18

Introduction to Spectral Techniques

The following three chapters deal with the representation and analysis of images in the frequency domain, based on the decomposition of image signals into sine and cosine functions using the well-known *Fourier transform*. Students often consider this a difficult topic, mainly because of its mathematical flavor and that its practical applications are not immediately obvious. Indeed, most common operations and methods in digital image processing can be sufficiently described in the original signal or image space without even mentioning spectral techniques. This is the reason why we pick up this topic relatively late in this text.

While spectral techniques were often used to improve the efficiency of image-processing operations, this has become increasingly less important due to the high power of modern computers. There exist, however, some important effects, concepts, and techniques in digital image processing that are considerably easier to describe in the frequency domain or cannot otherwise be understood at all. The topic should therefore not be avoided all together. Fourier analysis not only owns a very elegant (perhaps not always sufficiently appreciated) mathematical theory but interestingly enough also complements some important concepts we have seen earlier, in particular linear filters and linear *convolution* (see Chapter 5, Sec. 5.2). Equally important are applications of spectral techniques in many popular methods for image and video compression, and they provide valuable insight into the mechanisms of sampling (discretization) of continuous signals as well as the reconstruction and interpolation of discrete signals.

In the following, we first give a basic introduction to the concepts of frequency and spectral decomposition that tries to be minimally formal and thus should be easily "digestible" even for readers without previous exposure to this topic. We start with the representation of 1D signals and will then extend the discussion to 2D signals (images) in the next chapter. Subsequently, Chapter 20 briefly explains the *discrete cosine transform*, a popular variant of the discrete Fourier transform that is frequently used in image compression.

W. Burger, M.J. Burge, *Digital Image Processing*, Texts in Computer Science,
DOI 10.1007/978-1-4471-6684-9_18

18.1 The Fourier Transform

The concept of frequency and the decomposition of waveforms into elementary "harmonic" functions first arose in the context of music and sound. The idea of describing acoustic events in terms of "pure" sinusoidal functions does not seem unreasonable, considering that sine waves appear naturally in every form of oscillation (e.g., on a free-swinging pendulum).

18.1.1 Sine and Cosine Functions

The well-known *cosine* function,

$$f(x) = \cos(x), \tag{18.1}$$

has the value 1 at the origin ($\cos(0) = 1$) and performs exactly *one* full cycle between the origin and the point $x = 2\pi$ (Fig. 18.1(a)). We say that the function is periodic with a cycle length (period) $T = 2\pi$; that is,

$$\cos(x) = \cos(x + 2\pi) = \cos(x + 4\pi) = \cdots = \cos(x + k2\pi), \tag{18.2}$$

for any $k \in \mathbb{Z}$. The same is true for the corresponding *sine* function, except that its value is zero at the origin (since $\sin(0) = 0$).

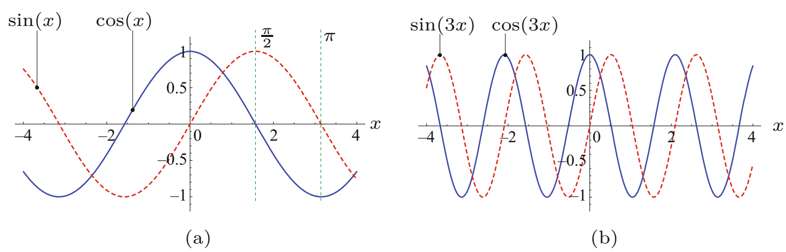

Fig. 18.1 Cosine and sine functions of different frequency. The expression $\cos(\omega x)$ describes a cosine function with angular frequency ω at position x. The angular frequency ω of this periodic function corresponds to a cycle length (period) $T = 2\pi/\omega$. For $\omega = 1$, the period is $T_1 = 2\pi$ (a), and for $\omega = 3$ it is $T_3 = 2\pi/3 \approx 2.0944$ (b). The same holds for the sine function $\sin(\omega x)$.

Frequency and amplitude

The number of oscillations of $\cos(x)$ over the distance $T = 2\pi$ is *one* and thus the value of the *angular frequency*

$$\omega = \frac{2\pi}{T} = 1\,. \tag{18.3}$$

If we modify the cosine function in Eqn. (18.1) to

$$f(x) = \cos(3x)\,, \tag{18.4}$$

we obtain a compressed cosine wave that oscillates three times faster than the original function $\cos(x)$ (see Fig. 18.1(b)). The function $\cos(3x)$ performs three full cycles over a distance of 2π and thus has the angular frequency $\omega = 3$ and a period $T = \frac{2\pi}{3}$. In general, the period T relates to the angular frequency ω as

$$T = \tfrac{2\pi}{\omega}, \tag{18.5}$$

for $\omega > 0$. A sine or cosine function oscillates between peak values $+1$ and -1, and its *amplitude* is 1. Multiplying by a constant $a \in \mathbb{R}$

changes the peak values of the function to $\pm a$ and its *amplitude* to a. In general, the expressions

$$a \cdot \cos(\omega x) \qquad \text{and} \qquad a \cdot \sin(\omega x)$$

denote a cosine or sine function, respectively, with amplitude a and angular frequency ω, evaluated at position (or point in time) x. The relation between the angular frequency ω and the "common" frequency f is given by

$$f = \frac{1}{T} = \frac{\omega}{2\pi} \quad \text{or} \quad \omega = 2\pi f, \tag{18.6}$$

respectively, where f is measured in cycles per length or time unit.[1] In the following, we use either ω or f as appropriate, and the meaning should always be clear from the symbol used.

Phase

Shifting a cosine function along the x axis by a distance φ,

$$\cos(x) \rightarrow \cos(x - \varphi),$$

changes the *phase* of the cosine wave, and φ denotes the *phase angle* of the resulting function. Thus a sine function is really just a cosine function shifted to the right[2] by a quarter period ($\varphi = \frac{2\pi}{4} = \frac{\pi}{2}$), so

$$\sin(\omega x) = \cos\left(\omega x - \tfrac{\pi}{2}\right). \tag{18.7}$$

If we take the cosine function as the reference with phase $\varphi_{\cos} = 0$, then the phase angle of the corresponding sine function is $\varphi_{\sin} = \frac{\pi}{2} = 90°$.

Cosine and sine functions are "orthogonal" in a sense and we can use this fact to create new "sinusoidal" functions with arbitrary frequency, phase, and amplitude. In particular, adding a cosine and a sine function with the identical frequencies ω and arbitrary amplitudes A and B, respectively, creates another sinusoid:

$$A \cdot \cos(\omega x) + B \cdot \sin(\omega x) = C \cdot \cos(\omega x - \varphi). \tag{18.8}$$

The resulting amplitude C and the phase angle φ are defined only by the two original amplitudes A and B as

$$C = \sqrt{A^2 + B^2} \quad \text{and} \quad \varphi = \tan^{-1}\left(\tfrac{B}{A}\right). \tag{18.9}$$

Figure 18.2(a) shows an example with amplitudes $A = B = 0.5$ and a resulting phase angle $\varphi = 45°$.

[1] For example, a temporal oscillation with frequency $f = 1000$ cycles/s (Hertz) has the period $T = 1/1000$ s and therefore the angular frequency $\omega = 2000\pi$. The latter is a unitless quantity.

[2] In general, the function $f(x-d)$ is the original function $f(x)$ shifted to the right by a distance d.

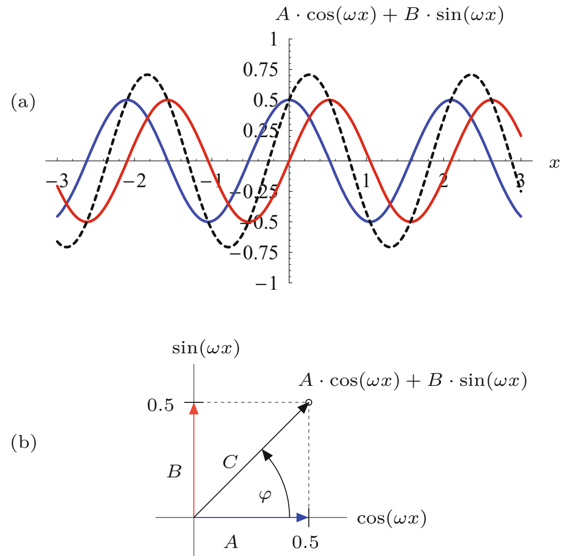

Fig. 18.2
Adding cosine and sine functions with identical frequencies, $A \cdot \cos(\omega x) + B \cdot \sin(\omega x)$, with $\omega = 3$ and $A = B = 0.5$. The result is a phase-shifted cosine function (dotted curve) with amplitude $C = \sqrt{0.5^2 + 0.5^2} \approx 0.707$ and phase angle $\varphi = 45°$ (a). If the cosine and sine components are treated as orthogonal vectors (A, B) in 2-space, the amplitude and phase of the resulting sinusoid (C) can be easily determined by vector summation (b).

Complex-valued sine functions—Euler's notation

Figure 18.2(b) depicts the contributing cosine and sine components of the new function as a pair of orthogonal vectors in 2-space whose *lengths* correspond to the amplitudes A and B. Not coincidentally, this reminds us of the representation of real and imaginary components of complex numbers,

$$z = a + \mathrm{i}\, b,$$

in the 2D plane $\mathbb{C}$, where i is the imaginary unit ($\mathrm{i}^2 = -1$). This association becomes even stronger if we look at Euler's famous notation of complex numbers along the unit circle,

$$z = e^{\mathrm{i}\theta} = \cos(\theta) + \mathrm{i} \cdot \sin(\theta), \tag{18.10}$$

where $e \approx 2.71828$ is the Euler number. If we take the expression $e^{\mathrm{i}\theta}$ as a function of the angle θ rotating around the unit circle, we obtain a "complex-valued sinusoid" whose real and imaginary parts correspond to a cosine and a sine function, respectively,

$$\begin{aligned} \mathrm{Re}(e^{\mathrm{i}\theta}) &= \cos(\theta), \\ \mathrm{Im}(e^{\mathrm{i}\theta}) &= \sin(\theta). \end{aligned} \tag{18.11}$$

Since $z = e^{\mathrm{i}\theta}$ is placed on the unit circle, the *amplitude* of the complex-valued sinusoid is $|z| = r = 1$. We can easily modify the amplitude of this function by multiplying it by some real value $a \geq 0$, that is,

$$|a \cdot e^{\mathrm{i}\theta}| = a \cdot |e^{\mathrm{i}\theta}| = a\,. \tag{18.12}$$

Similarly, we can alter the *phase* of a complex-valued sinusoid by adding a phase angle φ in the function's exponent or, equivalently, by multiplying it by a complex-valued constant $c = e^{\mathrm{i}\varphi}$,

$$e^{\mathrm{i}(\theta+\varphi)} = e^{\mathrm{i}\theta} \cdot e^{\mathrm{i}\varphi}. \tag{18.13}$$

In summary, multiplying by some real value affects only the *amplitude* of a sinusoid, while multiplying by some complex value c (with unit amplitude $|c| = 1$) modifies only the function's *phase* (without changing its amplitude). In general, of course, multiplying by some arbitrary complex value changes both the amplitude *and* the phase of the function (also see Sec. A.3 in the Appendix).

The complex notation makes it easy to combine orthogonal pairs of sine functions $\cos(\omega x)$ and $\sin(\omega x)$ with identical frequencies ω into a single expression,

$$e^{\mathrm{i}\theta} = e^{\mathrm{i}\omega x} = \cos(\omega x) + \mathrm{i} \cdot \sin(\omega x). \tag{18.14}$$

We will make more use of this notation later (in Sec. 18.1.4) to explain the Fourier transform.

18.1.2 Fourier Series Representation of Periodic Functions

As we demonstrated in Eqn. (18.8), sinusoidal functions of arbitrary frequency, amplitude, and phase can be described as the sum of suitably weighted cosine and sine functions. One may wonder if non-sinusoidal functions can also be decomposed into a sum of cosine and sine functions. The answer is yes, of course. It was Fourier[3] who first extended this idea to arbitrary functions and showed that (almost) any periodic function $g(x)$ with a fundamental frequency ω_0 can be described as a—possibly infinite—sum of "harmonic" sinusoids, that is,

$$g(x) = \sum_{k=0}^{\infty} A_k \cdot \cos(k\omega_0 x) + B_k \cdot \sin(k\omega_0 x). \tag{18.15}$$

This is called a *Fourier series*, and the constant factors A_k, B_k are the *Fourier coefficients* of the function $g(x)$. Notice that in Eqn. (18.15) the frequencies of the sine and cosine functions contributing to the Fourier series are integral multiples ("harmonics") of the fundamental frequency ω_0, including the zero frequency for $k = 0$. The corresponding coefficients A_k and B_k, which are initially unknown, can be uniquely derived from the original function $g(x)$. This process is commonly referred to as *Fourier analysis*.

18.1.3 Fourier Integral

Fourier did not want to limit this concept to periodic functions and postulated that nonperiodic functions, too, could be described as sums of sine and cosine functions. While this proved to be true in principle, it generally requires—beyond multiples of the fundamental frequency ($k\omega_0$)—infinitely many, densely spaced frequencies! The resulting decomposition,

$$g(x) = \int_{0}^{\infty} A_\omega \cdot \cos(\omega x) + B_\omega \cdot \sin(\omega x)\, \mathrm{d}\omega, \tag{18.16}$$

is called a *Fourier integral* and the coefficients A_ω, B_ω are again the weights for the corresponding cosine and sine functions with the

[3] Jean-Baptiste Joseph de Fourier (1768–1830).

(continuous) frequency ω. The Fourier integral is the basis of the Fourier spectrum and the Fourier transform, as will be described (for details, see, e.g., [35, Ch. 15, Sec. 15.3]).

In Eqn. (18.16), every coefficient A_ω and B_ω specifies the *amplitude* of the corresponding cosine or sine function, respectively. The coefficients thus define "how much of each frequency" contributes to a given function or signal $g(x)$. But what are the proper values of these coefficients for a given function $g(x)$, and can they be determined uniquely? The answer is yes again, and the "recipe" for computing the coefficients is amazingly simple:

$$\begin{aligned} A_\omega = A(\omega) &= \frac{1}{\pi} \cdot \int_{-\infty}^{\infty} g(x) \cdot \cos(\omega x) \, \mathrm{d}x, \\ B_\omega = B(\omega) &= \frac{1}{\pi} \cdot \int_{-\infty}^{\infty} g(x) \cdot \sin(\omega x) \, \mathrm{d}x. \end{aligned} \tag{18.17}$$

Since this representation of the function $g(x)$ involves infinitely many densely spaced frequency values ω, the corresponding coefficients $A(\omega)$ and $B(\omega)$ are indeed continuous functions as well. They hold the continuous distribution of frequency components contained in the original signal, which is called a "spectrum".

Thus the Fourier integral in Eqn. (18.16) describes the original function $g(x)$ as a sum of infinitely many cosine and sine functions, with the corresponding Fourier coefficients contained in the functions $A(\omega)$ and $B(\omega)$. In addition, a signal $g(x)$ is uniquely and fully represented by the corresponding coefficient functions $A(\omega)$ and $B(\omega)$. We know from Eqn. (18.17) how to compute the spectrum for a given function $g(x)$, and Eqn. (18.16) explains how to reconstruct the original function from its spectrum if it is ever needed.

18.1.4 Fourier Spectrum and Transformation

There is now only a small remaining step from the decomposition of a function $g(x)$, as shown in Eqn. (18.17), to the "real" Fourier transform. In contrast to the Fourier *integral*, the Fourier *transform* treats both the original signal and the corresponding spectrum as *complex-valued* functions, which considerably simplifies the resulting notation.

Based on the functions $A(\omega)$ and $B(\omega)$ defined in the Fourier integral (Eqn. (18.17)), the *Fourier spectrum* $G(\omega)$ of a function $g(x)$ is given as

$$\begin{aligned} G(\omega) &= \sqrt{\tfrac{\pi}{2}} \cdot \big[A(\omega) - \mathrm{i} \cdot B(\omega)\big] \\ &= \sqrt{\tfrac{\pi}{2}} \cdot \Big[\frac{1}{\pi} \int_{-\infty}^{\infty} g(x) \cdot \cos(\omega x) \, \mathrm{d}x \;-\; \mathrm{i} \cdot \frac{1}{\pi} \int_{-\infty}^{\infty} g(x) \cdot \sin(\omega x) \, \mathrm{d}x\Big] \\ &= \frac{1}{\sqrt{2\pi}} \cdot \int_{-\infty}^{\infty} g(x) \cdot \big[\cos(\omega x) - \mathrm{i} \cdot \sin(\omega x)\big] \, \mathrm{d}x \,, \end{aligned} \tag{18.18}$$

with $g(x), G(\omega) \in \mathbb{C}$. Using Euler's notation of complex values (see Eqn. (18.14)) yields the continuous Fourier spectrum in Eqn. (18.18) in its common form:

$$\boxed{\begin{aligned} G(\omega) &= \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} g(x) \cdot \big[\cos(\omega x) - \mathrm{i} \cdot \sin(\omega x)\big] \, \mathrm{d}x \\ &= \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} g(x) \cdot e^{-\mathrm{i}\omega x} \, \mathrm{d}x \,. \end{aligned}} \tag{18.19}$$

The transition from the function $g(x)$ to its Fourier spectrum $G(\omega)$ is called the *Fourier transform*[4] ($\mathcal{F}$). Conversely, the original function $g(x)$ can be reconstructed completely from its Fourier spectrum $G(\omega)$ using the *inverse Fourier transform*[5] ($\mathcal{F}^{-1}$), defined as

$$\boxed{\begin{aligned} g(x) &= \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} G(\omega) \cdot \big[\cos(\omega x) + \mathrm{i} \cdot \sin(\omega x)\big] \, \mathrm{d}\omega \\ &= \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} G(\omega) \cdot e^{\mathrm{i}\omega x} \, \mathrm{d}\omega \,. \end{aligned}} \tag{18.20}$$

In general, even if one of the involved functions ($g(x)$ or $G(\omega)$) is real-valued (which is usually the case for physical signals $g(x)$), the other function is complex-valued. One may also note that the forward transformation $\mathcal{F}$ (Eqn. (18.19)) and the inverse transformation $\mathcal{F}^{-1}$ (Eqn. (18.20)) are almost completely symmetrical, the sign of the exponent being the only difference.[6] The spectrum produced by the Fourier transform is a new representation of the signal in a space of frequencies. Apparently, this "frequency space" and the original "signal space" are *dual* and interchangeable mathematical representations.

18.1.5 Fourier Transform Pairs

The relationship between a function $g(x)$ and its Fourier spectrum $G(\omega)$ is unique in both directions: the Fourier spectrum is uniquely defined for a given function, and for any Fourier spectrum there is only one matching signal—the two functions $g(x)$ and

$$g(x) \circ\!\!-\!\!\bullet\, G(\omega).$$

Table 18.1 lists the transform pairs for some selected analytical functions, which are also shown graphically in Figs. 18.3 and 18.4.

The Fourier spectrum of a *cosine function* $\cos(\omega_0 x)$, for example, consists of two separate thin pulses arranged symmetrically at a distance ω_0 from the origin (Fig. 18.3(a, c)). Intuitively, this corresponds to our physical understanding of a spectrum (e.g., if we think of a pure monophonic sound in acoustics or the thin line produced by some extremely pure color in the optical spectrum). Increasing the frequency ω_0 would move the corresponding pulses in the spectrum

[4] Also called the "direct" or "forward" transformation.

[5] Also called "backward" transformation.

[6] Various definitions of the Fourier transform are in common use. They are contrasted mainly by the constant factors outside the integral and the signs of the exponents in the forward and inverse transforms, but all versions are equivalent in principle. The symmetric variant shown here uses the same factor ($1/\sqrt{2\pi}$) in the forward and inverse transforms.

Table 18.1
Fourier transforms of selected analytical functions; $\delta()$ denotes the "impulse" or *Dirac* function (see Sec. 18.2.1).

Function	*Transform pair* $g(x) \circ\!\!-\!\!\bullet\, G(\omega)$	*Figure*
Cosine function with frequency ω_0	$g(x) = \cos(\omega_0 x)$ $G(\omega) = \sqrt{\frac{\pi}{2}} \cdot \big(\delta(\omega+\omega_0) + \delta(\omega-\omega_0)\big)$	18.3(a,c)
Sine function with frequency ω_0	$g(x) = \sin(\omega_0 x)$ $G(\omega) = \mathrm{i}\sqrt{\frac{\pi}{2}} \cdot \big(\delta(\omega+\omega_0) - \delta(\omega-\omega_0)\big)$	18.3(b,d)
Gaussian function of width σ	$g(x) = \frac{1}{\sigma} \cdot e^{-\frac{x^2}{2\sigma^2}}$ $G(\omega) = e^{-\frac{\sigma^2\omega^2}{2}}$	18.4(a,b)
Rectangular pulse of width $2b$	$g(x) = \Pi_b(x) = \begin{cases} 1 & \lvert x\rvert \leq b \\ 0 & \text{sonst} \end{cases}$ $G(\omega) = \frac{2b\sin(b\omega)}{\sqrt{2\pi}\omega}$	18.4(c,d)

away from the origin. Notice that the spectrum of the cosine function is real-valued, the imaginary part being zero. Of course, the same relation holds for the sine function (Fig. 18.3(b, d)), with the only difference being that the pulses have different polarities and appear in the imaginary part of the spectrum. In this case, the real part of the spectrum $G(\omega)$ is zero.

The *Gaussian function* is particularly interesting because its Fourier spectrum is also a Gaussian function (Fig. 18.4(a, b))! It is one of the few examples where the function type in frequency space is the same as in signal space. With the Gaussian function, it is also clear to see that *stretching* a function in signal space corresponds to *shortening* its spectrum and vice versa.

The Fourier transform of a *rectangular pulse* (Fig. 18.4(c, d)) is the "Sinc" function of type $\sin(x)/x$. With increasing frequencies, this function drops off quite slowly, which shows that the components contained in the original rectangular signal are spread out over a large frequency range. Thus a rectangular pulse function exhibits a very wide spectrum in general.

18.1.6 Important Properties of the Fourier Transform

Symmetry

The Fourier spectrum extends over positive and negative frequencies and could, in principle, be an arbitrary complex-valued function. However, in many situations, the spectrum is symmetric about its origin (see, e.g., [43, p. 178]). In particular, the Fourier transform of a real-valued signal $g(x) \in \mathbb{R}$ is a so-called *Hermite* function with the property

$$G(\omega) = G^*(-\omega), \tag{18.21}$$

where G^* denotes the complex conjugate of G (see also Sec. A.3 in the Appendix).

Fig. 18.3
Fourier transform pairs—cosine and sine functions.

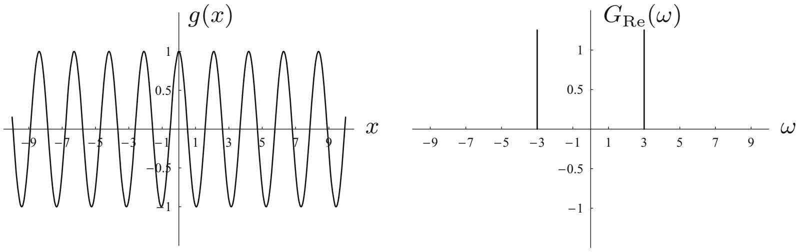

(a) Cosine ($\omega_0 = 3$): $g(x) = \cos(3x)$ $\circ\!\!-\!\!\bullet$ $G(\omega) = \sqrt{\frac{\pi}{2}} \cdot \big(\delta(\omega+3) + \delta(\omega-3)\big)$

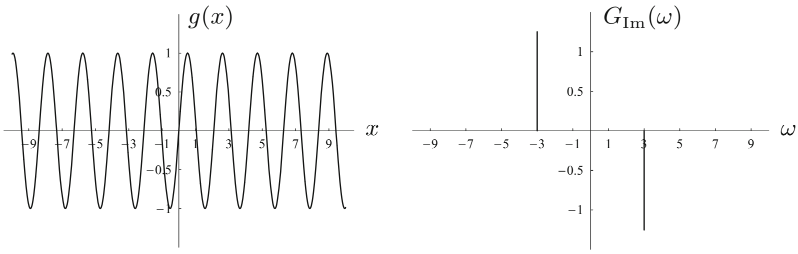

(b) Sine ($\omega_0 = 3$): $g(x) = \sin(3x)$ $\circ\!\!-\!\!\bullet$ $G(\omega) = \mathrm{i}\sqrt{\frac{\pi}{2}} \cdot \big(\delta(\omega+3) - \delta(\omega-3)\big)$

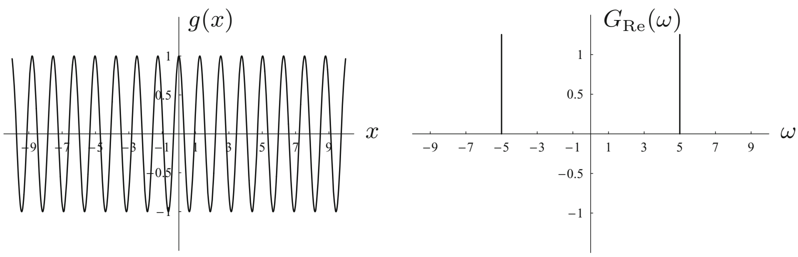

(c) Cosine ($\omega_0 = 5$): $g(x) = \cos(5x)$ $\circ\!\!-\!\!\bullet$ $G(\omega) = \sqrt{\frac{\pi}{2}} \cdot \big(\delta(\omega+5) + \delta(\omega-5)\big)$

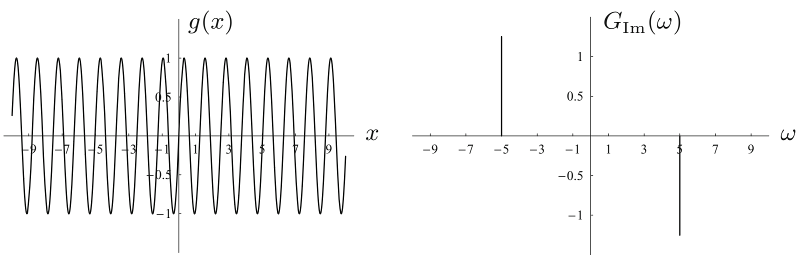

(d) Sine ($\omega_0 = 5$): $g(x) = \sin(5x)$ $\circ\!\!-\!\!\bullet$ $G(\omega) = \mathrm{i}\sqrt{\frac{\pi}{2}} \cdot \big(\delta(\omega+5) - \delta(\omega-5)\big)$

Fig. 18.4
Fourier transform pairs—Gaussian functions and square pulses.

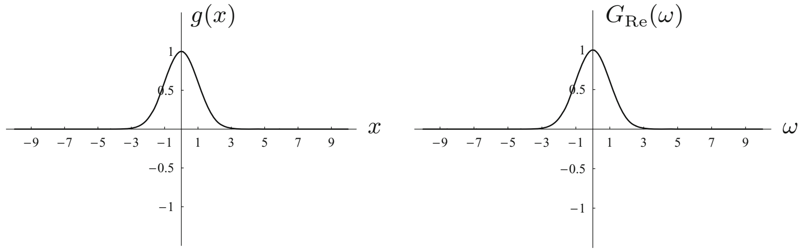

(a) Gauss. $(\sigma = 1)$: $g(x) = e^{-\frac{x^2}{2}}$ ○—● $G(\omega) = e^{-\frac{\omega^2}{2}}$

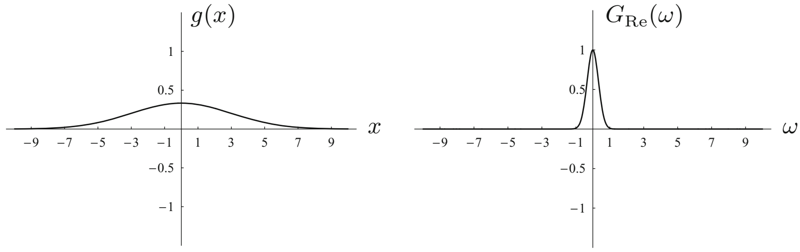

(b) Gauss. $(\sigma = 3)$: $g(x) = \frac{1}{3} \cdot e^{-\frac{x^2}{2 \cdot 9}}$ ○—● $G(\omega) = e^{-\frac{9\omega^2}{2}}$

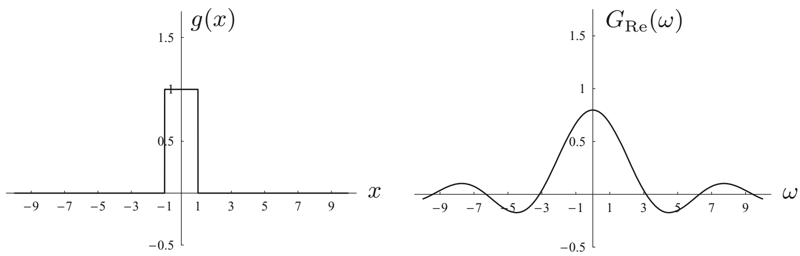

(c) Pulse $(b = 1)$: $g(x) = \Pi_1(x)$ ○—● $G(\omega) = \frac{2\sin(\omega)}{\sqrt{2\pi}\omega}$

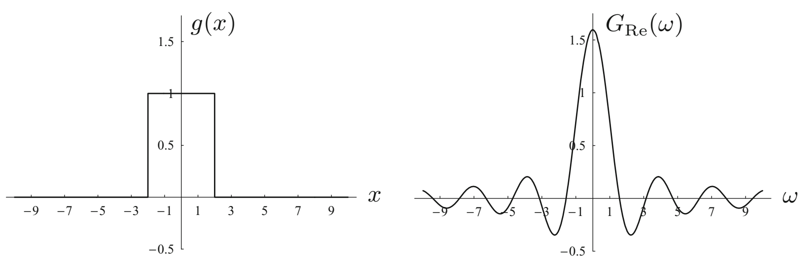

(d) Pulse $(b = 2)$: $g(x) = \Pi_2(x)$ ○—● $G(\omega) = \frac{4\sin(2\omega)}{\sqrt{2\pi}\omega}$

Linearity

The Fourier transform is also a *linear* operation such that multiplying the signal by a constant value $c \in \mathbb{C}$ scales the corresponding spectrum by the same amount,

$$a \cdot g(x) \circ\!\!-\!\!\bullet\, a \cdot G(\omega). \tag{18.22}$$

Linearity also means that the transform of the sum of two signals $g(x) = g_1(x) + g_2(x)$ is identical to the sum of their individual transforms $G_1(\omega)$ and $G_2(\omega)$ and thus

$$g_1(x) + g_2(x) \circ\!\!-\!\!\bullet\, G_1(\omega) + G_2(\omega). \tag{18.23}$$

Similarity

If the original function $g(x)$ is scaled in space or time, the opposite effect appears in the corresponding Fourier spectrum. In particular, as observed on the Gaussian function in Fig. 18.4, *stretching* a signal by a factor s (i.e., $g(x) \rightarrow g(sx)$) leads to a *shortening* of the Fourier spectrum:

$$g(sx) \circ\!\!-\!\!\bullet\, \tfrac{1}{|s|} \cdot G\left(\tfrac{\omega}{s}\right). \tag{18.24}$$

Similarly, the signal is shortened if the corresponding spectrum is stretched.

Shift property

If the original function $g(x)$ is shifted by a distance d along its coordinate axis (i.e., $g(x) \rightarrow g(x-d)$), then the Fourier spectrum multiplies by the complex value $e^{-\mathrm{i}\omega d}$ dependent on ω:

$$g(x-d) \circ\!\!-\!\!\bullet\, e^{-\mathrm{i}\omega d} \cdot G(\omega). \tag{18.25}$$

Since $e^{-\mathrm{i}\omega d}$ lies on the unit circle, the multiplication causes a phase shift on the spectral values (i.e., a redistribution between the real and imaginary components) without altering the magnitude $|G(\omega)|$. Obviously, the amount (angle) of phase shift (ωd) is proportional to the angular frequency ω.

Convolution property

From the image-processing point of view, the most interesting property of the Fourier transform is its relation to linear convolution (see Ch. 5, Sec. 5.3.1). Let us assume that we have two functions $g(x)$ and $h(x)$ and their corresponding Fourier spectra $G(\omega)$ and $H(\omega)$, respectively. If the original functions are subject to linear convolution (i.e., $g(x) * h(x)$), then the Fourier transform of the result equals the (pointwise) product of the individual Fourier transforms $G(\omega)$ and $H(\omega)$:

$$g(x) * h(x) \circ\!\!-\!\!\bullet\, G(\omega) \cdot H(\omega). \tag{18.26}$$

Due to the duality of signal space and frequency space, the same also holds in the opposite direction; i.e., a pointwise multiplication of two signals is equivalent to convolving the corresponding spectra:

$$g(x) \cdot h(x) \circ\!\!-\!\!\bullet\, G(\omega) * H(\omega). \tag{18.27}$$

A multiplication of the functions in *one* space (signal or frequency space) thus corresponds to a linear convolution of the Fourier spectra in the *opposite* space.

18.2 Working with Discrete Signals

The definition of the continuous Fourier transform in Sec. 18.1 is of little use for numerical computation on a computer. Neither can arbitrary continuous (and possibly infinite) functions be represented in practice. Nor can the required integrals be computed. In reality, we must always deal with *discrete* signals, and we therefore need a new version of the Fourier transform that treats signals and spectra as finite data vectors—the "discrete" Fourier transform. Before continuing with this issue we want to use our existing wisdom to take a closer look at the process of discretizing signals in general.

18.2.1 Sampling

We first consider the question of how a continuous function can be converted to a discrete signal in the first place. This process is usually called "sampling" (i.e., taking samples of the continuous function at certain points in time (or in space), usually spaced at regular distances). To describe this step in a simple but formal way, we require an inconspicuous but nevertheless important piece from the mathematician's toolbox.

The impulse function $\delta(x)$

We casually encountered the impulse function (also called the *delta* or *Dirac* function) earlier when we looked at the impulse response of linear filters (see Ch. 5, Sec. 5.3.4) and in the Fourier transforms of the cosine and sine functions (Fig. 18.3). This function, which models a continuous "ideal" impulse, is unusual in several respects: its value is zero everywhere except at the origin, where it is nonzero (though undefined), but its integral is one, that is,

$$\delta(x) = 0 \;\text{ for } x \neq 0 \qquad \text{and} \qquad \int_{-\infty}^{\infty} \delta(x)\, \mathrm{d}x = 1\,. \tag{18.28}$$

One could imagine $\delta(x)$ as a single pulse at position $x = 0$ that is infinitesimally narrow but still contains finite energy (1). Also remarkable is the impulse function's behavior under scaling along the time (or space) axis (i.e., $\delta(x) \rightarrow \delta(sx)$), with

$$\delta(sx) = \frac{1}{|s|} \cdot \delta(x), \tag{18.29}$$

for $s \neq 0$. Despite the fact that $\delta(x)$ does not exist in physical reality and cannot be plotted (the corresponding plots in Fig. 18.3 are for illustration only), this function is a useful mathematical tool for describing the sampling process, as will be shown.

Sampling with the impulse function

Using the concept of the ideal impulse, the sampling process can be described in a straightforward and intuitive way.[7] If a continuous

[7] The following description is intentionally a bit superficial (in a mathematical sense). See, for example, [43, 128] for more precise coverage of these topics.

function $g(x)$ is multiplied with the impulse function $\delta(x)$, we obtain a new function

$$\bar{g}(x) = g(x) \cdot \delta(x) = \begin{cases} g(0) & \text{for } x = 0, \\ 0 & \text{otherwise.} \end{cases} \tag{18.30}$$

The resulting function $\bar{g}(x)$ consists of a single pulse at position 0 whose height corresponds to the original function value $g(0)$ (at position 0). Thus, by multiplying the function $g(x)$ by the impulse function, we obtain a single discrete sample value of $g(x)$ at position $x = 0$. If the impulse function $\delta(x)$ is shifted by a distance x_0, we can sample $g(x)$ at an arbitrary position $x = x_0$,

$$\bar{g}(x) = g(x) \cdot \delta(x{-}x_0) = \begin{cases} g(x_0) & \text{for } x = x_0, \\ 0 & \text{otherwise.} \end{cases} \tag{18.31}$$

Here $\delta(x{-}x_0)$ is the impulse function shifted by x_0, and the resulting function $\bar{g}(x)$ is zero except at position x_0, where it contains the original function value $g(x_0)$. This relationship is illustrated in Fig. 18.5 for the sampling position $x_0 = 3$.

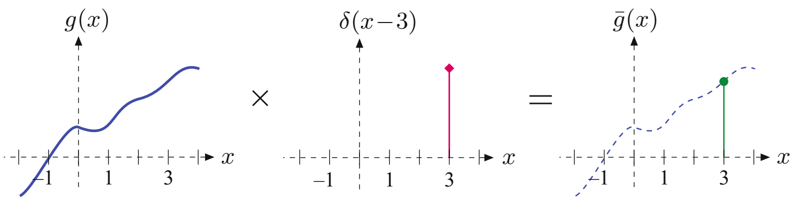

Fig. 18.5 Sampling with the impulse function. The continuous signal $g(x)$ is sampled at position $x_0 = 3$ by multiplying $g(x)$ by a shifted impulse function $\delta(x{-}3)$.

To sample the function $g(x)$ at more than one position simultaneously (e.g., at positions x_1 and x_2), we use two separately shifted versions of the impulse function, multiply $g(x)$ by both of them, and simply add the resulting function values. In this particular case, we get

$$\bar{g}(x) = g(x) \cdot \delta(x{-}x_1) + g(x) \cdot \delta(x{-}x_2) \tag{18.32}$$
$$= g(x) \cdot \bigl[\delta(x{-}x_1) + \delta(x{-}x_2)\bigr] \tag{18.33}$$
$$= \begin{cases} g(x_1) & \text{for } x = x_1, \\ g(x_2) & \text{for } x = x_2, \\ 0 & \text{otherwise.} \end{cases} \tag{18.34}$$

From Eqn. (18.33), sampling a continuous function $g(x)$ at N positions $x_i = 1, 2, \ldots, N$ can thus be described as the sum of the N individual samples, that is,

$$\bar{g}(x) = g(x) \cdot \bigl[\delta(x{-}1) + \delta(x{-}2) + \ldots + \delta(x{-}N)\bigr]$$
$$= g(x) \cdot \sum_{i=1}^{N} \delta(x{-}i)\,. \tag{18.35}$$

The comb function

The sum of shifted impulses $\sum_{i=1}^{N} \delta(x{-}i)$ in Eqn. (18.35) is called a *pulse sequence* or *pulse train*. Extending this sequence to infinity in both directions, we obtain the "comb" or "Shah" function

$$\text{III}(x) = \sum_{i=-\infty}^{\infty} \delta(x - i)\ . \tag{18.36}$$

The process of discretizing a continuous function by taking samples at regular integral intervals can thus be written simply as

$$\bar{g}(x) = g(x) \cdot \text{III}(x), \tag{18.37}$$

that is, as a pointwise multiplication of the original signal $g(x)$ with the comb function $\text{III}(x)$. As Fig. 18.6 illustrates, the function values of $g(x)$ at integral positions $x_i \in \mathbb{Z}$ are transferred to the discrete function $\bar{g}(x_i)$ and ignored at all non-integer positions.

Of course, the sampling interval (i.e., the distance between adjacent samples) is not restricted to 1. To take samples at regular but *arbitrary* intervals τ, the sampling function $\text{III}(x)$ is simply scaled along the time or space axis; that is,

$$\bar{g}(x) = g(x) \cdot \text{III}\left(\tfrac{x}{\tau}\right), \quad \text{for } \tau > 0. \tag{18.38}$$

Effects of sampling in frequency space

Despite the elegant formulation made possible by the use of the comb function, one may still wonder why all this math is necessary to describe a process that appears intuitively to be so simple anyway. The Fourier spectrum gives one answer to this question. Sampling a continuous function has massive—though predictable—effects upon the frequency spectrum of the resulting (discrete) signal. Using the comb function as a formal model for the sampling process makes it relatively easy to estimate and interpret those spectral effects. Similar to the Gaussian (see Sec. 18.1.5), the comb function features the special property that its Fourier transform

$$\text{III}(x) \circ\!\!-\!\!\bullet\ \text{III}(\tfrac{1}{2\pi}\omega) \tag{18.39}$$

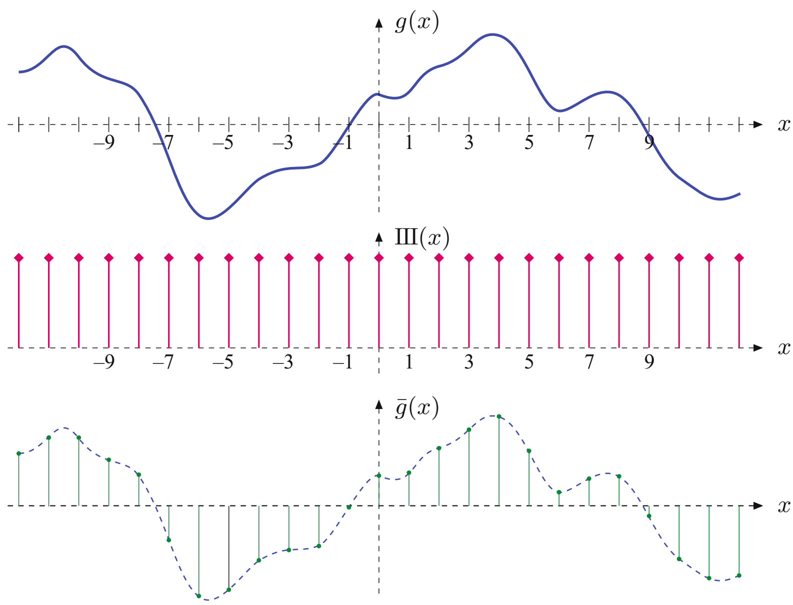

Fig. 18.6
Sampling with the comb function. The original continuous signal $g(x)$ is multiplied by the comb function $\text{III}(x)$. The function value $g(x)$ is transferred to the resulting function $\bar{g}(x)$ only at integral positions $x = x_i \in \mathbb{Z}$ and ignored at all non-integer positions.

is again a comb function (i.e., the same type of function). In general, the Fourier transform of a comb function scaled to an arbitrary sampling interval τ is

$$\text{III}(\tfrac{x}{\tau}) \;\circ\!\!-\!\!\bullet\; \tau \text{III}\left(\tfrac{\tau}{2\pi}\omega\right), \tag{18.40}$$

due to the similarity property of the Fourier transform (Eqn. (18.24)). Figure 18.7 shows two examples of the comb function $\text{III}_\tau(x)$ with sampling intervals $\tau = 1$ and $\tau = 3$ and the corresponding Fourier transforms.

Now, what happens to the Fourier spectrum during discretization, that is, when we multiply a function in signal space by the comb function $\text{III}(\frac{x}{\tau})$? We get the answer by recalling the convolution property of the Fourier transform (Eqn. (18.26)): the product of two functions in one space (signal or frequency space) corresponds to the linear convolution of the transformed functions in the opposite space, and thus

$$g(x) \cdot \text{III}(\tfrac{x}{\tau}) \;\circ\!\!-\!\!\bullet\; G(\omega) * \tau \cdot \text{III}\left(\tfrac{\tau}{2\pi}\omega\right). \tag{18.41}$$

We already know that the Fourier spectrum of the sampling function is a comb function again and therefore consists of a sequence of regularly spaced pulses (Fig. 18.7). In addition, we know that convolving an arbitrary function with the impulse $\delta(x)$ returns the original function; that is, $f(x) * \delta(x) = f(x)$ (see Ch. 5, Sec. 5.3.4). Convolving with a *shifted* pulse $\delta(x-d)$ also reproduces the original function $f(x)$, though shifted by the same distance d:

$$f(x) * \delta(x-d) = f(x-d). \tag{18.42}$$

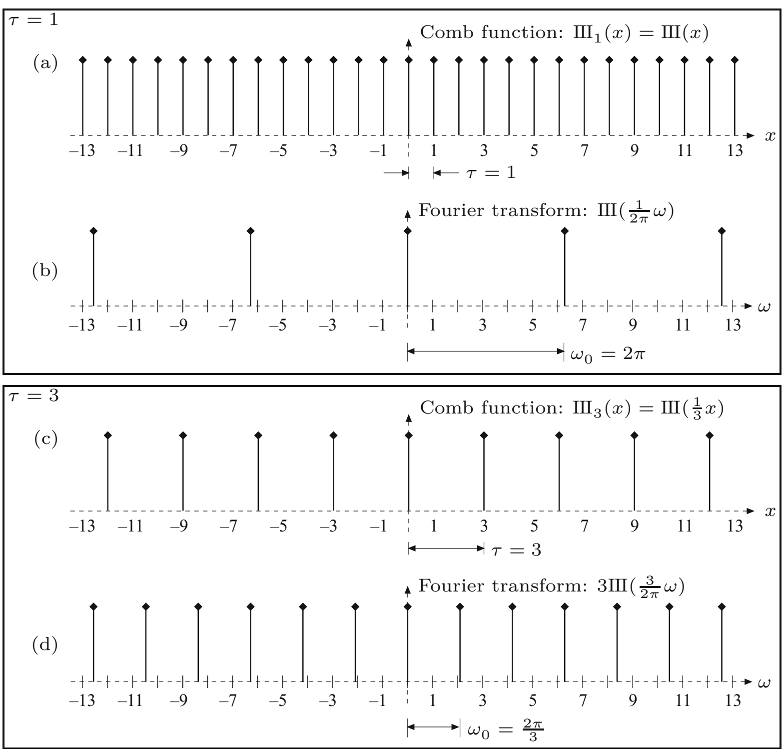

Fig. 18.7
Comb function and its Fourier transform. Comb function $\text{III}_\tau(x)$ for the sampling interval $\tau = 1$ (a) and its Fourier transform. Comb function for $\tau = 3$ (c) and its Fourier transform (d). Note that the actual height of the δ-pulses is undefined and shown only for illustration.

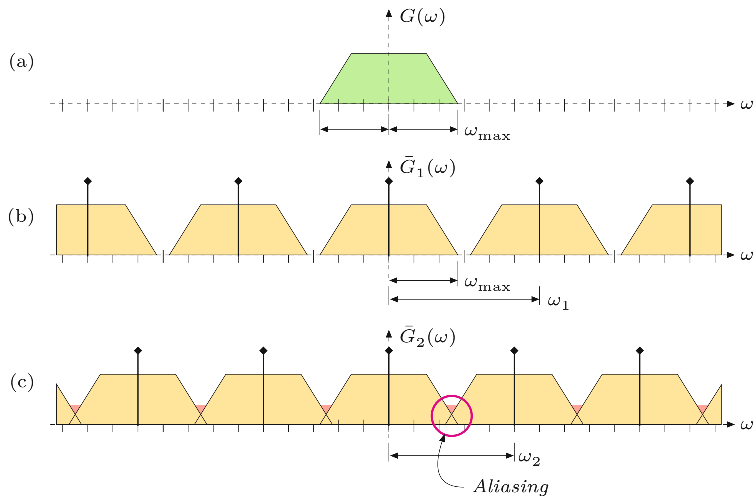

Fig. 18.8
Spectral effects of sampling. The spectrum $G(\omega)$ of the original continuous signal is assumed to be band-limited within the range $\pm\omega_{\max}$ (a). Sampling the signal at a rate (sampling frequency) $\omega_s = \omega_1$ causes the signal's spectrum $G(\omega)$ to be replicated at multiples of ω_1 along the frequency (ω) axis (b). Obviously, the replicas in the spectrum do not overlap as long as $\omega_s > 2\omega_{\max}$. In (c), the sampling frequency $\omega_s = \omega_2$ is less than $2\omega_{\max}$, so there is overlap between the replicas in the spectrum, and frequency components are mirrored at $2\omega_{\max}$ and superimpose the original spectrum. This effect is called "aliasing" because the original spectrum (and thus the original signal) cannot be reproduced from such a corrupted spectrum.

As a consequence, the spectrum $G(\omega)$ of the original continuous signal becomes *replicated* in the Fourier spectrum $\bar{G}(\omega)$ of a sampled signal at every pulse of the sampling function's spectrum; that is, infinitely many times (see Fig. 18.8(a, b))! Thus the resulting Fourier spectrum is repetitive with a period $\frac{2\pi}{\tau}$, which corresponds to the sampling frequency ω_s.

Aliasing and the sampling theorem

As long as the spectral replicas in $\bar{G}(\omega)$ created by the sampling process do not overlap, the original spectrum $G(\omega)$—and thus the original continuous function—can be reconstructed without loss from any isolated replica of $G(\omega)$ in the periodic spectrum $\bar{G}(\omega)$. As we can see in Fig. 18.8, this requires that the frequencies contained in the original signal $g(x)$ be within some upper limit $\omega_{\max}$; that is, the signal contains no components with frequencies greater than $\omega_{\max}$. The maximum allowed signal frequency $\omega_{\max}$ depends upon the sampling frequency ω_s used to discretize the signal, with the requirement

$$\omega_{\max} \leq \tfrac{1}{2} \cdot \omega_s \qquad \text{or} \qquad \omega_s \geq 2 \cdot \omega_{\max}. \tag{18.43}$$

Discretizing a continuous signal $g(x)$ with frequency components in the range $0 \leq \omega \leq \omega_{\max}$ thus requires a sampling frequency ω_s of at least twice the maximum signal frequency $\omega_{\max}$. If this condition is not met, the replicas in the spectrum of the sampled signal overlap (Fig. 18.8(c)) and the spectrum becomes corrupted. Consequently, the original signal cannot be recovered flawlessly from the sampled signal's spectrum. This effect is commonly called "aliasing".

What we just said in simple terms is nothing but the essence of the famous "sampling theorem" formulated by Shannon and Nyquist (see, e.g., [43, p. 256]). It actually states that the sampling frequency must be at least twice the *bandwidth*[8] of the continuous signal to avoid aliasing effects. However, if we assume that a signal's frequency range

[8] This may be surprising at first because it allows a signal with high frequency—but low bandwidth—to be sampled (and correctly recon-

starts at zero, then bandwidth and maximum frequency are the same anyway.

18.2.2 Discrete and Periodic Functions

Assume that we are given a continuous signal $g(x)$ that is periodic with a period of length T. In this case, the corresponding Fourier spectrum $G(\omega)$ is a sequence of thin spectral lines equally spaced at a distance $\omega_0 = 2\pi/T$. As discussed in Sec. 18.1.2, the Fourier spectrum of a periodic function can be represented as a Fourier series and is therefore *discrete*. Conversely, if a continuous signal $g(x)$ is *sampled* at regular intervals τ, then the corresponding Fourier spectrum becomes *periodic* with a period of length $\omega_s = 2\pi/\tau$.

Sampling in signal space thus leads to periodicity in frequency space and vice versa. Figure 18.9 illustrates this relationship and the transition from a continuous nonperiodic signal to a discrete periodic function, which can be represented as a finite vector of numbers and thus easily processed on a computer.

Thus, in general, the Fourier spectrum of a continuous, nonperiodic signal $g(x)$ is also continuous and nonperiodic (Fig. 18.9(a, b)). However, if the signal $g(x)$ is *periodic*, then the corresponding spectrum is *discrete* (Fig. 18.9(c,d)). Conversely, a discrete—but not necessarily periodic—signal leads to a periodic spectrum (Fig. 18.9(e, f)). Finally, if a signal is discrete *and* periodic with M samples per period, then its spectrum is also discrete and periodic with M values (Fig. 18.9(g, h)). Note that the particular signals and spectra in Fig. 18.9 were chosen for illustration only and do not really correspond with each other.

18.3 The Discrete Fourier Transform (DFT)

In the case of a discrete periodic signal, only a finite sequence of M sample values is required to completely represent either the signal $g(u)$ itself or its Fourier spectrum $G(m)$.[9] This representation as finite vectors makes it straightforward to store and process signals and spectra on a computer. What we still need is a version of the Fourier transform applicable to discrete signals.

18.3.1 Definition of the DFT

The discrete Fourier transform is, just like its continuous counterpart, identical in both directions. For a discrete signal $g(u)$ of length M ($u = 0 \ldots M-1$), the forward transform (DFT) is defined as

structed) at a relatively low sampling frequency, even well below the maximum signal frequency. This is possible because one can also use a filter with suitably low bandwidth for reconstructing the original signal. For example, it may be sufficient to strike (i.e., "sample") a church bell (a low-bandwidth oscillatory system with small internal damping) to uniquely generate a sound wave of relatively high frequency.

[9] Notation: We use $g(x)$, $G(\omega)$ for a *continuous* signal or spectrum, respectively, and $g(u)$, $G(m)$ for the *discrete* versions.

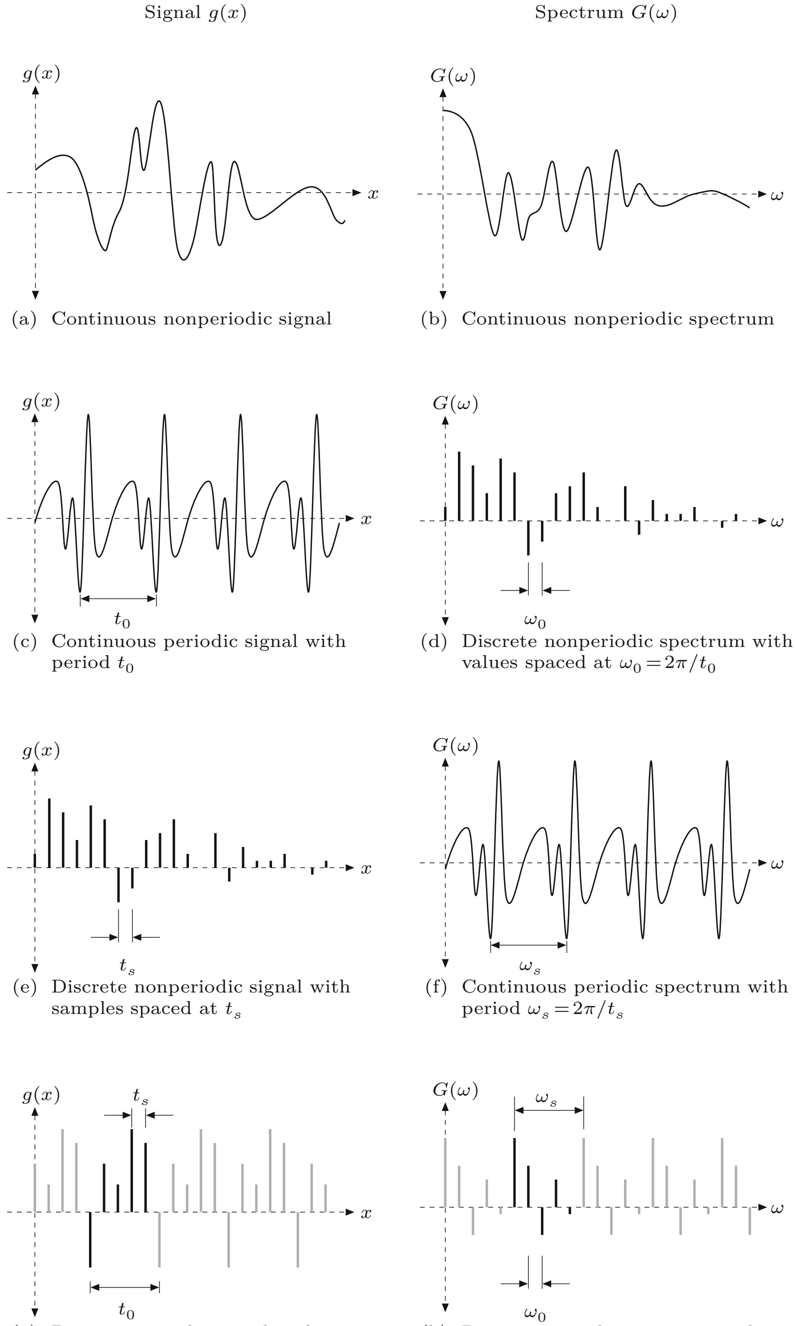

(a) Continuous nonperiodic signal

(b) Continuous nonperiodic spectrum

(c) Continuous periodic signal with period t_0

(d) Discrete nonperiodic spectrum with values spaced at $\omega_0 = 2\pi/t_0$

(e) Discrete nonperiodic signal with samples spaced at t_s

(f) Continuous periodic spectrum with period $\omega_s = 2\pi/t_s$

(g) Discrete periodic signal with samples spaced at $t_0 = t_s M$

(h) Discrete periodic spectrum with values spaced at $\omega_0 = 2\pi/t_0$ and period $\omega_s = 2\pi/t_s = \omega_0 M$

Fig. 18.9
Transition from continuous to discrete periodic functions (illustration only).

$$G(m) = \frac{1}{\sqrt{M}} \sum_{u=0}^{M-1} g(u) \cdot \left[\cos\left(2\pi\frac{mu}{M}\right) - \mathrm{i}\cdot\sin\left(2\pi\frac{mu}{M}\right)\right] \quad (18.44)$$

$$= \frac{1}{\sqrt{M}} \sum_{u=0}^{M-1} g(u) \cdot e^{-\mathrm{i}2\pi\frac{mu}{M}}, \quad (18.45)$$

for $0 \leq m < M$, and the *inverse* transform (DFT^{-1}) is[10]

$$g(u) = \frac{1}{\sqrt{M}} \sum_{m=0}^{M-1} G(m) \cdot \Bigl[\cos\Bigl(2\pi\frac{mu}{M}\Bigr) + \mathrm{i}\cdot\sin\Bigl(2\pi\frac{mu}{M}\Bigr)\Bigr] \tag{18.46}$$

$$= \frac{1}{\sqrt{M}} \sum_{m=0}^{M-1} G(m) \cdot e^{\mathrm{i}2\pi\frac{mu}{M}}, \tag{18.47}$$

for $0 \leq u < M$. Note that both the *signal* $g(u)$ and the discrete *spectrum* $G(m)$ are complex-valued vectors of length M, that is,

$$\begin{aligned} g(u) &= g_{\mathrm{Re}}(u) + \mathrm{i}\cdot g_{\mathrm{Im}}(u),\\ G(m) &= G_{\mathrm{Re}}(m) + \mathrm{i}\cdot G_{\mathrm{Im}}(m), \end{aligned} \tag{18.48}$$

for $u, m = 0, \ldots, M{-}1$. A numerical example for a DFT with $M = 10$ is shown in Fig. 18.10. Converting Eqn. (18.44) from Euler's exponential notation (Eqn. (18.10)) we obtain the discrete Fourier spectrum in component notation as

$$G(m) = \frac{1}{\sqrt{M}} \cdot \sum_{u=0}^{M-1} \Bigl[\underbrace{g_{\mathrm{Re}}(u) + \mathrm{i}\cdot g_{\mathrm{Im}}(u)}_{g(u)}\Bigr] \cdot \Bigl[\underbrace{\cos\bigl(2\pi\tfrac{mu}{M}\bigr)}_{\boldsymbol{C}_m^M(u)} - \mathrm{i}\cdot\underbrace{\sin\bigl(2\pi\tfrac{mu}{M}\bigr)}_{\boldsymbol{S}_m^M(u)}\Bigr], \tag{18.49}$$

where we denote as $\boldsymbol{C}_m^M$ and $\boldsymbol{S}_m^M$ the discrete (cosine and sine) basis functions, as described in the next section. Applying the usual complex multiplication,[11] we obtain the real and imaginary parts of the discrete Fourier spectrum as

$$G_{\mathrm{Re}}(m) = \frac{1}{\sqrt{M}} \cdot \sum_{u=0}^{M-1} g_{\mathrm{Re}}(u) \cdot \boldsymbol{C}_m^M(u) + g_{\mathrm{Im}}(u) \cdot \boldsymbol{S}_m^M(u), \tag{18.50}$$

$$G_{\mathrm{Im}}(m) = \frac{1}{\sqrt{M}} \cdot \sum_{u=0}^{M-1} g_{\mathrm{Im}}(u) \cdot \boldsymbol{C}_m^M(u) - g_{\mathrm{Re}}(u) \cdot \boldsymbol{S}_m^M(u), \tag{18.51}$$

for $m = 0, \ldots, M{-}1$. Analogously, the *inverse* DFT in Eqn. (18.46) expands to

$$g_{\mathrm{Re}}(u) = \frac{1}{\sqrt{M}} \cdot \sum_{m=0}^{M-1} G_{\mathrm{Re}}(m) \cdot \boldsymbol{C}_u^M(m) - G_{\mathrm{Im}}(m) \cdot \boldsymbol{S}_u^M(m), \tag{18.52}$$

$$g_{\mathrm{Im}}(u) = \frac{1}{\sqrt{M}} \cdot \sum_{m=0}^{M-1} G_{\mathrm{Im}}(m) \cdot \boldsymbol{C}_u^M(m) + G_{\mathrm{Re}}(m) \cdot \boldsymbol{S}_u^M(m), \tag{18.53}$$

for $u = 0, \ldots, M{-}1$.

[10] Compare these definitions with the corresponding expressions for the *continuous* forward and inverse Fourier transforms in Eqns. (18.19) and (18.20), respectively.

[11] See also Sec. A.3 in the Appendix.

Fig. 18.10
Complex-valued result of the DFT for a signal of length $M = 10$ (example). In the discrete Fourier transform (DFT), both the original signal $g(u)$ and its spectrum $G(m)$ are complex-valued vectors of length M; $*$ indicates values with $|G(m)| < 10^{-15}$.

u	$g(u)$ Re	$g(u)$ Im		$G(m)$ Re	$G(m)$ Im	m
0	1.0000	0.0000		14.2302	0.0000	0
1	3.0000	0.0000	DFT	−5.6745	−2.9198	1
2	5.0000	0.0000	⟶	*0.0000	*0.0000	2
3	7.0000	0.0000		−0.0176	−0.6893	3
4	9.0000	0.0000		*0.0000	*0.0000	4
5	8.0000	0.0000		0.3162	0.0000	5
6	6.0000	0.0000		*0.0000	*0.0000	6
7	4.0000	0.0000	DFT^{-1}	−0.0176	0.6893	7
8	2.0000	0.0000	⟵	*0.0000	*0.0000	8
9	0.0000	0.0000		−5.6745	2.9198	9

18.3.2 Discrete Basis Functions

The inverse DFT (Eqn. (18.46)) performs the decomposition of the discrete function $g(u)$ into a finite sum of M discrete cosine and sine functions ($\boldsymbol{C}_m^M$, $\boldsymbol{S}_m^M$) whose weights (or "amplitudes") are determined by the DFT coefficients in $G(m)$. Each of these 1D basis functions (first used in Eqn. (18.49)),

$$\boldsymbol{C}_m^M(u) = \boldsymbol{C}_u^M(m) = \cos\bigl(2\pi \tfrac{mu}{M}\bigr), \tag{18.54}$$

$$\boldsymbol{S}_m^M(u) = \boldsymbol{S}_u^M(m) = \sin\bigl(2\pi \tfrac{mu}{M}\bigr), \tag{18.55}$$

is periodic with M and has a discrete frequency (wave number) m, which corresponds to the angular frequency

$$\omega_m = 2\,\pi \cdot \frac{m}{M}. \tag{18.56}$$

For example, Figs. 18.11 and 18.12 show the discrete basis functions (with integer ordinate values $u \in \mathbb{Z}$) for the DFT of length $M = 8$ as well as their continuous counterparts (with ordinate values $x \in \mathbb{R}$).

For wave number $m = 0$, the cosine function $\boldsymbol{C}_0^M(u)$ (Eqn. (18.54)) has the constant value 1. The corresponding DFT coefficient $G_{\text{Re}}(0)$—the real part of $G(0)$—thus specifies the constant part of the signal or the average value of the signal $g(u)$ in Eqn. (18.52). In contrast, the zero-frequency sine function $\boldsymbol{S}_0^M(u)$ is zero for any value of u and thus cannot contribute anything to the signal. The corresponding DFT coefficients $G_{\text{Im}}(0)$ in Eqn. (18.52) and $G_{\text{Re}}(0)$ in Eqn. (18.53) are therefore of no relevance. For a real-valued signal (i.e., $g_{\text{Im}}(u) = 0$ for all u), the coefficient $G_{\text{Im}}(0)$ in the corresponding Fourier spectrum must also be zero.

As seen in Fig. 18.11, the wave number $m = 1$ relates to a cosine or sine function that performs exactly one full cycle over the signal length $M = 8$. Similarly, the wave numbers $m = 2, \ldots, 7$ correspond to $2, \ldots, 7$ complete cycles over the signal length M (see Figs. 18.11 and 18.12).

18.3.3 Aliasing Again!

A closer look at Figs. 18.11 and 18.12 reveals an interesting fact: the sampled (discrete) cosine and sine functions for $m = 3$ and $m = 5$ are *identical*, although their continuous counterparts are different! The same is true for the frequency pairs $m = 2, 6$ and $m = 1, 7$. What we

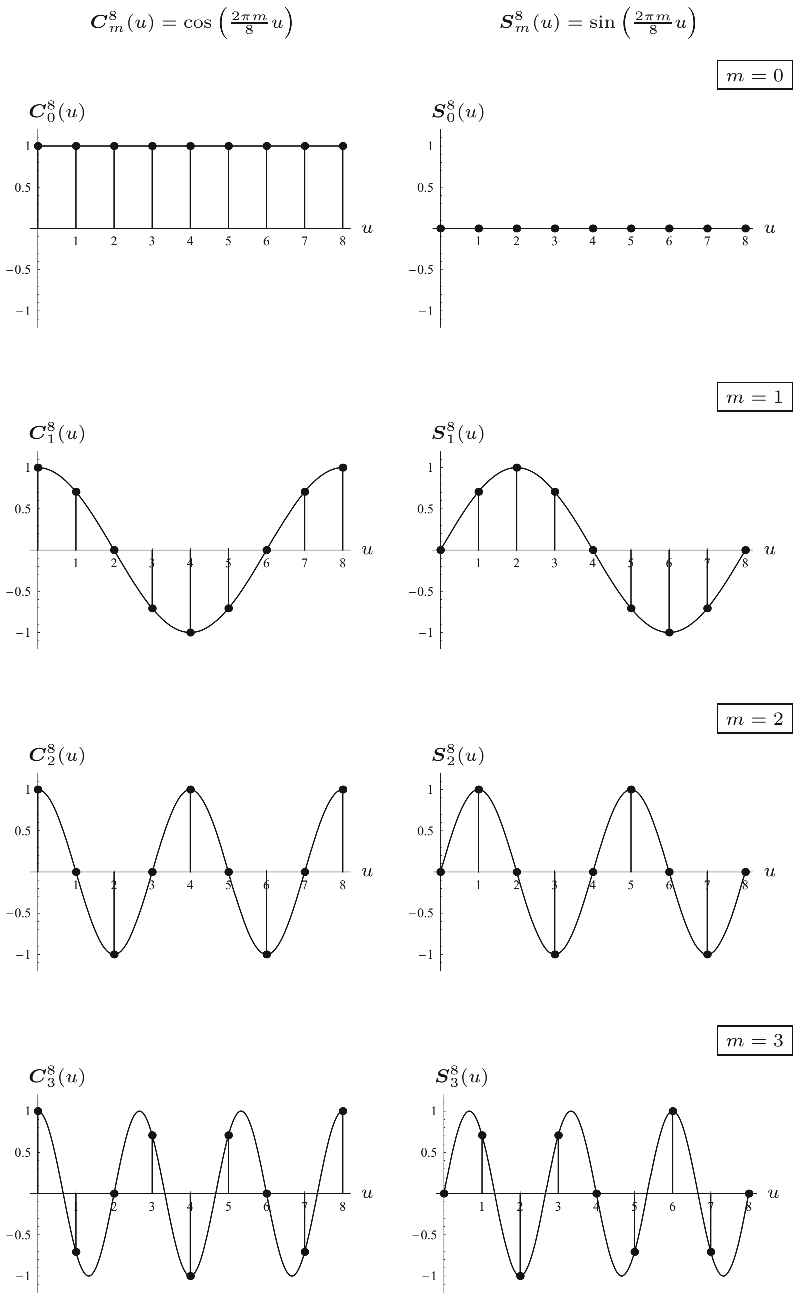

Fig. 18.11
Discrete basis functions $C_m^M(u)$ and $S_m^M(u)$ for the signal length $M = 8$ and wave numbers $m = 0, \ldots, 3$. Each plot shows both the discrete function (round dots) and the corresponding continuous function.

see here is another manifestation of the sampling theorem—which we had originally encountered (Sec. 18.2.1) in frequency space—in *signal space*. Obviously, $m = 4$ is the maximum frequency component that can be represented by a discrete signal of length $M = 8$. Any discrete function with a higher frequency ($m = 5, \ldots, 7$ in this case) has an identical counterpart with a lower wave number and thus cannot be reconstructed from the sampled signal (see also Fig. 18.13)!

If a continuous signal is sampled at a regular distance τ, the corresponding Fourier spectrum is repeated at multiples of $\omega_s = 2\pi/\tau$,

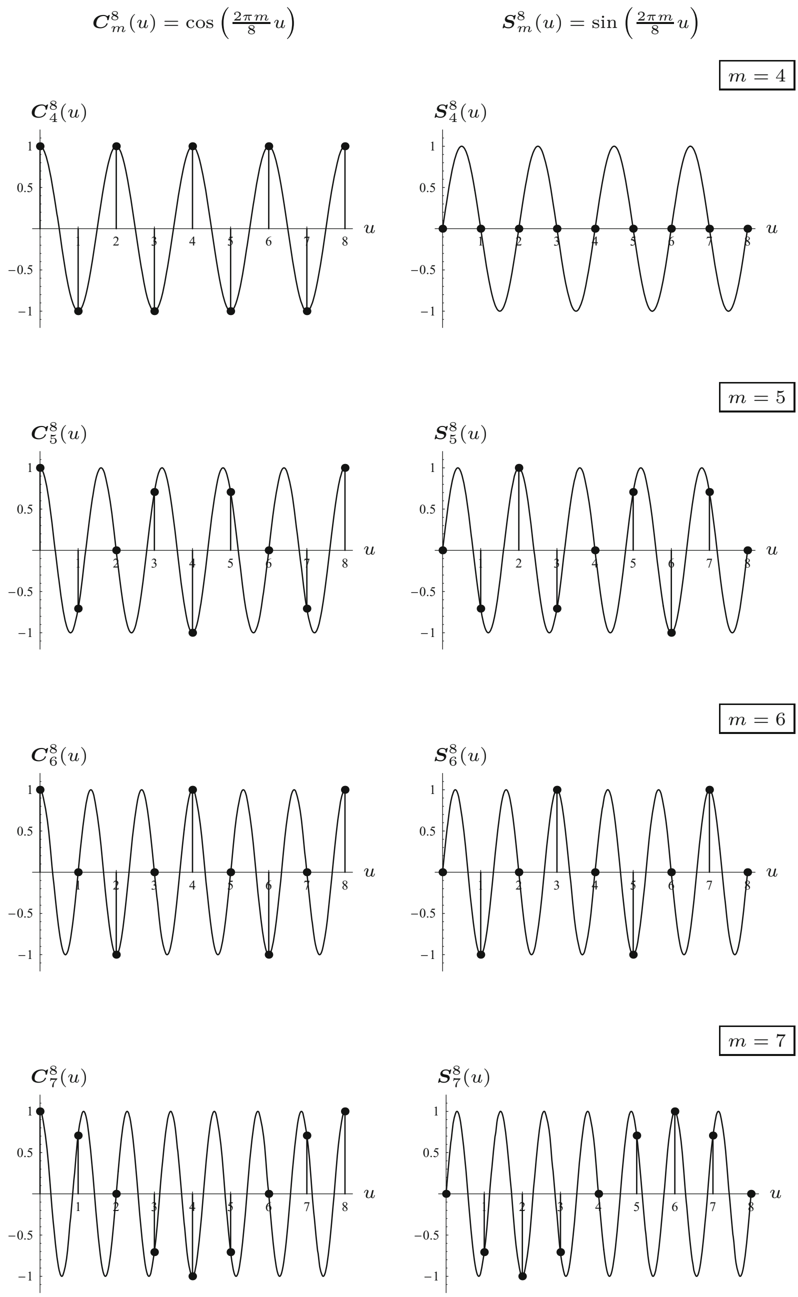

Fig. 18.12 Discrete basis functions (continued). Signal length $M = 8$ and wave numbers $m = 4, \ldots, 7$. Notice that, for example, the discrete functions for $m = 5$ and $m = 3$ (Fig. 18.11) are identical because $m = 4$ is the maximum wave number that can be represented in a discrete spectrum of length $M = 8$.

as we have shown earlier (Fig. 18.8). In the discrete case, the spectrum is periodic with length M. Since the Fourier spectrum of a real-valued signal is symmetric about the origin (Eqn. (18.21)), there is for every coefficient with wave number m an equal-sized duplicate with wave number $-m$. Thus the spectral components appear pairwise and mirrored at multiples of M; that is,

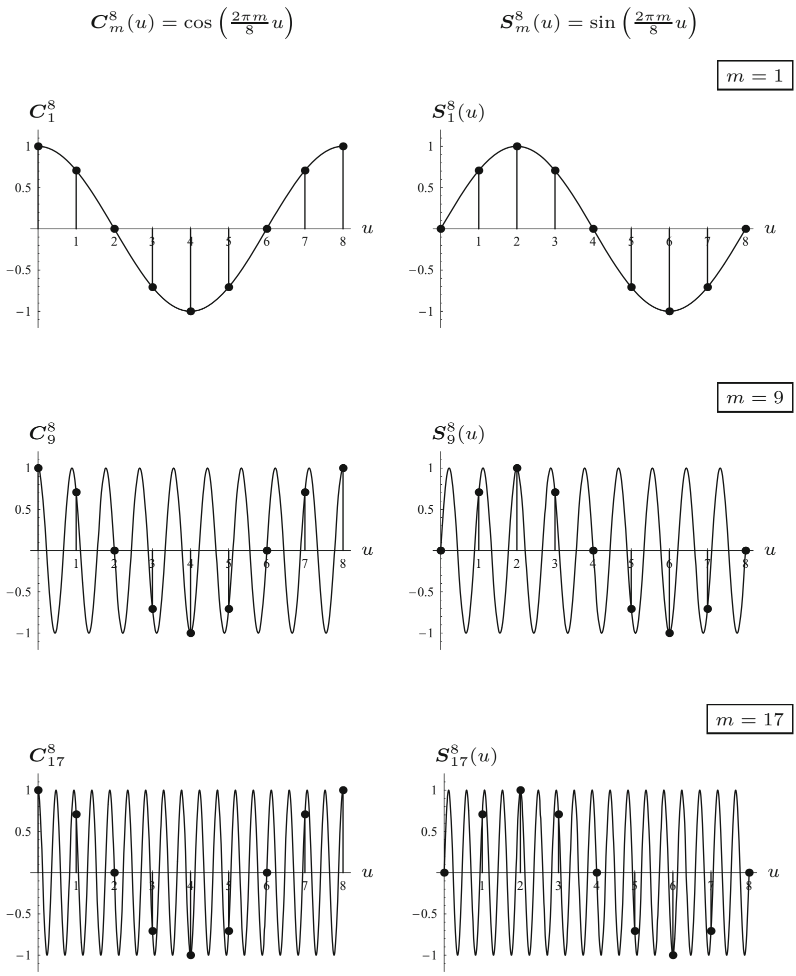

Fig. 18.13
Aliasing in signal space. For the signal length $M = 8$, the discrete cosine and sine basis functions for the wave numbers $m = 1, 9, 17, \ldots$ (round dots) are all identical. The sampling frequency itself corresponds to the wave number $m = 8$.

$$
\begin{aligned}
|G(m)| &= |G(M-m)| &= |G(M+m)| \\
&= |G(2M-m)| &= |G(2M+m)| \\
&\ldots \\
&= |G(kM-m)| &= |G(kM+m)|,
\end{aligned}
\tag{18.57}
$$

for all $k \in \mathbb{Z}$. If the original continuous signal contains "energy" at the frequencies

$$\omega_m > \omega_{M/2}$$

(i.e., signal components with wave numbers $m > M/2$), then, according to the sampling theorem, the overlapping parts of the spectra are superimposed in the resulting periodic spectrum of the discrete signal.

18.3.4 Units in Signal and Frequency Space

The relation between the units in signal and frequency space and the interpretation of wave numbers m is a common cause of confusion. While the discrete signal and its spectrum are simple numerical vectors and units of measurement are irrelevant for computing the DFT itself, it is nevertheless important to understand how the coordinates in the spectrum relate to physical dimensions in the real world.

Clearly, every complex-valued spectral coefficient $G(m)$ corresponds to one pair of cosine and sine functions with a particular frequency in signal space. Assume a continuous signal is sampled at M consecutive positions spaced at τ (an interval in time or distance in space). The *wave number* $m = 1$ then corresponds to the *fundamental period* of the discrete signal (which is now assumed to be periodic) with a period of length $M\tau$; that is, to the *frequency*

$$f_1 = \frac{1}{M\tau}\,. \tag{18.58}$$

In general, the wave number m of a discrete spectrum relates to the physical frequency as

$$f_m = m\frac{1}{M\tau} = m \cdot f_1 \tag{18.59}$$

for $0 \leq m < M$, which is equivalent to the angular frequency

$$\omega_m = 2\pi f_m = m\frac{2\pi}{M\tau} = m \cdot \omega_1. \tag{18.60}$$

Obviously then, the sampling frequency $f_s = 1/\tau = M \cdot f_1$ corresponds to the wave number $m_s = M$. As expected, the maximum nonaliased wave number in the spectrum is

$$m_{\max} = \frac{M}{2} = \frac{m_s}{2}\,, \tag{18.61}$$

that is, half the sampling frequency index m_s.

Example 1: time-domain signal

We assume for this example that $g(u)$ is a signal in the time domain (e.g., a discrete sound signal) that contains $M = 500$ sample values taken at regular intervals $\tau = 1$ ms $= 10^{-3}$ s. Thus the sampling frequency is $f_s = 1/\tau = 1000$ Hertz (cycles per second) and the total duration (fundamental period) of the signal is $M\tau = 0.5$ s.

The signal is implicitly periodic, and from Eqn. (18.58) we obtain its fundamental frequency as $f_1 = \frac{1}{500 \cdot 10^{-3}} = \frac{1}{0.5} = 2$ Hertz. The wave number $m = 2$ in this case corresponds to a real frequency $f_2 = 2f_1 = 4$ Hertz, $f_3 = 6$ Hertz, etc. The maximum frequency that can be represented by this discrete signal without aliasing is $f_{\max} = \frac{M}{2} f_1 = \frac{1}{2\tau} = 500$ Hertz, exactly half the sampling frequency f_s.

Example 2: space-domain signal

Assume we have a 1D print pattern with a resolution (i.e., spatial sampling frequency) of 120 dots per cm, which equals approximately 300 dots per inch (dpi) and a total signal length of $M = 1800$ samples. This corresponds to a spatial sampling interval of $\tau = 1/120$ cm $\approx 83\,\mu$m and a physical signal length of $(1800/120)$ cm $= 15$ cm.

The fundamental frequency of this signal (again implicitly assumed to be periodic) is $f_1 = \frac{1}{15}$, expressed in cycles per cm. The sampling frequency is $f_s = 120$ cycles per cm and thus the maximum signal frequency is $f_{\max} = \frac{f_s}{2} = 60$ cycles per cm. The maximum signal frequency specifies the finest structure ($\frac{1}{60}$ cm) that can be reproduced by this print raster.

18.3.5 Power Spectrum

The *magnitude* of the complex-valued Fourier spectrum,

$$|G(m)| = \sqrt{G^2_{\mathrm{Re}}(m) + G^2_{\mathrm{Im}}(m)}, \tag{18.62}$$

is commonly called the "power spectrum" of a signal. It specifies the energy that individual frequency components in the spectrum contribute to the signal. The power spectrum is real-valued and positive and thus often used for graphically displaying the results of Fourier transforms (see also Ch. 19, Sec. 19.2).

Since all phase information is lost in the power spectrum, the original signal cannot be reconstructed from the power spectrum alone. However, because of the missing phase information, the power spectrum is insensitive to shifts of the original signal and can thus be efficiently used for comparing signals. To be more precise, the power spectrum of a circularly shifted signal is identical to the power spectrum of the original signal. Thus, given a discrete periodic signal $g_1(u)$ of length M and a second signal $g_2(u)$ shifted by some offset d, such that

$$g_2(u) = g_1(u-d) \tag{18.63}$$

the corresponding power spectra are the same, that is,

$$|G_2(m)| = |G_1(m)|, \tag{18.64}$$

although in general the complex-valued spectra $G_1(m)$ and $G_2(m)$ are different. Furthermore, from the symmetry property of the Fourier spectrum, it follows that

$$|G(m)| = |G(-m)|, \tag{18.65}$$

for real-valued signals $g(u) \in \mathbb{R}$.

18.4 Implementing the DFT

18.4.1 Direct Implementation

Based on the definitions in Eqns. (18.50) and (18.51) the DFT can be directly implemented, as shown in Prog. 18.1. The main method `DFT()` transforms a signal vector of arbitrary length M (not necessarily a power of 2). It requires roughly M^2 operations (multiplications and additions); that is, the time complexity of this DFT algorithm is $\mathcal{O}(M^2)$.

One way to improve the efficiency of the DFT algorithm is to use lookup tables for the sin and cos functions (which are relatively "expensive" to compute) since only function values for a set of M different angles φ_m are ever needed. The angles $\varphi_m = 2\pi\frac{m}{M}$ corresponding to $m = 0, \ldots, M-1$ are evenly distributed over the full 360° circle. Any integral multiple $\varphi_m \cdot u$ (for $u \in \mathbb{Z}$) can only fall onto one of these angles again because

Prog. 18.1
Direct implementation of the DFT based on the definition in Eqns. (18.50) and (18.51). The method `DFT()` returns a complex-valued vector with the same length as the complex-valued input (signal) vector `g`. This method implements both the forward and the inverse transforms, controlled by the Boolean parameter `forward`. The class `Complex` (bottom) defines the structure of the complex-valued vector elements.

```
class Complex {
    double re, im;
    Complex(double re, double im) { //constructor method
        this.re = re;
        this.im = im;
    }
}
```

```
Complex[] DFT(Complex[] g, boolean forward) {
    int M = g.length;
    double s = 1 / Math.sqrt(M); //common scale factor
    Complex[] G = new Complex[M];
    for (int m = 0; m < M; m++) {
        double sumRe = 0;
        double sumIm = 0;
        double phim = 2 * Math.PI * m / M;
        for (int u = 0; u < M; u++) {
            double gRe = g[u].re;
            double gIm = g[u].im;
            double cosw = Math.cos(phim * u);
            double sinw = Math.sin(phim * u);
            if (!forward) // inverse transform
                sinw = -sinw;
            // complex multiplication: [gRe+i·gIm]·[cos(ω)+i·sin(ω)]
            sumRe += gRe * cosw + gIm * sinw;
            sumIm += gIm * cosw - gRe * sinw;
        }
        G[m] = new Complex(s * sumRe, s * sumIm);
    }
    return G;
}
```

$$\varphi_m \cdot u = 2\pi \tfrac{mu}{M} \;\equiv\; \tfrac{2\pi}{M} \cdot (\underbrace{mu \mod M}_{0 \leq k < M}) = 2\pi \tfrac{k}{M} = \varphi_k, \qquad (18.66)$$

where mod denotes the "modulus" operator.[12] Thus we can set up two constant tables (floating-point arrays) W_C and W_S of size M with the values

$$\mathsf{W}_C(k) \leftarrow \cos(\omega_k) = \cos\left(2\pi \tfrac{k}{M}\right), \qquad (18.67)$$

$$\mathsf{W}_S(k) \leftarrow \sin(\omega_k) = \sin\left(2\pi \tfrac{k}{M}\right), \qquad (18.68)$$

for $0 \leq k < M$. For computing the DFT, the necessary cosine and sine values (Eqn. (18.49)) can be read from these tables as

$$\boldsymbol{C}_k^M(u) = \cos\left(2\pi \tfrac{mu}{M}\right) \equiv \mathsf{W}_C(mu \bmod M), \qquad (18.69)$$

$$\boldsymbol{S}_k^M(u) = \sin\left(2\pi \tfrac{mu}{M}\right) \equiv \mathsf{W}_S(mu \bmod M), \qquad (18.70)$$

for arbitrary values of $m, u \in \mathbb{Z}$, without any additional computation. The necessary modification of the `DFT()` method in Prog. 18.1 is left as an exercise (Exercise 18.5).

Despite this significant improvement, the direct implementation of the DFT remains computationally intensive. As a matter of fact,

[12] See also Sec. F.1.2 in the Appendix.

it has been impossible for a long time to compute this form of DFT in sufficiently short time on off-the-shelf computers, and this is still true today for many real applications.

18.4.2 Fast Fourier Transform (FFT)

Fortunately, for computing the DFT in practice, fast algorithms exist that lay out the sequence of computations in such a way that intermediate results are only computed once and optimally reused many times. This "fast Fourier transform", which exists in many variations, generally reduces the time complexity of the computation from $\mathcal{O}(M^2)$ to $\mathcal{O}(M \log_2 M)$. The benefits are substantial, in particular for longer signals. For example, with a signal of length $M = 10^3$, the FFT leads to a speedup by a factor of 100 over the direct DFT implementation and an impressive gain of 10,000 times for a signal of length $M = 10^6$. Since its invention, the FFT has therefore become an indispensable tool in almost any application of spectral signal analysis [34].

Most FFT algorithms, including the one described in the famous publication by Cooley and Tukey in 1965 (see [88, p. 156] for a historic overview), are designed for signals of length $M = 2^k$ (i.e., powers of 2). However, FFT algorithms have also been developed for other lengths, including several small prime numbers [25]. Efficient Java implementations are available, for example, as part of the *JTransform* library[13] by Piotr Wendykier [255] or the *Apache Commons Math* libary.[14]

It is important to remember, though, that the DFT and FFT compute exactly the *same* result and the FFT is only a special—though ingenious—method for *implementing* the discrete Fourier transform (Eqn. (18.44)).

18.5 Exercises

Exercise 18.1. Calculate the values of the cosine function $f(x) = \cos(\omega x)$ with angular frequency $\omega = 5$ for the positions $x = -3, -2, \ldots, 2, 3$. What is the length of this function's period?

Exercise 18.2. Determine the phase angle φ of the function $f(x) = A \cdot \cos(\omega x) + B \cdot \sin(\omega x)$ for $A = -1$ and $B = 2$.

Exercise 18.3. Calculate the real part, the imaginary part, and the magnitude of the complex value $z = 1.5 \cdot e^{-\mathrm{i}\,2.5}$.

Exercise 18.4. A 1D optical scanner for sampling film transparencies is supposed to resolve image structures with a precision of 4,000 dpi. What spatial distance (in mm) between samples is required such that no aliasing occurs?

Exercise 18.5. Modify the direct implementation of the 1D DFT given in Prog. 18.1 by using lookup tables for the cos and sin functions as described in Eqns. (18.69)–(18.70).

[13] http://sites.google.com/site/piotrwendykier/software/jtransforms.

[14] http://commons.apache.org/math/ (class `FastFourierTransformer`).