alpha(y-x)}\beta(y)dy = \int_{-\infty}^{\infty}\overline{\alpha(y)}\beta(x+y)dy \tag{4.51}$$

$$(\alpha*\beta)(x) = \int_{-\infty}^{\infty}\alpha(y)\beta(x-y)dy = \int_{-\infty}^{\infty}\alpha(y+x)\beta(-y)dy \tag{4.52}$$

Now, consider some important properties of, and relations between, these functional operators.

$$\delta*\alpha = \alpha \tag{4.53}$$

$$\alpha*\beta = \beta*\alpha \tag{4.54}$$

$$(\alpha*\beta)*\gamma = \alpha*(\beta*\gamma) \tag{4.55}$$

$$\alpha*(\beta+\gamma) = (\alpha*\beta)+(\alpha*\gamma) \tag{4.56}$$

$$\delta\star\alpha = \alpha \tag{4.57}$$

$$\alpha\star\beta = \overline{\beta}\star([x]\alpha(-x)) \tag{4.58}$$

$$(\alpha\star\beta)\star\gamma = \alpha\star(\beta\star\gamma) \tag{4.59}$$

$$\alpha\star(\beta+\gamma) = (\alpha\star\beta)+(\alpha\star\gamma) \tag{4.60}$$

$$\alpha\star\beta = ([x]\overline{\alpha}(-x))*\beta = \overline{\alpha}*([x]\beta(-x)) \tag{4.61}$$

$$\alpha*\beta = ([x]\overline{\alpha}(-x))\star\beta = \overline{\alpha}\star([x]\beta(-x)) \tag{4.62}$$

$$(\alpha*\beta)' = \alpha'*\beta \tag{4.63}$$

$$(\alpha*\beta)' = \alpha*\beta' \tag{4.64}$$

$$\delta\alpha = (\alpha(0))\delta \tag{4.65}$$

$$([x]\delta(x-A))*\alpha = [x]\alpha(x-A) \tag{4.66}$$

$$([x]\delta(x-A))\star\alpha = [x]\alpha(x-A) \tag{4.67}$$

$$([x]\delta(x-A))\alpha = \alpha(A)\delta \tag{4.68}$$

To calculate Fourier transforms, use the following rules that show how the transform and other functional operators interact.

$$\mathcal{F}(\alpha+\beta) = \mathcal{F}(\alpha)+\mathcal{F}(\beta) \tag{4.69}$$

$$\mathcal{F}(\overline{\alpha}) = [f]\overline{\mathcal{F}(\alpha)(-f)} \tag{4.70}$$

$$\mathcal{F}(\boldsymbol{\alpha}\beta) = \boldsymbol{\alpha}\mathcal{F}(\beta) \text{ for constant } \boldsymbol{\alpha} \tag{4.71}$$

$$\mathcal{F}([t]\beta(t-A)) = \mathcal{F}(\beta)([f]\xi(Af)) = \mathcal{F}(\beta)([f]e^{\boldsymbol{i}2\pi Af}) \text{ for real constant } A \tag{4.72}$$

$$\mathcal{F}([t]\beta(At)) = ([f]\mathcal{F}(\beta)(f/A))/|A| \text{ for real constant } A \tag{4.73}$$

$$\mathcal{F}([t]\beta(-t)) = [f]\mathcal{F}(\beta)(-f) \tag{4.74}$$

$$\mathcal{F}(\alpha*\beta) = \mathcal{F}(\alpha)\mathcal{F}(\beta) \tag{4.75}$$

$$\mathcal{F}(\alpha\beta) = \mathcal{F}(\alpha)*\mathcal{F}(\beta) \tag{4.76}$$

$$\mathcal{F}(\alpha\star\beta) = -\overline{[f]\mathcal{F}(\alpha)(-f)}\mathcal{F}(\beta) \tag{4.77}$$

$$\mathcal{F}(\alpha([t]e^{\boldsymbol{i}2\pi Ft})) = [f]\mathcal{F}(\alpha)(f-F) \tag{4.78}$$

$$\mathcal{F}(\alpha\star\alpha) = |\mathcal{F}(\alpha)|^2 \tag{4.79}$$

$$\mathcal{F}(\alpha') = [f]\boldsymbol{i}2\pi f\mathcal{F}(\alpha)(f) \tag{4.80}$$

4.4 The Meaning of Convolution and Multiplication

Table 4.1 shows how various signal-processing operations can be expressed as various multiplications and convolutions in the time and frequency domains. In each case, the subscript t indicates a sound signal in the time domain, and the subscript f indicates its Fourier transform in the frequency domain. In the first five cases, the correspondence of the multiplicative formula in one domain to the signal-processing operation is intuitively clear—the appropriateness of the convolution formula follows from Equations 4.75 and 4.76. In the last case (time-shifting), the convolution formula in the time domain is intuitively clear, and the multiplication formula in the frequency domain is derived from it.

The Fourier transform provides a number of useful insights into the nature of steady-state sound. Amplitude modulation is seen to introduce frequency components at the sums and differences of components in the carrier and the modulator, since it translates from multiplication in the time domain to convolution in the frequency domain. Filtering to boost certain frequencies in relation to others clearly corresponds to multiplication in the frequency domain, so it may be implemented by convolution in the time domain. Since sampling in the time domain is essentially multiplication by III, it has the effect of convolution with III in the frequency domain, and the replication of frequency components at regular intervals resulting from convolution with III is precisely aliasing. Finally, notice how Equation 4.79 tells us that the relative magnitudes of components in the frequency spectrum $\mathcal{F}(\sigma)$ depend only on the statistical correllation of σ with the time-shifted versions $[t]\sigma(t-A)$. This autocorrellation property is what allows random functions to generate continuous spreads of frequencies, at the cost of randomizing the phases of the frequency components.

Operation	**Time Domain**	**Frequency Domain**
Amplitude modulation	$A_t\beta_t$	$A_f * \beta_f$
Amplitude/phase modulation	$\alpha_t\beta_t$	$\alpha_f * \beta_f$
Sampling	$\mathrm{III}\beta_t$	$\mathrm{III} * \beta_f$
Filtering	$A_t * \beta_t$	$A_f\beta_f$
Filtering/phase-shifting	$\alpha_t * \beta_t$	$\alpha_f\beta_f$
Time-shifting	$([t]\delta(t-A)) * \beta_t$	$([f]e^{\mathbf{i}2\pi Af})\beta_f$

Table 4.1: The meanings of convolutions and multiplications of signals.

Chapter 5

Additive Spectral Synthesis

5.1 Steady-State Sound

5.2 Amplitude Modulation, Enveloping

5.2.1 Enveloping a Multifrequency Sound

5.2.2 Enveloping Individual Spectral Components

5.3 Frequency Modulation

Chapter 6

Time-Varying Spectral Analysis

- Dual use of time—sonic time vs. variation time.

6.1 Shortcomings of the Fourier Transform

- Implications for sampling rate.

6.2 Time-Varying Spectral Analysis with the Continuous Wavelet Transform

- Windowing.
- Window shape.
- Constant Q.
- Causality.
- Phase vocorder.

Chapter 7

Synthesis by Resonant Modes

- Shortcomings of additive synthesis w.r.t. simplicity, transformability.

7.1 Filters

In general, a filter is just something that converts an input signal to an output signal: that is, a function $\mathcal{T}$ from signals to signals. Most of filter theory deals with restricted classes of filters with qualities that are physically reasonable and/or convenient for analysis. The important restricted classes of filters include:

Time-invariant. Although we express time as a real-valued parameter t, it is usually unrealistic to attach any special significance to time $t = 0$, nor to any other particular value of t. Normally, only the time *differences* between events have perceptual significance in a sound signal. A *time-invariant* filter is one whose output at time t does not depend on the specific value of t, but only on the time difference between t and various events in the input signal. The output of a time-invariant filter normally varies over time, but that variation depends entirely on input variation, not on the value of t. That is, the filter has no internal clock, and no access to information other than the input signal. Technically, this means that if the input is shifted in time, the output gets shifted in exactly the same way.

Strictly speaking, if we change the behavior of a filter by manipulating some knobs, sliders, or other controls, then the filter is not time-invariant, since its response to an input signal depends on which parts of the signal arrive before our adjustmets and whcih parts arrive afterwards. For example, the tuning circuitry on a radio is not a time-invariant filter when we turn the dial. But, we normally understand such an adjustable filter as a parameterized sequence of time-invariant filters corresponding to the different control settings. This view works well enough as long as we are willing to ignore transient behavior close to the times when the controls change.

Causal. A *causal* filter is one whose output at time t depends only on the input at times $\leq t$. That is, present output depends on past and present, but not future, input. In the strictest sense, only causal filters are physically realizable when both input and output are signals in the time domain. But, it may be convenient to view a causal filter as an approximation to a mathematically simpler noncausal filter. Also, it is sometimes convenient to callibrate input and output times to different 0 points, in which case a filter that is physically causal may be modelled by one that appears to look into the future for a limited interval. Causality is largely irrelevant when we filter signals that are arranged in space rather than time. Noncausal filtering of sound signals in real time is impossible, but noncausal filters are perfectly feasible, and quite useful, for image processing and off-line sound processing.

Finite Impulse Response. A filter has *finite impulse response* if, whenever the input becomes 0 and remains 0 forever, the output eventually becomes and remains 0.

Memoryless. A filter is *memoryless* if the output at time t depends only on the input and its derivatives at time t—not on the future nor past behavior of the input signal.

Pointwise. A filter is *pointwise* if the output at time t depends only on the input *value* at time t—no memory and no derivatives.

Stable. A *stable* filter is one whose output cannot run off to infinity without some sort of infinite input stimulus. Physically realizable filters are always stable in some sense, but it may be convenient to model a situation where the filter gets destroyed as an instability. There are many possible variants of stability depending on the precise physical interpretation of signals, and on the limitations imposed on inputs. In these notes, I require that the output of a stable filter cannot go off to infinity unless the input does so. Physically realizable filters often satisfy stronger restrictions, such as conservation of energy.

Linear. A *linear* filter is one for which sums and scalings of input produce corresponding sums and scalings of outputs. Linear filters are *much* easier to analyze than general nonlinear filters, since we may break the input up into helical components, filter each component, and reconstruct the output from the filtered components. We have a thorough mathematical explanation of linear filters, and very little mathematical information about nonlinear ones. Physically realizable filters are never precisely linear, but they are often very nearly linear within reasonable limits on the input signal. In such cases, it is *very* helpful to consider the linear approximation to a filter, instead of a more exact but less analyzable form.

Each restrictions on filters may be expressed precisely as a mathematical property of a filter $\mathcal{T}$:

Time-invariant: $\mathcal{T}([s]\sigma(s+A))(t) = \mathcal{T}(\sigma)(t+A)$ for all signals σ, real constants A.

Causal: $\mathcal{T}(\sigma)(t) = \mathcal{T}(\sigma[s]H(t-s))(t)$.

Finite Impulse Response: For all signals σ, real constants A, there is a real constant B such that $\mathcal{T}([s](H(s-A)) * \sigma)(t) = 0$ for all $t > B$.

Memoryless: $\mathcal{T}(\sigma)(t) = f(\sigma(t), \frac{d}{dt}\sigma(t), \frac{d^2}{dt^2}\sigma(t), \ldots)$ for some function f.

Pointwise: $\mathcal{T}(\sigma)(t) = f(\sigma(t))$ for some function f.

Stable: For all signals σ, if there is a real constant A such that $|\sigma(t)| < A$ for all t, then there is another real constant B such that $|\mathcal{T}(\sigma)(t)| < B$ for all t.

Linear: $\mathcal{T}(\gamma\sigma + \boldsymbol{\delta}\beta) = \gamma\mathcal{T}(\sigma) + \boldsymbol{\delta}\mathcal{T}(\beta)$ for all time signals σ and β, complex constants γ, $\boldsymbol{\delta}$.

In this section, I review the theory of linear, time-invariant filters. Nonlinear filters are very important in sound modelling, but I cannot find a general theory for them.

7.1.1 Unit Resonances and Antiresonances

Just as sound signals may be decomposed into simple helices, linear time-invariant filters may be composed into simple resonances, which boost signals near a certain frequency, and antiresonances, which suppress signals near a certain frequency.

Unit Resonances

The *unit resonance* filter is essentially the ideal spring system of Section 1.1.1, with an idealized frictional force opposing the spring tension. The tendency of a spring to vibrate at a particular frequency produces the desired resonance—in fact a frictionless spring, once moved away from 0, vibrates forever at its natural frequency. Friction is needed to keep the system stable when it is stimulated by an input signal. Because of the form of analysis described in Section 7.1.3, this type of basic filter is also called a *single-pole filter*. Since current behavior depends only on present and past behavior, it is causal. But, it is neither memoryless, nor finite-impulse response.

Recall Equation 1.51 from Section 1.2.2, the complex-number form of the differential equation for a rotor:

$$\rho' = \boldsymbol{i} A \rho$$

The real constant A determines the frequency of the rotor; $\boldsymbol{i}A\rho$ is a vector perpendicular to ρ with length scaled by A, so $\boldsymbol{i}A\rho$ represents a motion around a circle. Idealized friction causes the rotor to decay toward 0, by adding a displacement in the opposite direction to ρ, and of proportional length. So, the differential equation for a decaying rotor is

$$\rho' = (B + \boldsymbol{i}A)\rho \tag{7.1}$$

with $B < 0$. If $B = 0$, then this reduces to the frictionless rotor. For $B > 0$, we get a rotor with antifriction, which is rather explosive. We might as well take full advantage of complex arithmetic, and rewrite Equation 7.1 as

$$\rho' = \boldsymbol{\alpha}\rho \tag{7.2}$$

where $\Re(\boldsymbol{\alpha}) < 0$ and $\Im(\boldsymbol{\alpha}) > 0$.

With initial condition $\rho(0) = 1$, the solution to Equation 7.2 is $\rho(t) = e^{\boldsymbol{\alpha}t}$, just as if $\boldsymbol{\alpha}$ were real. Using Euler's formula (Equation 1.61), this is

$$\rho(t) = e^{\boldsymbol{\alpha}t} = e^{\Re(\boldsymbol{\alpha})t}e^{\boldsymbol{i}\Im(\boldsymbol{\alpha})t} = e^{\Re(\boldsymbol{\alpha})t}(\cos(\Im(\boldsymbol{\alpha})t) + \boldsymbol{i}\sin(\Im(\boldsymbol{\alpha})t))$$

So, the behavior of a decaying rotor, when released at $\rho(0) = 1$, is to follow a spiral that turns at $\Im(\boldsymbol{\alpha})/(2\pi)$ cycles per unit time, and decays exponentially with amplitude $e^{\Re(\boldsymbol{\alpha})t}$ at time t.

To use a decaying rotor as a unit resonance, let the input to the filter contribute an additional component to the rotor derivative, and take the output from the rotor state. Let $\mathcal{R}_{\boldsymbol{\alpha}}$ be the filter constructed from the decaying rotor with constant $\boldsymbol{\alpha}$. Then $\mathcal{R}_{\boldsymbol{\alpha}}(\sigma) = \rho$, where

$$\rho' \;=\; \sigma + \boldsymbol{\alpha}\rho \qquad (7.3)$$

This filter resonates to frequencies near $\Im(\boldsymbol{\alpha})/(2\pi)$, and the strength of the resonance decreases as $\Re(\boldsymbol{\alpha})$ decreases.

First, consider the behavior of $\mathcal{R}_{\boldsymbol{\alpha}}$ when the input is a unit helix at the resonant frequency, that is $\mathcal{R}_{\boldsymbol{\alpha}}(\sigma)$ where $\sigma(t) = e^{\boldsymbol{i}\Im(\boldsymbol{\alpha})t}$. Assume that the input continues for all time, so that the rotor configuration reaches some sort of steady state. In that steady state, the rotor state and the input have exactly the same angle (they are precisely in phase), and the input cancels the decay component of the rotor equation, so that the rotor state is also a helix at frequency $\Im(\boldsymbol{\alpha})/(2\pi)$, but with possibly a different amplitude than the input. Look at Figure (see Maple manuscript) for the relationship between the input and rotor state.

The picture in Figure (Maple manuscript) is probably the best tool for understanding the essential quality of the rotor filter with input at its resonant frequency. But, if you really want to work through the mathematics (which is elementary, but icky), here it is. Since ρ has the same frequency and phase as the input, but some unknown amplitude A, we may write $\rho(t) = Ae^{\boldsymbol{i}\Im(\boldsymbol{\alpha})t}$, which implies that

$$A\boldsymbol{i}\Im(\boldsymbol{\alpha})e^{\boldsymbol{i}\Im(\boldsymbol{\alpha})t} = \rho' = \sigma + \boldsymbol{\alpha}\rho = e^{\boldsymbol{i}\Im(\boldsymbol{\alpha})t} + \boldsymbol{\alpha}Ae^{\boldsymbol{i}\Im(\boldsymbol{\alpha})t} = (1 + A\boldsymbol{\alpha})e^{\boldsymbol{i}\Im(\boldsymbol{\alpha})t}$$

Comparing the first and last terms above, we find that $\boldsymbol{i}A\Im(\boldsymbol{\alpha}) = (1 + A\boldsymbol{\alpha})$, so

$$0 = \Re(\boldsymbol{i}A\Im(\boldsymbol{\alpha})) = \Re(1 + A\boldsymbol{\alpha}) = 1 + A\Re(\boldsymbol{\alpha})$$

so $A = -1/\Re(\boldsymbol{\alpha})$. A is undefined when $\Re(\boldsymbol{\alpha}) = 0$. In fact, the only way to have a steady state with $\Re(\boldsymbol{\alpha}) = 0$ is to make $\sigma = 0$. The derivation yields a steady state when $\Re(\boldsymbol{\alpha}) > 0$. This steady state is correct in a sense. It corresponds to input exactly $\frac{1}{2}$ a cycle out of phase with the rotor state. But, this configuration is highly unstable: the least deviation from opposite phase will send the rotor state off to infinity. By contrast, when $\Re(\boldsymbol{\alpha}) < 0$ every initial state heads asymptotically toward the steady state derived above.

Now, consider an input unit helix at an arbitrary frequency, $\sigma(t) = e^{\boldsymbol{i}Bt}$. The picture of the steady state is a bit more complicated than before. The frequency of the output must still be the same as the frequency of the input, so that the rotor state and input remain at the same angle. If the input frequency is greater than the resonant frequency, then the phase of the input advances ahead of the phase of the rotor state, so that part of the input contributes to moving the rotor at a higher frequency than the resonance. If the input frequency is less than the resonant frequency, then the phase of the input falls behind the phase of the rotor state, so that part of the input retards the rotor. As a result, there is less of the input available to cancel the decay, and the output has a smaller amplitude than it does at the resonant frequency. Look at Figures (Maple manuscript) to see the relationship between the input and rotor state when the input frequency is greater (respectively, less) than the resonant frequency.

If you really want to see the mathematical derivation, let the output be $\rho(t) = Ae^{\boldsymbol{i}(Bt+C)}$. Then

$$A\boldsymbol{i}Be^{\boldsymbol{i}(Bt+C)} = \rho' = \sigma + \boldsymbol{\alpha}\rho = e^{\boldsymbol{i}Bt} + \boldsymbol{\alpha}Ae^{\boldsymbol{i}(Bt+C)} = (e^{-\boldsymbol{i}C} + A\boldsymbol{\alpha})e^{\boldsymbol{i}(Bt+C)}$$

Comparing the first and last terms, $\boldsymbol{i}AB = e^{-\boldsymbol{i}C} + A\boldsymbol{\alpha}$, so

$$0 = \Re(\boldsymbol{i}AB) = \Re(e^{-siC} + A\boldsymbol{\alpha}) == -\cos(C) + A\Re(\boldsymbol{\alpha})$$

and

$$AB = \Im(\boldsymbol{i}AB) = \Im(e^{\boldsymbol{i}C} + A\boldsymbol{\alpha}) = -\sin(C) + A\Im(\boldsymbol{\alpha})$$

From the real part, $\cos(C) = A\Re(\boldsymbol{\alpha})$. From the imaginary part, $\sin(C) = A(\Im(\boldsymbol{\alpha}) - B)$, so $C = \arctan((\Im(\boldsymbol{\alpha}) - B)/\Re(\boldsymbol{\alpha}))$ and $A = \cos(C)/\Re(\boldsymbol{\alpha}) = \sin(C)/(\Im(\boldsymbol{\alpha}) - B) = 1/|\boldsymbol{\alpha} - \boldsymbol{i}B|$. It is easy to see that A is maximum, and C is 0, when $B = \Im(\boldsymbol{\alpha})$. Notice also that the amplitude A falls off more rapidly the closer that $\Re(\boldsymbol{\alpha})$ gets to 0. When $B > \Im(\boldsymbol{\alpha})$, then $C < 0$, which means that the input phase is ahead of the output. When $B < \Im(\boldsymbol{\alpha})$, then $C > 0$, so the input phase is behind the output. When $\Re(\boldsymbol{\alpha}) = 0$, and when $\Re(\boldsymbol{\alpha}) > 0$, behavior is analogous to the case where input frequency equals resonant frequency: we get no steady state unless the input is 0, and a steady but unstable case with input nearly $\frac{1}{2}$ cycle out of phase with rotor state, respectively.

One mathematical complication has been swept under the rug. If you didn't notice it, don't worry about. Just in case you did: the steady-state output derived above is not the only solution satisfying the differential equation for the unit resonance. There are an infinite number of solutions. We may select any complex number value for $\mathcal{R}_{\boldsymbol{\alpha}}(\sigma)(0)$, and there is a unique solution with that 0 value. The solution that makes physical sense is the only one in which $\lim_{t \to -\infty} \mathcal{R}_{\boldsymbol{\alpha}}(\sigma)(t)$ is bounded. Intuitively, this corresponds to setting the filter state to 0 at time $-\infty$.

Even though you skipped the mathematical derivations, you should review the following crucial qualitative observations.

- $\mathcal{R}_{\boldsymbol{\alpha}}$ has a resonant frequency at $\Im(\boldsymbol{\alpha})/(2\pi)$.
- Set to some nonzero initial state, and left to itself, the decaying rotor with constant $\boldsymbol{\alpha}$ spins at the resonant frequency, and its amplitude decays in the form of $e^{i\Re(\boldsymbol{\alpha})t}$.
- When $\Re(\boldsymbol{\alpha}) < 0$, the rotor behavior is stable.
- When $\Re(\boldsymbol{\alpha}) = 0$ and the resonant frequency is not 0, the rotor behavior is unstable for all nonzero helical inputs. On zero input, it rotates at the resonant frequency without decay. The peculiar special case where $\boldsymbol{\alpha} = 0$ is stable for helical inputs with nonzero frequency.
- When $\Re(\boldsymbol{\alpha}) > 0$, the rotor behavior has an unstable equilibrium where the input is approximately opposite to the rotor state, and cancels the increase. Any deviation from that unstable equilibrium sends the rotor off to infinity.
- For helical inputs, output amplitude is highest when the input frequency equals the resonant frequency.
- Input is in phase with output at the resonant frequency. At higher frequencies, input phase advances ahead of output phase, at lower frequencies it retards.
- Larger values of $\Re(\boldsymbol{\alpha})$ (i.e., closer to 0, since they should be negative) give higher amplitude of output, and also a sharper peak at the resonant frequency.

Although a signal at frequency 0 is not perceptible as sound, a filter with resonance at frequency 0 may be quite useful. $\mathcal{R}_D$, where D is a real constant, is a *low-pass* filter—it boosts low frequencies and suppresses high frequencies, and D determines how steeply its response rolls of with increasing frequency. $\mathcal{R}_0$ just integrates the input, which is the simplest form of linear lowpass filtering. Negative-frequency filters also act as low-pass when the input has only positive frequencies. When working with real-valued signals, instead of complex-valued, each resonance at frequency F is usually paired with another at frequency $-F$. Even a resonance at infinite frequency makes sense—it is a *high-pass* filter, which boosts high frequencies and suppresses low frequencies—but resonance at infinite frequency cannot be implemented with a decaying rotor. For now, just think of it as the limiting filter as frequency goes infinite.

It is not difficult to show that every unit resonance filter $\mathcal{R}_{\boldsymbol{\alpha}}$ is linear and time-invariant. So, the steady-state behavior of a resonance is entirely determined by its response to unit helices. The amplitude A of the output to a unit helix at frequency F is called the *gain* at F, and the phase shift C is called the *phase delay* at F. Many people abbreviate this to "delay," but that practice is hazardous, since there is another completely different sort of delay in filter behavior. Section 7.1.3 shows how to analyze and describe the entire behavior of a linear time-invariant filter, including transient behavior as well as steady state.

Resonances have very interesting transition effects, in addition to their steady-state behavior. When a frequency component starts or increases in the input, the output component at that frequency responds gradually, and approaches its steady-state value asymptotically. When an input component reduces or disappears, the resonance continues to vibrate, or "ring", at the resonant frequency. Higher values for $\Re(\boldsymbol{\alpha})$ lead to slower response when a component increases, and quicker decay of the ringing. When filters are used in sound reproduction systems and conventional signal-processing applications, these transient effects are often considered undesirable, and extensive design effort is expended to minimize them. But, for sound synthesis, transient effects, and particularly ringing, are often the main point of using a filter.

Just as all sound signals may be expressed as a combination of unit helices (perhaps infinitely many) scaled by complex constants, all causal linear time-invariant filters may be expressed as a combination of scaled resonances filters (noncausal filters require time-reversed versions of the resonances). But, many natural and useful filters may be expressed more compactly and more insightfully if we also alo antiresonances, suppressing a given frequency. A particular filter may be expressible by a discrete or even finite combination of resonances and antiresonances, while requiring the integration of a continuous spread of resonances. So, the next subsection investigates antiresonances.

Unit Antiresonances

In terms of input/output behavior, a unit antiresonance filter is the inverse of a unit resonance. Unfortunately, the rotor mechanism does not explain antiresonances. Notice that, since the input to a decaying rotor is added to the derivative of the rotor state, the effect of the rotor is to *integrate* input, scaled in some way determined by the resonance, over time. Antiresonance depends on *differentiating* the input. A unit antiresonance filter is causal, finite-impulse-response, and even memoryless (output at time t depends only on the input and its derivative at time t), but not pointwise. Because of the form of analysis in Section 7.1.3, a unit antiresonance is usually called a *single-zero* filter.

Recall the derivative of a unit helix: $\frac{d}{dt}e^{\boldsymbol{i}Bt} = \boldsymbol{i}Be^{\boldsymbol{i}Bt}$. That is, the derivative of a unit helix is another helix at the same frequency, but advanced in phase by $\frac{1}{4}$ cycle (from the multiplication by $\boldsymbol{i}$), and scaled by B (frequency times 2π). The crucial quality of the derivative is the scaling by B, since this gives a mathematical tool for suppressing a particular frequency. We must work around the phase advance, but it is uniform for all frequencies. So, the general form for an antiresonance is to cancel a scaled and rotated version of the input signal against the derivative of the input.

Let $\mathcal{A}_{\boldsymbol{\alpha}}$ be the unit antiresonance with antiresonance characterized by α. Let σ be an input, and let $\rho = \mathcal{A}_{\boldsymbol{\alpha}}(\sigma)$ be the corresponding output, defined by

$$\rho \quad = \quad \sigma' - \boldsymbol{\alpha}\sigma \qquad (7.4)$$

This is the inverse of the resonance filter behavior in Equation 7.3. Solving for σ' yields

$$\sigma' \quad = \quad \rho + \boldsymbol{\alpha}\sigma \qquad (7.5)$$

which has exactly the form of Equation 7.3, with the roles of σ and ρ reversed. So, the output $\mathcal{A}_{\boldsymbol{\alpha}}(\sigma)$ from an antiresonance filter, when presented as input to the corresponding resonance filter $\mathcal{R}_{\boldsymbol{\alpha}}$, yields the original input σ, and the same thing happens when the output of the resonance is supplied as input to the antiresonance.

$$\mathcal{R}_{\boldsymbol{\alpha}}(\mathcal{A}_{\boldsymbol{\alpha}}(\sigma)) = \sigma \tag{7.6}$$

$$\mathcal{A}_{\boldsymbol{\alpha}}(\mathcal{R}_{\boldsymbol{\alpha}}(\sigma)) = \sigma \tag{7.7}$$

Consider the behavior of $\mathcal{A}_{\boldsymbol{\alpha}}$ on a unit helix $e^{\boldsymbol{i}Bt}$. Since $\boldsymbol{\alpha}\sigma$ and σ' have the same frequency as σ, the output is certainly of the form $\rho = Ae^{\boldsymbol{i}(Bt+C)}$, and we need to determine the amplitude A and phase shift C. $\boldsymbol{i}\Im(\boldsymbol{\alpha})\sigma$ is a vector in the same direction as σ', so by subtracting this from σ' we cancel entirely when $B = \Im(\boldsymbol{\alpha})$, and nearly cancel when B is close to the antiresonant frequency. We are also left with $-\Re(\boldsymbol{\alpha})\sigma$, which is just a scaling of the input. Look at Figures (Maple manuscript) to see how this works.

Here's the mathematical derivation, but you may skip it as before.

$$Ae^{\boldsymbol{i}C}e^{\boldsymbol{i}Bt} = Ae^{\boldsymbol{i}(Bt+C)} = \rho = \sigma' - \boldsymbol{\alpha}\sigma = \boldsymbol{i}Be^{\boldsymbol{i}Bt} - \boldsymbol{\alpha}e^{\boldsymbol{i}Bt} = (\boldsymbol{i}B - \boldsymbol{\alpha})e^{\boldsymbol{i}Bt}$$

Comparing the first and last terms, $Ae^{\boldsymbol{i}C} = \boldsymbol{i}B - \boldsymbol{\alpha}$, so

$$A\cos(C) = \Re(Ae^{\boldsymbol{i}C}) = \Re(\boldsymbol{i}B - \boldsymbol{\alpha}) = \Re(\boldsymbol{\alpha})$$

and

$$A\sin(C) = \Im(Ae^{\boldsymbol{i}C}) = \Im(\boldsymbol{i}B - \boldsymbol{\alpha}) = B - \Im(\boldsymbol{\alpha})$$

From the real part, $\cos(C) = \Re(\boldsymbol{\alpha})/A$. From the imaginary part, $\sin(C) = (B - \Im(\boldsymbol{\alpha}))/A$, so $C = \arctan((B - \Im(\boldsymbol{\alpha}))/\Re(\boldsymbol{\alpha})$ and $A = \Re(\boldsymbol{\alpha})/\cos(C) = (B - \Im(\boldsymbol{\alpha}))/\sin(C) = |\boldsymbol{\alpha} - \boldsymbol{i}B|$. It is easy to see that A is minimum, and C is 0, when $B = \Im(\boldsymbol{\alpha})$. The amplitude A increases more rapidly the closer that $\Re(\boldsymbol{\alpha})$ gets to 0. When $\Re(\boldsymbol{\alpha}) = 0$, the antiresonant frequency $\Im(\boldsymbol{\alpha})/(2\pi)$ is completely suppressed. When $B > \Im(\boldsymbol{\alpha})$, then $C > 0$, so the phase of the output lags behind the input. When $B < \Im(\boldsymbol{\alpha})$, then $C < 0$, so the phase of the output advances ahead of the input. This is the opposite of the phase behavior for a resonance. Unlike a resonant filter, an antiresonant filter is stable even when $\Re(\boldsymbol{\alpha}) = 0$ and $\Re(\boldsymbol{\alpha}) > 0$. When the real part is 0, we get complete suppression of the antiresonant frequency. When the real part is positive, we get the same output amplitude as if the real part were negated, but the opposite phase.

Even though you skipped the mathematical derivations, please review these critical qualitative observations.

- $\mathcal{A}_{\boldsymbol{\alpha}}$ has an antiresonant frequency at $\Im(\boldsymbol{\alpha})/(2\pi)$.
- $\mathcal{A}_{\boldsymbol{\alpha}}$ has no internal state, so it does not produce sonically interesting output when released in some initial state. Section 7.1.3 shows the mathematical analogue to the decay of the rotor.
- $\mathcal{A}_{\boldsymbol{\alpha}}$ behaves stably for all values of α.
- When $\Re(\boldsymbol{\alpha}) = 0$, the antiresonant frequency is eliminated entirely.
- As $|\Re(\boldsymbol{\alpha})|$ grows, the antiresonant frequency is reduced less and less, and the increase in amplitude away from the antiresonant frequency is less sharp.
- Negating $\Re(\boldsymbol{\alpha})$ corresponds to reversing the phase of the output.
- When $\Re(\boldsymbol{\alpha}) < 0$, frequencies higher than the antiresonant frequency are retarded in phase, lower frequencies are advanced.

????????????? An antiresonance at frequency 0 is the same as a resonance at ∞: that is, $\mathcal{A}_D$ for real constant D is a high-pass filter. In particular, $\mathcal{A}_0$ just differentiates its input, which is the simplest form of linear highpass filtering. An antiresonance at ∞ makes perfect sense, but it is usually more convenient to represent it as a resonance at 0. ARE THESE REALLY THE SAME?????? I THINK NOT.

The antiresonances, like the resonances, are linear and time-invariant. So, their steady-state behaviors are determined entirely by their responses to unit helices. As with resonances, the amplitude A of the output to a unit helix at frequency F is the gain, and the shift C is the phase delay at F. Antiresonant filters are memoryless, and present output depends only on the value and derivative of present input, so they do not create the transient effects, particularly ringing, of resonant filters. But, because of the dependence on input derivative, antiresonances react in interesting ways to transients in the input. A sudden change in a frequency component of the input boosts the amplitude of that component in the output. Higher values of $\Re(\boldsymbol{\alpha})$ lead to stronger reactions to changes in the input.

Identity, Constant Output, Time Shifting

There are a few linear time-invariant behaviors that do not develop naturally from resonances and antiresonances. But, they are quite simple to understand directly. The *identity* filter produces output exactly the same as its input.

$$\mathcal{I}(\sigma) \;=\; \sigma \tag{7.8}$$

$\mathcal{I}$ is neither a resonance nor an antiresonance, but when we consider combinations of resonances and antiresonances, it is the natural starting point—the combination of 0 resonances and 0 antiresonances. The constant output filters produce the same output no matter what the input.

$$\mathcal{C}_{\boldsymbol{\alpha}}(\sigma) \;=\; \boldsymbol{\alpha} \tag{7.9}$$

Notice that the output must be a complex constant, independent of time, else the filter would not be time-invariant.

The remaining time-invariant filtering behavior that does not appear to develop naturally from resonances and antiresonances is time shifting.

$$\mathcal{S}_E(\sigma)(t) \;=\; \sigma(t-E) \tag{7.10}$$

When $E > 0$, the time shifter $\mathcal{S}_E$ is causal, and represents a delay. The delay induced by $\mathcal{S}_E$ is called *pulse delay.* It is totally different from the phase delay induced by resonances and antiresonances.

Combining Filters in Sequence and In Parallel

Given the unit resonances, unit antiresonances, identity, constants, and time shifters, a lot of interesting filter behaviors may be constructed by the intuitive equivalent of scaling the filters and wiring them together. The *scaling* of a filter $\mathcal{T}$ by $\boldsymbol{\alpha}$ multiplies the output by $\boldsymbol{\alpha}$. When the filter is linear and time-invariant, this has the same result as multiplying the input by $\boldsymbol{\alpha}$.

$$(\boldsymbol{\alpha}\mathcal{T})(\sigma) \;=\; \boldsymbol{\alpha}(\mathcal{T}(\sigma)) \tag{7.11}$$

$$(\mathcal{T}\boldsymbol{\alpha})(\sigma) \;=\; \mathcal{T}(\boldsymbol{\alpha}\sigma) \tag{7.12}$$

$$\boldsymbol{\alpha}\mathcal{T} \;=\; \mathcal{T}\boldsymbol{\alpha} \text{ when } \mathcal{T} \text{ is linear time-invariant} \tag{7.13}$$

The *sum* of two filters $\mathcal{T}_1$ and $\mathcal{T}_2$ corresponds to splitting the input signal, feeding it separately to $\mathcal{T}_\infty$ and $\mathcal{T}_2$, and adding the two outputs.

$$(\mathcal{T}_1+\mathcal{T}_2)(\sigma) \;=\; \mathcal{T}_1(\sigma)+\mathcal{T}_2(\sigma) \tag{7.14}$$

The *composition*, or *cascade*, of $\mathcal{T}_1$ and $\mathcal{T}_2$ corresponds to feeding the input to $\mathcal{T}_1$, then taking that output and feeding it to $\mathcal{T}_2$. When both filters are linear and time-invariant, the order makes no difference.

$$(\mathcal{T}_1 \circ \mathcal{T}_2)(\sigma) = \mathcal{T}_1(\mathcal{T}_2(\sigma)) \quad (7.15)$$

$$\mathcal{T}_1 \circ \mathcal{T}_2 = \mathcal{T}_2 \circ \mathcal{T}_1 \text{ when } \mathcal{T}_1, \mathcal{T}_2 \text{ are linear time-invariant} \quad (7.16)$$

With suitable notions of infinitary sum and composition, all causal linear time-invariant filters may be constructed from scaled versions of resonances, antiresonances, identity, constant, and time shifters $\mathcal{S}_E$ with $E > 0$. Along with the negative time shifters, they suffice to construct all linear time-invariant filters, although the forms of filters that depend on the indefinite future are rather peculiar, and more natural presentations require time-reversed versions of the resonances. Section **??** explores some of the special forms for representing filters using various operations on the the basic filters.

7.1.2 Forms for Representing Filters

There are a number of different ways to represent linear time-invariant filters, each of them useful in different ways. The best representation for the purpose of implementation is not always the best for understanding and analysis.

Impulse Response and the Convolution Form

The *impulse response* for a filter $\mathcal{T}$ is its output whent the input is an infinitely sharp pulse at time 0, that is $\mathcal{T}(\delta)$. Since the only nonzero input is at time 0, $\mathcal{T}(\delta)(t)$ gives the contribution of an input at time 0 to the output at time t. For a time-invariant filter, this is the same as the contribution of input at time s to output at time $s + t$, for all s. For a linear filter, the contribution of an input point to an output point is proportional to the input value, and the output at a given time is the sum of the contributions from other times. So, $\mathcal{T}(\delta)$ completely determines the behavior of $\mathcal{T}$. At first, it may seem surprising that a single input determines the behavior for the infinitely many different unit helices, but recall that $\mathcal{F}(\delta) = 1$, so δ contains helical components at every frequency.

The description of $\mathcal{T}(\sigma)$ in terms of the impulse response is particularly simple.

$$\mathcal{T}(\sigma) = \mathcal{T}(\delta) * \sigma \quad (7.17)$$

This is the *convolution form* of the filter $\mathcal{T}$. Recall that $(\mathcal{T}(\delta) * \sigma)(t) = \int_{-\infty}^{\infty} \mathcal{T}(\delta)(s)\sigma(t-s)ds$, so the convolution is just the integration of the contributions of each $\sigma(t-s)$ to $\sigma(t)$. This is particularly clear mathematically, but it may be very expensive to compute, particularly if the impulse response goes to 0 slowly.

Interesting properties of a linear time-invariant filter $\mathcal{T}$ may be expressed in terms of the impulse response.

Causal: $\mathcal{T}(\delta)(t) = 0$ for all $t < 0$.

Finite Impulse Response: There is some real constant A such that $\mathcal{T}(\delta)(t) = 0$ for all $t > A$.

Memoryless: $\mathcal{T}(\delta)(t) = 0$ for all $t \neq 0$.

Pointwise: $\mathcal{T}(\delta) = \boldsymbol{\alpha}\delta$ for some complex constant $\boldsymbol{\alpha}$.

Stable: $\lim_{t\to\infty} \mathcal{T}(\delta)(t) = 0$.

Since convolution is a form of weighted integration, it is surprising at first that the convolution form can express filters, such as the antiresonances $\mathcal{A}_{\boldsymbol{\alpha}}$, that depend on derivatives of the input. But, δ', the derivative of the Dirac impulse, has the strange property that it produces derivatives of other functions by convolution.

$$\delta' * \sigma \quad = \quad \sigma' \tag{7.18}$$

$\delta'(t) = 0$ for all real numbers t, but δ' shoots up to ∞ infinitesimally before 0, and down to $-\infty$ infinitesimally after 0, so by convolution it produces the slope of an infinitesimal region about each point in the signal σ. The convolution form suggests another interesting restricted class of filters: a filter is *convolutional* if its impulse response is a normal function from reals to complex numbers, rather than a generalized function.

If we object to convolution with generalized functions, then we may express all filters as sums of convolutions with derivatives

$$\mathcal{T}(\sigma) \quad = \quad \sum_{k=0}^{\infty} \iota_k * \frac{d^k}{dt^k}\sigma \tag{7.19}$$

This form is not unique—there are many different sequences $\iota_0, \iota_1, \ldots$ representing the same filter $\mathcal{T}$—but all of the ι_ks may be normal functions from reals to complex numbers. Perhaps it is more interesting to let the ι_ks be sums of normal functions and time-shifted Diracs $[t]\delta(t-A)$, so that they combine integrals and discrete values of their respective derivatives.

Differential Equation Forms

The unit resonances were all defined by ordinary differential equations. This representation generalizes to define all linear time-invariant filters in the form of ODEs of any chosen order k. That is, $\mathcal{T}(\sigma)=\rho$, where ρ is the solution to the differential equation of the form

$$\frac{d^k}{dt^k}\rho \quad = \quad \iota * \sigma + \eta * \rho \tag{7.20}$$

Appropriate initial conditions are required to select the right one of the infinitely many solutions. For causal filters, the right idea is to initialize ρ and all of its derivatives to 0 at time $-\infty$; formally we require that $\lim_{t\to-\infty} fracd^l dt^l \rho$ is bounded for all l. For noncausal filters, things are more complicated. Essentially, a noncausal filter is the sum of a causal component that starts in value 0 at time $-\infty$ and an anticausal component that ends in value 0 at time ∞.

With $k=0$, $\iota=\mathcal{T}(\delta)$ and $\eta=0$, the ODE form is the same as the convolution form. For every filter $\mathcal{T}$ and positive integer k, we can get a kth order ODE by letting $\iota=\frac{d^k}{dt^k}\mathcal{T}(\delta)=\mathcal{T}(\frac{d^k}{dt^k}\delta)$. But, the ODE form is useful when ι and η are both nonzero, and both simpler or easier to compute with than all of the $\mathcal{T}(\frac{d^k}{dt^k}\delta)$ functions.

Ideally, α and β may be finite sums of impulses and derivatives of impulses. For example, $\mathcal{R}_{\boldsymbol{\alpha}}(\delta)=H[t]e^{-\alpha t}$ is an infinitely long, continuous impulse response, but $\mathcal{R}_{\boldsymbol{\alpha}}$ may also be presented as a first-order ODE with $\iota=1$ and $\eta=\boldsymbol{\alpha}$. The composition $\mathcal{R}_{\boldsymbol{\alpha}_1}\circ\cdots\circ\mathcal{R}_{\boldsymbol{\alpha}_k}$ of k unit resonances may be defined by a finite kth order ODE. For example, $\mathcal{R}_{\boldsymbol{\alpha}_1}\circ\mathcal{R}_{\boldsymbol{\alpha}_2}$ has the 2nd order ODE

$$\frac{d^2}{dt^2}\rho \quad = \quad \sigma+(\boldsymbol{\alpha}_1+\boldsymbol{\alpha}_2)\frac{d}{dt}\rho-\boldsymbol{\alpha}_1\boldsymbol{\alpha}_2\rho \tag{7.21}$$

with $k=2$, $\iota=1$, $\eta=(\boldsymbol{\alpha}_1+\boldsymbol{\alpha}_2)\delta'-\boldsymbol{\alpha}_1\boldsymbol{\alpha}_2$.

Interesting properties of the filter $\mathcal{T}$ may be expressed in terms of the parameters of an ODE form for $\mathcal{T}$.

Causal: $\eta(t)=0$ for all $t>0$.

Finite Impulse Response: $\eta = 0$.

Memoryless: $\eta = 0$ and $\iota = \sum_{l=0}^{\infty} \boldsymbol{\alpha}_l \frac{d^l}{dt^l}\delta$, for some infinite sequence $\boldsymbol{\alpha}_0, \boldsymbol{\alpha}_1, \ldots$ of complex constants.

Pointwise: $\eta = 0$ and $\iota = \sum_{l=0}^{\infty} \boldsymbol{\alpha}_l \delta$, for some infinite sequence $\boldsymbol{\alpha}_0, \boldsymbol{\alpha}_1, \ldots$ of complex constants.

Stable: $|\int_{-\infty}^{\infty} \eta| < 1$????? Or is it $\int_{-\infty}^{\infty} |\eta| < 1$?????

Every differential equation of order k may be reduced to lower orders $k-1, k-2, \ldots 1$ by introducing extra variables. When defining a filter, these extra variables may be understood as the internal state of the filter. For example, $\mathcal{R}_{\boldsymbol{\alpha}_1} \circ \mathcal{R}_{\boldsymbol{\alpha}_2}$, which was represented by a 2nd order ODE above, may also be represented by a 1st order ODE to be solved simultaneously for 2 variables. That is, $\mathcal{R}_{\boldsymbol{\alpha}_1} \circ \mathcal{R}_{\boldsymbol{\alpha}_2}(\sigma) = \rho_1$, where ρ_1, ρ_2 are simultaneous solutions to the ODE

$$\begin{aligned} \rho_1' &= \rho_2 + \boldsymbol{\alpha}_1 \rho_1 \\ \rho_2' &= \sigma + \boldsymbol{\alpha}_2 \rho_2 \end{aligned} \tag{7.22}$$

Multivariate low-order (usually 1st order) ODEs often provide the most direct understanding of a filter as a device with input, internal state that develops as a function of input and previous state, and output derived from the internal state. For example, a multirotor system may be described by an ODE with one complex variable keeping the state of each rotor. The final output may be given as a linear combination of the rotor states.

7.1.3 Analyzing Filters with the Laplace Transform

Just as we analyze every sound signal in terms of unit helices, we may analyze every linear time-invariant filter in terms of unit resonances. Information about antiresonances comes as a byproduct of the analysis of resonances. In fact, the Fourier transform of the impulse response already provides enough information to characterize the behavior of a filter, but that information is revealed more conveniently by an extension of the Fourier transform, called the *Laplace transform*. The Laplace transform of a signal σ is

$$\mathcal{L}(\sigma)(\boldsymbol{\phi}) = \int_{-\infty}^{\infty} \sigma(t) e^{-\boldsymbol{\phi} f t} dt \tag{7.23}$$

The Laplace transform takes a function from real to complex values and produces a function from complex to complex. It generalizes the Fourier transform by allowing the multiplier $\boldsymbol{i}2\pi f$ of t in the exponent to be an arbitrary complex number $\boldsymbol{\phi}$. The Laplace transform of σ contains the Fourier transform in a simple way.

$$\mathcal{F}(\sigma)(f) \quad = \quad \mathcal{L}(\sigma)(\boldsymbol{i}2\pi f) \tag{7.24}$$

In fact, the whole Laplace transform may be defined as a combination of Fourier transforms of σ times a real-valued exponential.

$$\mathcal{L}(\sigma)(\boldsymbol{\phi}) \quad = \quad \mathcal{F}(e^{-\Re(\boldsymbol{\phi})})(\Im(\boldsymbol{\phi})/(2\pi)) \tag{7.25}$$

Fixing an arbitrary constant value D for the real component of $\boldsymbol{\phi}$, the real-to-complex function $[f]\mathcal{L}(\sigma)(D+\boldsymbol{i}f)$ carries enough information to reconstruct σ, so the Laplace transform is highly redundant. Most functions from complex to complex are not the Laplace transform of any signal.

Different books and papers use variations on the definition of the Laplace transform. Many use a *one-sided Laplace transform*, with integral from 0 to ∞. Of course, the one-sided transform of σ is just the two-sided transform of $H\sigma$. For greater generality, and closer correspondence with the Fourier transform, the two-sided form is better.

Roughly, $\mathcal{L}(\sigma)(\boldsymbol{\phi})$ matches σ against the decaying (or expanding) helical pattern $[t]e^{\boldsymbol{\phi}t}$. The negation of $\boldsymbol{\phi}$ in the exponent within the integral defining the transform serves the purpose of conjugating the helix, as in the Fourier transform, and it also reverses the decay/expansion of the helix. This reversal of decay/expansion causes the product $[t]\sigma(t)e^{-\boldsymbol{\phi}t}$ to be the constant 1 precisely when $\sigma = [t]e^{\boldsymbol{\phi}t}$. But, while the helical patterns in the Fourier transform are linearly independent, all decaying/expanding helices at the same frequency are dependent. So, a resonance $\boldsymbol{\alpha}$ affects the Laplace transform at $\boldsymbol{\beta}$ as long as $\Im(\boldsymbol{\alpha}) = \Im(\boldsymbol{\beta})$, even when the real components differ.

The precise understanding of the Laplace transform as pattern matching is rather subtle, and beyond the mathematical powers of these notes. In particular, for many functions of interest, the defining integral diverges, and $\mathcal{L}(\sigma)(\boldsymbol{\phi})$ is undefined, except for a small band of values of $\boldsymbol{\phi}$. Some of the divergence may be understood through generalized functions, but much of it may not. Some applications of Laplace transform require thorough understanding of the regions of convergence. For our purposes, the Laplace transform of most useful signals produces a nice symbolic expression, and the function defined by that expression gives the insight that we need, even though strictly speaking the expression is only correct within the region of convergence. This is spooky, but we'll have to live with it. I am sure that there is a way of recasting the mathematics to explain the apparent values of Laplace transforms in the divergent regions, but I haven't found it yet.

The Laplace transform has properties analogous to those of the Fourier transform, but with some differences on inputs with nonzero real components. Most tables of Laplace transforms leave out δs and other functional components with infinite values, since by definition they are outside of the region of convergence. This confuses me, and obscures the correspondence between Laplace transform and Fourier transform, so I have included these infinite components.

$$\mathcal{L}(\delta) = 1 \tag{7.26}$$

$$\mathcal{L}(\delta') = [\phi]\phi \tag{7.27}$$

$$\mathcal{L}(\frac{d^k}{dt^k}\delta) = [\phi]\phi^k \tag{7.28}$$

$$\mathcal{L}(\delta' - \boldsymbol{\alpha}\delta) = [\phi]\phi - \boldsymbol{\alpha} \tag{7.29}$$

$$\mathcal{L}(\delta'' - 2\boldsymbol{\alpha}\delta' + \boldsymbol{\alpha}^2\delta) = [\phi](\phi - \boldsymbol{\alpha})^2 \tag{7.30}$$

$$\mathcal{L}([t]\delta(t - A)) = [\phi]e^{A\phi} \tag{7.31}$$

$$\mathcal{L}(H) = [\phi](\phi^{-1} + \delta(\phi)/2) \tag{7.32}$$

$$\mathcal{L}(H - \frac{1}{2}) = [\phi]\phi^{-1} \tag{7.33}$$

$$\mathcal{L}(Z) = [\phi]e^{\phi^2/(4\pi)} \tag{7.34}$$

$$\begin{aligned}\mathcal{L}(\Pi) &= [\phi](2\sinh(\phi/2)/\phi) \\ &= [\phi]((e^{\phi/2} - e^{-\phi/2})/\phi)\end{aligned} \tag{7.35}$$

$$\mathcal{L}(\mathrm{sinc}) = [\phi](1 + (\arctan(-\phi/\pi) - \arctan(\phi/\pi))/\pi() \tag{7.36}$$

$$\tag{7.37}$$

To calculate Laplace transforms, use the following rules that show how the transform and other functional operators interact.

$$\mathcal{L}(\alpha + \beta) = \mathcal{L}(\alpha) + \mathcal{L}(\beta) \tag{7.38}$$

$$\mathcal{L}(\overline{\alpha}) = [\phi]\overline{\mathcal{L}(\alpha)((\overline{}\phi))} \tag{7.39}$$

$$\mathcal{L}(\boldsymbol{\alpha}\beta) = \boldsymbol{\alpha}\mathcal{L}(\beta) \text{ for constant } \boldsymbol{\alpha} \tag{7.40}$$

$$\begin{aligned}
\mathcal{L}([t]\beta(t-A)) &= \mathcal{L}(\beta)([\phi]e^{-A\phi}) \text{ for real constant } A && (7.41)\\
\mathcal{L}([t]\beta(At)) &= ([\phi]\mathcal{L}(\beta)(\phi/A))/|A| \text{ for real constant } A && (7.42)\\
\mathcal{L}([t]\beta(-t)) &= ([\phi]\mathcal{L}(\beta)(-\phi)) && (7.43)\\
\mathcal{L}(\alpha * \beta) &= \mathcal{L}(\alpha)\mathcal{L}(\beta) && (7.44)\\
\mathcal{L}(\alpha\beta) &= [\phi](\boldsymbol{i}2\pi)^{-1}\int_{C-\boldsymbol{i}\infty}^{C+\boldsymbol{i}\infty}\mathcal{L}(\alpha)(\psi)\mathcal{L}(\phi-\psi)d\psi \\
&\quad \text{for certain } C \text{ depending on } \phi && (7.45)\\
\mathcal{L}(\alpha \star \beta) &= [\phi]\mathcal{L}(\alpha)(-\phi)\mathcal{L}(\beta) && (7.46)\\
\mathcal{L}(\alpha \star \alpha) &= [\phi]\mathcal{L}(\alpha)(\phi)\mathcal{L}(\alpha)(-\phi) && (7.47)\\
\mathcal{L}(\alpha') &= ([\phi]\phi)\mathcal{L}(\alpha) && (7.48)
\end{aligned}$$

In the convolution theorem (Equation **??**), the complicated-looking integral is a convolution of traces in the imaginary direction through α and β, determined by the real value C. α is traced at C, and β at $\Re(\phi) - C$. If possible, C is chosen so that both traces converge. A different C may be chosen for each ϕ. If more than one choice of C yields convergence, then all such values of C yield the same value.

The Laplace transform is most often applied to signals that are 0 for all negative times, so it is worth reviewing those cases in particular.

$$\begin{aligned}
\mathcal{L}(1H) &= [\phi](1/\phi + \delta(\phi)/2) && (7.49)\\
\mathcal{L}(\xi H) &= [\phi]1/(\phi - \boldsymbol{i}2\pi) && (7.50)\\
\mathcal{L}(\delta H) &= 1 && (7.51)\\
\mathcal{L}(\text{III}H) &= [\phi]\coth(\phi/2) && (7.52)\\
\mathcal{L}(([t]e^{\boldsymbol{\alpha}t})H) &= [\phi](\phi - \boldsymbol{\alpha})^{-1} && (7.53)
\end{aligned}$$

7.1.4 Discrete Filters and the Z Transform

7.2 Sound Creation by Modal Synthesis

7.2.1 Continuously Driven Resonators

7.2.2 Ringing Resonators

7.3 Sound Modification with Formant Filters