Chapter 1

Introduction

1.1 A Bit of History

1.2 An Application

1.3 Problems

The material essential for a student's mathematical laboratory is very simple. Each student should have a copy of Barlow's tables of squares, etc., a copy of Crelle's Calculation Tables, *and a seven-place table of logarithms. Further it is necessary to provide a stock of computing paper, . . . and lastly, a stock of computing forms for practical Fourier analysis. With this modest apparatus nearly all of the computations hereafter described may be performed, although time and labour may often be saved by the use of multiplying and adding machines when these are available.*

- E. T. Whittaker and G. Robinson

The Calculus of Observations, 1924

1.1. A Bit of History

Rather than beginning with the definition of the discrete Fourier transform on the first page of this book, we thought that a few pages of historical and practical introduction might be useful. Some valuable perspective comes with the realization that the DFT was not discovered ten years ago, nor was it invented with the fast Fourier transform (FFT) thirty years ago. It has a fascinating history, spanning over two centuries, that is closely associated with the development of applied mathematics and numerical analysis. Therefore, this chapter will be devoted, in part, to the history of the DFT. However, before diving into the technicalities of the subject, there is insight to be gained from seeing the DFT in action on a specific problem. Therefore, a few pages of this chapter will also be spent extracting as much understanding as possible from a very practical example. With these excursions into history and applications behind us, the path to the DFT should be straight and clear.

Let's begin with some history. Fourier analysis is over 200 years old, and its history is filled with both controversy and prodigious feats. Interwoven throughout it is the thread of discrete or practical Fourier analysis which is most pertinent to this book. In order to appreciate the complete history, one must retreat some 60 years prior to the moment in 1807 when Jean Baptiste Joseph Fourier presented the first version of his paper on the theory of heat conduction to the Paris Academy. The year 1750 is a good starting point: George II was king of England and the American colonies were in the midst of the French and Indian War; Voltaire, Rousseau, and Kant were writing in Europe; Bach had just died, Mozart was soon to be born; and the calculus of Newton and Leibnitz, published 75 years earlier, was enabling the creation of powerful new theories of celestial and continuum mechanics.

There were two outstanding problems of the day that focused considerable mathematical energy, and formed the seeds that ultimately became Fourier analysis. The first problem was to describe the vibration of a taut string anchored at both ends (or equivalently the propagation of sound in an elastic medium). Remarkably, the wave equation as we know it today had already been formulated, and the mathematicians Jean d'Alembert, Leonhard Euler, Daniel Bernoulli, and Joseph-Louis Lagrange had proposed methods of solution around 1750. Bernoulli's solution took the form of a trigonometric series

$$y = A \sin x \cos at + B \sin 2x \cos 2at + \cdots,$$

in which x is the spatial coordinate and t is the time variable. This solution already anticipated the continuous form of a Fourier series. It appears that both Euler and Lagrange actually *discretized* the vibrating string problem by imagining the string to consist of a finite number of connected particles. The solution of this discrete problem required finding *samples* of the function that describes the displacement of the string. A page from Lagrange's work on this problem (see Figure 1.1), published in 1759 [90], contains all of the ingredients of what we would today call a discrete Fourier sine series.

The second problem which nourished the roots of Fourier analysis, particularly in its discrete form, was that of determining the orbits of celestial bodies. Euler, Lagrange, and Alexis Claude Clairaut made fundamental contributions by proposing that data taken from observations be approximated by linear combinations of periodic functions. The calculation of the coefficients in these trigonometric expansions led to a computation that we would call a discrete Fourier transform. In fact, a paper

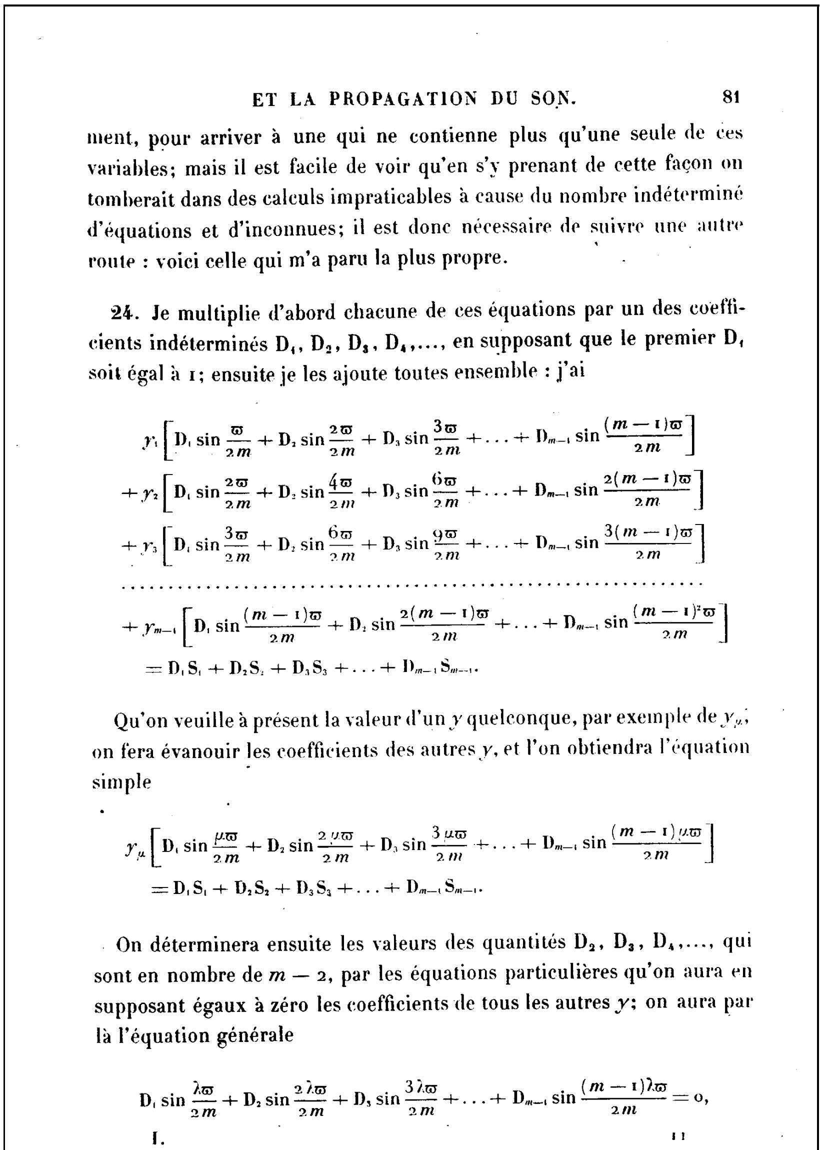

ET LA PROPAGATION DU SON. 81

ment, pour arriver à une qui ne contienne plus qu'une seule de ces variables; mais il est facile de voir qu'en s'y prenant de cette façon on tomberait dans des calculs impraticables à cause du nombre indéterminé d'équations et d'inconnues; il est donc nécessaire de suivre une autre route : voici celle qui m'a paru la plus propre.

24. Je multiplie d'abord chacune de ces équations par un des coefficients indéterminés D_1, D_2, D_3, D_4,..., en supposant que le premier D_1 soit égal à 1; ensuite je les ajoute toutes ensemble : j'ai

$$
\begin{aligned}
&y_1\left[D_1 \sin\frac{\varpi}{2m} + D_2 \sin\frac{2\varpi}{2m} + D_3 \sin\frac{3\varpi}{2m} + \ldots + D_{m-1} \sin\frac{(m-1)\varpi}{2m}\right]\\
+&y_2\left[D_1 \sin\frac{2\varpi}{2m} + D_2 \sin\frac{4\varpi}{2m} + D_3 \sin\frac{6\varpi}{2m} + \ldots + D_{m-1} \sin\frac{2(m-1)\varpi}{2m}\right]\\
+&y_3\left[D_1 \sin\frac{3\varpi}{2m} + D_2 \sin\frac{6\varpi}{2m} + D_3 \sin\frac{9\varpi}{2m} + \ldots + D_{m-1} \sin\frac{3(m-1)\varpi}{2m}\right]\\
&\ldots\ldots\ldots\ldots\ldots\ldots\ldots\ldots\ldots\ldots\ldots\ldots\ldots\ldots\\
+&y_{m-1}\left[D_1 \sin\frac{(m-1)\varpi}{2m} + D_2 \sin\frac{2(m-1)\varpi}{2m} + \ldots + D_{m-1} \sin\frac{(m-1)^2\varpi}{2m}\right]\\
&= D_1S_1 + D_2S_2 + D_3S_3 + \ldots + D_{m-1}S_{m-1}.
\end{aligned}
$$

Qu'on veuille à présent la valeur d'un y quelconque, par exemple de y_μ; on fera évanouir les coefficients des autres y, et l'on obtiendra l'équation simple

$$
\begin{aligned}
&y_\mu\left[D_1 \sin\frac{\mu\varpi}{2m} + D_2 \sin\frac{2\mu\varpi}{2m} + D_3 \sin\frac{3\mu\varpi}{2m} + \ldots + D_{m-1} \sin\frac{(m-1)\mu\varpi}{2m}\right]\\
&= D_1S_1 + D_2S_2 + D_3S_3 + \ldots + D_{m-1}S_{m-1}.
\end{aligned}
$$

On déterminera ensuite les valeurs des quantités D_2, D_3, D_4,..., qui sont en nombre de $m-2$, par les équations particulières qu'on aura en supposant égaux à zéro les coefficients de tous les autres y; on aura par là l'équation générale

$$
D_1 \sin\frac{\lambda\varpi}{2m} + D_2 \sin\frac{2\lambda\varpi}{2m} + D_3 \sin\frac{3\lambda\varpi}{2m} + \ldots + D_{m-1} \sin\frac{(m-1)\lambda\varpi}{2m} = 0,
$$

I. 11

FIG. 1.1. *This page from* The Nature and Propagation of Sound, *written by Lagrange in* 1759 [90], *dealt with the solution of the vibrating string problem. Lagrange assumed that the string consists of* $m-1$ *particles whose displacements,* y_n, *are sums of sine functions of various frequencies. This representation is essentially a discrete Fourier sine transform. Note that* ϖ *means* π.

published in 1754 by Clairaut contains what has been described as the first explicit formula for the DFT [74].

The story follows two paths at the beginning of the nineteenth century. Not surprisingly, we might call one path *continuous* and the other *discrete*. On the continuous path, in 1807 Fourier presented his paper before the Paris Academy, in which he asserted that an arbitrary function can be represented as an infinite series of sines and cosines. The paper elicited only mild encouragement from the judges and the suggestion that Fourier refine his work and submit it for the grand prize in 1812. The Academy's panel of judges for the grand prize included Lagrange, Laplace, and Legendre, and they did award Fourier the grand prize in 1812, but not without reservations. Despite the fact that Euler and Bernoulli had introduced trigonometric representations of functions, and that Lagrange had already produced what we would call a Fourier series solution to the wave equation, Fourier's broader claim that an *arbitrary function* could be given such a representation aroused skepticism, if not outrage. The grand prize came with the deflating assessment that

> the way in which the author arrives at his equations is not exempt from difficulties, and his analysis still leaves something to be desired, be it in generality, or be it even in rigor.

Historians are divided over how much credit is due to Lagrange for the discovery of Fourier series. One historian [154] has remarked that

> certainly Lagrange could find nothing new in Fourier's theorem except the sweeping generality of its statement and the preposterous legerdemain advanced as a proof.

Fourier's work on heat conduction and the supporting theory of trigonometric series culminated in the *Théorie analytique de la chaleur* [60], which was published in 1822. Since that year, the work of Fourier has been given more generous assessments. Clerk Maxwell called it a "great mathematical poem." The entire career of William Thompson (Lord Kelvin) was impacted by Fourier's theory of heat, and he proclaimed that

> it is difficult to say whether their uniquely original quality, or their transcendently intense mathematical interest, or their perennially important instructiveness for the physical sciences, is most to be praised.

Regardless of the originality and rigor of the work when it was first presented, there can be little doubt that for almost 200 years, the subject of Fourier analysis has changed the entire landscape of mathematics and its applications [156], [162].

The continuous path did not end with Fourier's work. The remainder of the nineteenth century was an incubator of mathematical thought in Europe. Some of the greatest mathematicians of the period such as Poisson, Dirichlet, and Riemann advanced the theory of trigonometric series and addressed the challenging questions of convergence. The campaign continued into the twentieth century when Lebesgue, armed with his new theory of integration, was able to produce even more general statements about the convergence of trigonometric series.

Let's return to the beginning of the nineteenth century and follow the second path, with all of its intrigue. As mentioned earlier, Clairaut and Lagrange had considered the problem of fitting astronomical data, and because those data had periodic patterns, it was natural to use approximating functions consisting of sines and cosines. Since the data represented *discrete* samples of an unknown function, and since the approximating functions were *finite* sums of trigonometric functions, this work also led to some of the earliest expressions of the discrete Fourier transform.

The work of Lagrange on interpolation was undoubtedly known to the German mathematician Carl Friedrich Gauss, whose prolific stream of mathematics originated in Göttingen. Almost a footnote to Gauss' vast output was his own contribution to trigonometric interpolation, which also contained the discrete Fourier transform. Equally significant is a small calculation buried in his treatise on interpolation [61] that appeared posthumously in 1866 as an unpublished paper. This work has been dated to 1805, and it contains the first clear and indisputable use of the **fast Fourier transform (FFT)**, which is generally attributed to Cooley and Tukey in 1965. Ironically, Gauss' calculation was cited in 1904 in the mathematical encyclopedia of Burkhardt [25] and again in 1977 by Goldstine [67]. The entire history of the FFT was recorded yet again in 1985 in a fascinating piece of mathematical sleuthing by Heideman, Johnson, and Burrus [74], who remark that "Burkhardt's and Goldstine's works went almost as unnoticed as Gauss' work itself."

To introduce the discrete Fourier transform and to recognize a historical landmark, it would be worthwhile to have a brief look at the problem that Gauss was working on when he resorted to his own fast Fourier transform. Around 1800 Gauss became interested in astronomy because of the discovery of the asteroid Ceres, the orbit of which he was able to determine with great accuracy. At about the same time, Gauss obtained the data presented in Table 1.1 for the position of the asteroid Pallas.

Table 1.1

Data on the asteroid Pallas used by Gauss, reproduced by Goldstine from tables of Baron von Zach. Reprinted here, by permission, from [H. H. Goldstine, A History of Numerical Analysis from the 16th through the 19th Century*, Springer–Verlag, Berlin,* 1977]. ©1977*, Springer–Verlag.*

Ascension, θ (degrees)	0	30	60	90	120	150
Declination, X (minutes)	408	89	−66	10	338	807
Ascension, θ (degrees)	180	210	240	270	300	330
Declination, X (minutes)	1238	1511	1583	1462	1183	804

Departing from Gauss' notation, we will let θ and X represent the ascension (still measured in degrees) and the declination, respectively. We will also let (θ_n, X_n) denote the actual data pairs where $n = 0, \ldots, 11$. Since the values of X appear to have a periodic pattern as they vary with θ, the goal is to fit these twelve data points with a trigonometric expression of the form

$$X = f(\theta) = a_0 + \sum_{k=1}^{5} \left[(a_k \cos\left(\frac{2\pi k\theta}{360}\right) + b_k \sin\left(\frac{2\pi k\theta}{360}\right)\right] + a_6 \cos\left(\frac{2\pi \cdot 6 \cdot \theta}{360}\right).$$

Notice that there are twelve unknown coefficients (the a_k's and b_k's), each of which multiplies a sine or cosine function (called **modes**) with a certain period or frequency. The fundamental period (represented by the coefficients a_1 and b_1) is 360 degrees; the modes with $k > 1$ have smaller periods of $360/k$ degrees.

Twelve conditions are needed to determine the twelve coefficients. Those conditions are simply that the function f match the data; that is,

$$f(\theta_n) = X_n$$

for $n = 0 : 11$ (throughout the book the notation $n = N_1 : N_2$ will be used to indicate that the integer index n runs through all consecutive integers between and including the integers N_1 and N_2). These conditions amount to a system of twelve linear equations that even in Gauss' day could have been solved without undue exertion. Nevertheless Gauss, either in search of a shortcut or intrigued by the symmetries of the sine and cosine functions, discovered a way to collect the coefficients and equations in three subproblems that were easily solved. The solutions to these subproblems were then combined into a final solution. Goldstine [67] has reproduced this remarkable computation, and it is evident that the splitting process discovered by Gauss lies at the heart of the modern fast Fourier transform method.

At the moment we are most interested in the outcome and interpretation of Gauss' calculation. Either by solving a 12×12 linear system or by doing a small FFT, the coefficients of the interpolation function are given in Table 1.2.

TABLE 1.2
Coefficients in Gauss' trigonometric interpolating function

k	0	1	2	3	4	5	6
a_k	780.6	−411.0	43.4	−4.3	−1.1	.3	.1
b_k	–	−720.2	-2.2	5.5	−1.0	−.3	–

A picture is far more illuminating than the numbers. Figure 1.2 shows the data (marked as $\times$'s) and the function f plotted as a smooth curve. Notice that the interpolation function does its job: it passes through the data points. Furthermore, the function and its coefficients exhibit the "frequency structure" of the data. The coefficients with the largest magnitude belong to the constant mode (corresponding to a_0), and to the fundamental mode (corresponding to a_1 and b_1). However, contributions from the higher frequency modes are needed to represent the data exactly.

His use of the FFT aside, Gauss' fitting of the data with a sum of trigonometric functions is an example of a DFT and is a prototype problem for this book. As we will see shortly, the "input" for such a problem may originate as data points (as in this example) or as a function defined on an interval. In the latter case, the process of finding the coefficients in the interpolating function is equivalent to approximating either the Fourier series coefficients or the Fourier transform of the given function. We will explore all of these uses of the DFT in the remainder of this book.

1.2. An Application

We now turn to a practical application of the DFT to a problem of **data analysis** or, more specifically, **time series analysis**. Even with many details omitted, the example offers a vivid demonstration of some of the fundamental uses of the DFT, namely spectral decomposition and filtering. On a recent research expedition to the Trapridge Glacier in the Yukon Territory, Canada, data were collected by sensors in the

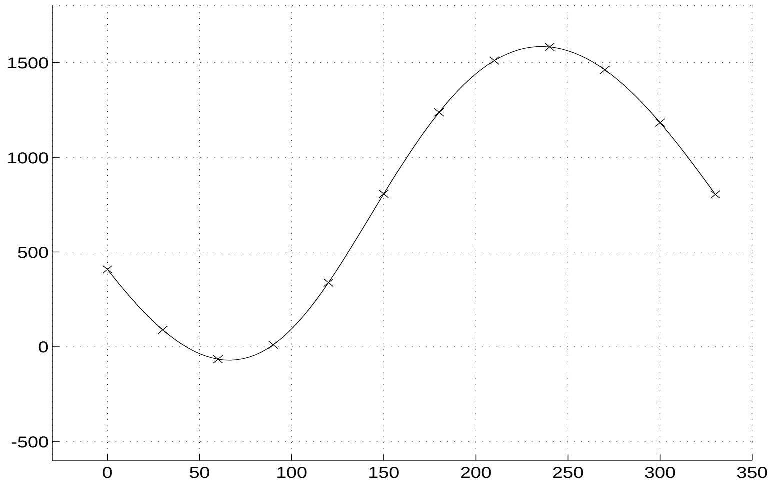

FIG. 1.2. *The* 12 *data points* (*marked by* $\times$*'s*) *describe the position of the asteroid Pallas. The declination* (*in minutes*) *of Pallas is plotted along the vertical axis against the ascension* (*in degrees*). *A weighted sum of* 12 *sine and cosine functions with different periods is used to fit the data* (*smooth curve*). *The coefficients in that weighted sum are determined by a discrete Fourier transform.*

bed of the glacier, 80 meters below the surface [132]. In particular, measurements of the turbidity (amount of suspended material) of the subglacial water were taken every $\Delta t = 10$ minutes $\approx .0069$ days. When plotted, these data produce the jagged curve shown in the upper left graph of Figure 1.3 (time increases to the right and turbidity values are plotted on the vertical axis). The curve actually consists of $N = 368$ data points, which represent $N\Delta t = 2.55$ days of data.

Notice that the data exhibit both patterns *and* irregularities. On the largest scale we see a wave-like pattern with a period of approximately one day; this is an explicable diurnal variation in the water quality variables of the glacier that has a relatively *low frequency.* On a smaller time scale the data appear to be infected with a *high frequency* oscillation which is attributable partly to instrument noise. Between these high and low frequency patterns, there may be other effects that occur on intermediate time scales.

One of the most revealing analyses that can be done on a set of data, or any "signal" for that matter, is to decompose the data according to its various frequencies. This procedure is often called **spectral analysis** or **spectral decomposition**, and it provides an alternative picture of the data in the **frequency domain**. This frequency picture tells how much of the variability of the data is comprised of low frequency waves and how much is due to high frequency patterns. Using the glacier data, here is how it works. The first question that needs to be answered concerns the range of frequencies that appear in a data set such as that shown in the top left graph of Figure 1.3. Just a few physical arguments provide the answer. We will let the total time interval spanned by the data set be denoted A; for the glacier data, we have $A = 2.55$ days. As mentioned earlier, we will let the spacing between data points be denoted Δt. These two parameters, the length of the time interval A, and the so-called **sampling rate** Δt are absolutely crucial in the following arguments. Here are the key observations that give us the range of frequencies.

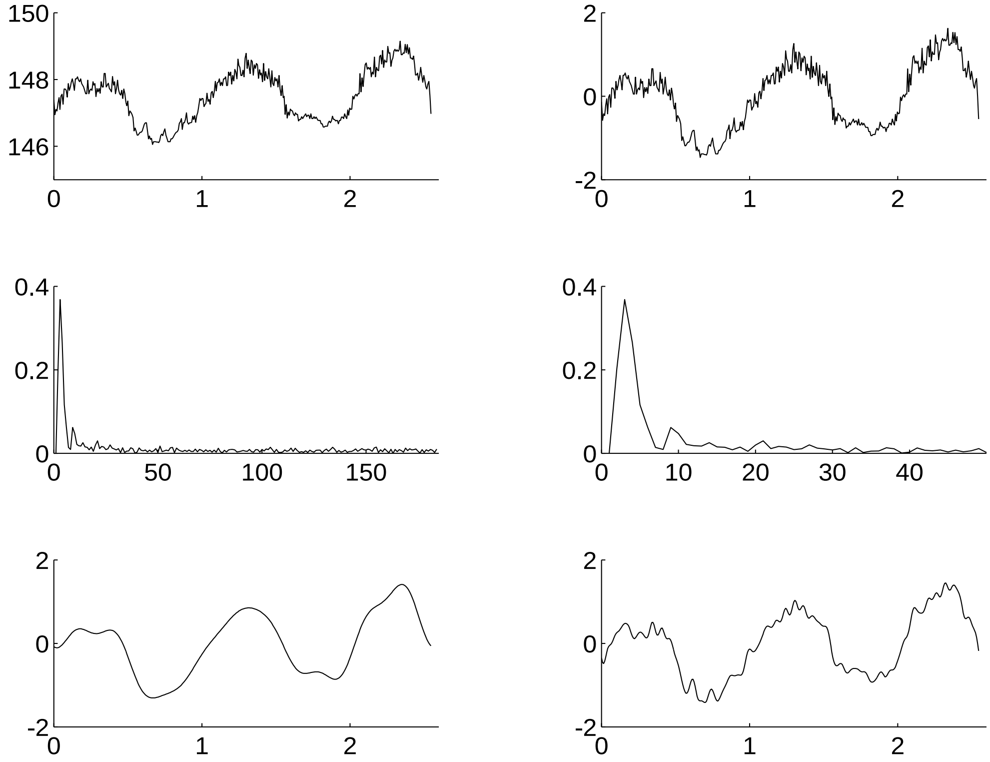

FIG. 1.3. *The above graphs show how the DFT can be used to carry out a spectral analysis of a data set. The original data set consisting of* $N = 368$ *subglacial water turbidity measurements taken every* $\Delta t = 10$ *minutes is shown in the top left figure. The same data set with the "mean subtracted out" is shown in the top right figure. The center left figure shows the spectrum of the data after applying the DFT. The horizontal axis now represents frequency with values between* $\omega_1 = .39$ *and* $\omega_{184} = 72$ *days*$^{-1}$. *The relative weight of each frequency component in the overall structure of the turbidity data is plotted on the vertical axis. The center right figure is a close-up of the spectrum showing the low frequency contributions. By filtering the high frequency components, the data can be smoothed. The lower left and right figures show reconstructions of the data using the lowest* 10 *and lowest* 50 *frequencies, respectively.*

If we collect data for 2.55 days, the longest complete oscillation that we can resolve lasts for precisely 2.55 days. Patterns with periods greater than 2.55 days cannot be fully resolved within a 2.55 day time frame. What is the frequency of this longest wave? Using the units of periods per day for frequency, we see that this wave has a frequency of

$$\omega_1 = \frac{1}{A} = \frac{1 \text{ period}}{2.55 \text{ days}} = .39 \text{ days}^{-1}.$$

The subscript 1 on ω_1 means that this is the first (and lowest) of an entire set of frequencies. Therefore, the length of the time interval A days determines the lowest frequency in the system.

We now consider periodic patterns with higher frequencies. A wave with two full periods over A days has a frequency of $\omega_2 = 2/A$ days^{-1}. Three full periods over the interval corresponds to a frequency of $\omega_3 = 3/A$ days^{-1}. Clearly, the full set of frequencies is given by

$$\omega_k = \frac{k}{2.55} = \frac{k}{A} = \text{days}^{-1}.$$

Having determined where the range of frequencies begins, an even more important consideration is where it ends. Not surprisingly, the sampling rate Δt determines the maximum frequency. One might argue that there should be no maximum frequency, since oscillations continue on ever-decreasing time scales all the way to molecular and atomic levels. However, when we take samples every $\Delta t = 10$ minutes, there *is* a limit to the time scales that can be resolved. In pages to come, we will demonstrate in many ways the following fundamental fact: if a signal (or phenomenon of any kind) is sampled once every Δt units, then a wave with a period less than $2\Delta t$ cannot be resolved accurately. The most detail that we can see with a "shutter speed" of Δt is a wave that has a peak at one sample time, a valley at the next, and a peak at the next. This "up-down-up" pattern has a period of $2\Delta t$ units, and a frequency of

$$\omega_{\max} = 1/(2\Delta t) \quad \text{units}^{-1}.$$

Where does this maximum frequency fit into the set of frequencies that we denoted $\omega_k = k/A$? Notice that $\Delta t = A/N$, where N is the total number of data points. Therefore

$$\omega_{\max} = \frac{1}{2\Delta t} = \frac{1}{2 \cdot A/N} = \frac{N}{2} \cdot \frac{1}{A} = \omega_{\frac{N}{2}}.$$

In other words, the highest frequency that can be resolved with a sample rate of $\Delta t = A/N$ is the $N/2$ frequency in the set of frequencies. (If N is odd, a modified argument can be made, as shown in Chapter 3.) In summary, the $N/2$ frequencies that can be resolved over an interval of A units with a sample rate of Δt units are given by

$$\omega_1 = \frac{1}{A}, \ldots, \omega_k = \frac{k}{A}, \ldots, \omega_{\frac{N}{2}} = \frac{1}{2\Delta t} = \frac{N}{2A}.$$

For the glacier data set, the numerical values of these frequencies that can be resolved are

$$\omega_1 = .39, \ \ \omega_2 = .78, \ldots, \omega_k = .39\,k, \ldots, \omega_{184} = 72 \quad \text{days}^{-1}.$$

We can now turn to an analysis of the turbidity data. A common first step in data or signal analysis is to subtract out the mean. This maneuver amounts to computing the average of the turbidity values and subtracting this average from all of the turbidity values. The result of this rescaling is shown in the top right plot of Figure 1.3; the data values appear as fluctuations about the mean value of zero. With this adjusted data set, we can perform the frequency or spectral decomposition of the data. The details of this calculation will be omitted, since it involves the discrete Fourier transform, which is the subject of the remainder of the book! However, we can examine the output of this calculation and appreciate the result. The center left graph of Figure 1.3 shows what is commonly called the **spectrum** or **power spectrum** of the data set. The horizontal axis of this graph is no longer *time*, but rather *frequency*. Notice that $N/2 = 184$ frequencies are represented on this axis. On the vertical axis is a measure of the relative weight of each of the frequencies in the overall structure of the turbidity data. Clearly, most of the "energy" in the data resides in the lower frequencies, that is, the ω_k's with $k < 10$. The center right graph of Figure 1.3 is a close-up of the spectrum for the lower frequencies. In this picture it is evident that the dominant frequencies are ω_2 and ω_3, which correspond to periods of 1.25 days and 0.83 days. These periods are the nearest to the prominent diurnal (daily) oscillation that we observed in the data at the outset. Thus the spectral decomposition does capture both the obvious patterns in the data, as well as other hidden frequency patterns that are not so obvious visually.

What about the experimental "noise" that also appears in the data? This noise shows up as all of the high frequency ($k >> 10$) contributions in the spectrum plot. This observation brings us to another strategy. *If* we believe that the high frequency oscillations really are spurious and contain no useful physical information, we might be tempted to **filter** these oscillations out of the data. Using the original data set, this could be a delicate task, since the high frequency part of the signal is spread throughout the data set. However, with the spectral decomposition plot in front of us, it is an easy task, since the high frequency noise is localized in the high frequency part of the plot.

The process of filtering is a science and an art that goes far beyond the meat-cleaver approach that we are about to adopt. Readers who wish to see the more refined aspects of filtering should consult one of many excellent sources [20], [70], [108]. With apologies aside, the idea of filtering can be demonstrated by simply removing all of the high frequency contributions in the spectrum above a chosen frequency. This method corresponds to using a sharp low-pass filter (low frequencies are allowed to pass) that truncates the spectrum.

With this new truncated spectrum, it is now possible to reconstruct the data set in an inverse operation that also requires the use of the DFT. Omitting details, we can observe the outcome of this reconstruction in the lower plots of Figure 1.3. The first of these two figures shows the result of removing all frequencies above ω_{10}. Notice that all of the noise in the data has been eliminated, and indeed even the low frequency oscillations have been smoothed considerably. The second of the reconstructions uses all frequencies below ω_{50}; this filter results in a slightly sharper replica of the low frequency oscillations, with more of the high frequency noise also evident.

This discussion of spectral decomposition and filtering has admittedly been rather sparse in details and technicalities. Hopefully it has illustrated a few fundamental concepts that will recur endlessly in the pages to come. In summary, given a function or data set defined in a **time** or **spatial domain**, the DFT serves as a tool that gives a frequency picture of that function or data set. The frequency picture depends on the length of the interval A on which the original data set is defined, and it also depends on the rate Δt that is used to sample the data. Once the data set is viewed in the frequency domain, there are operations such as filtering that can be applied if necessary. Equally important is the fact that if the problem calls for it, there is always a path back to the original domain; this passage also uses the DFT. All of these ideas will be elaborated and explored in the pages to come. It is now time to dispense with generalities and history, and look at the DFT in all of its power and glory.

1.3. Problems

1. Gauss' problem revisited. Consider the following simplified version of Gauss' data fitting problem. Assume that the following $N = 4$ data pairs (x_n, y_n)

are collected:

$$(0, y_0), \quad \left(\frac{\pi}{2}, y_1\right), \quad (\pi, y_2), \quad \left(\frac{3\pi}{2}, y_3\right),$$

where the values y_0, y_1, y_2, and y_3 are known, but will be left unspecified. The data set is also assumed to be periodic with period of 2π, so one could imagine another data point $(2\pi, y_4) = (2\pi, y_0)$ if necessary. This data set is to be fit with a function of the form

$$y = f(x) = a_0 + a_1 \cos x + b_1 \sin x + a_2 \cos 2x, \tag{1.1}$$

where the four coefficients a_0, a_1, b_1, a_2 are unknown, but will be determined from the data.

(a) Verify that each of the individual functions in the representation (1.1), namely, $f_0(x) = 1$, $f_1(x) = \cos x$, $f_2(x) = \sin x$, and $f_3(x) = \cos(2x)$ has a period of 2π (that is, it satisfies $f(x) = f(x + 2\pi)$ for all x).

(b) Conclude that the entire function f also has a period of 2π.

(c) Plot each of the individual functions f_0, f_1, f_2, and f_3, and note the period and frequency of each function.

(d) To find the coefficients a_0, a_1, b_1, and a_2 in the representation (1.1), impose the **interpolating conditions**

$$f(x_n) = y_n$$

for $n = 0, 1, 2, 3$, where $x_n = n\pi/2$. Write down the resulting system of linear equations for the unknown coefficients (recalling that y_0, y_1, y_2, and y_3 are assumed to be known).

(e) Use any available means (pencil and paper will suffice!) to solve for the coefficients a_0, a_1, b_1, and a_2 given the data values

$$y_0 = 2, \;\; y_1 = 1, \;\; y_2 = 3, \;\; y_3 = 2.$$

(f) Plot the resulting function f and verify that it does pass through each of the four data points.

2. The complex exponential. Given the Euler relation

$$e^{i\theta} = \cos\theta + i \sin\theta,$$

verify that

(a) $e^{im\theta} = \cos(m\theta) + i \sin(m\theta)$,

(b) $e^{i\theta} e^{i\phi} = e^{i(\theta+\phi)} = \cos(\theta + \phi) + i \sin(\theta + \phi)$,

(c) $\cos\theta = \dfrac{e^{i\theta} + e^{-i\theta}}{2}$ and $\sin\theta = \dfrac{e^{i\theta} - e^{-i\theta}}{2i}$ (also called Euler relations),

(d) $e^{i2\pi m/N} = 1$ if m is any integer multiple of the integer N,

(e) $e^{i2\pi nk/N} = e^{i2\pi(n+mN)k/N}$ where m and N are any integers.

3. Modes, freqencies and periods. Consider the functions (or modes) $v_k(x) = \cos(\pi k x/2)$ and $w_k(x) = \sin(\pi k x/2)$ on the interval $[-2, 2]$.

(a) Plot on separate sets of axes (hand sketches will do) the functions v_k and w_k for $k = 0, 1, 2, 3, 4$. In each case, note the period of the mode and the frequency of the mode. How are the period and frequency of a given mode related to the index k?

(b) Now consider the grid with grid spacing (or sampling rate) $\Delta x = 1/2$ consisting of the points $x_n = n/2$ where $n = -4 : 4$. Mark this grid on the plots of part (a) and indicate the sampled values of the mode.

(c) What are the values of v_0 and w_0 at the grid points?

(d) What are the values of v_4 and w_4 at the grid points?

(e) Plot v_5 on the interval $[-2, 2]$. What is the effect of sampling v_5 at the grid points? Compare the values of v_5 and the values of v_3 at the sample points.

4. Sampling and frequencies. Throughout this book, we have, somewhat arbitrarily, used Δx to denote a grid spacing or sampling rate. This notation suggests that the sampling takes place on a *spatial* domain. However, sampling can also take place on *temporal* domains, as illustrated by the glacier data example in the text. In this case Δt is a more appropriate notation for the grid spacing.

(a) Imagine that you have collected $N = 140$ equally spaced data values (for example, temperatures or tidal heights) over a time interval of duration $A = 14$ days. What is the grid spacing (or sampling rate) Δt? What is the minimum frequency ω_1 that is resolved by this sampling? What is the maximum frequency ω_{70} that is resolved by this sampling? What are the units of these frequencies?

(b) Imagine that you have collected $N = 120$ equally spaced data values (for example, material densities or beam displacements) over a spatial interval of length $A = 30$ centimeters. What is the grid spacing Δx? What is the minimum frequency ω_1 that is resolved by this sampling? What is the maximum frequency ω_{60} that is resolved by this sampling? What are the units of these frequencies? Note that in spatial problems frequencies are often referred to as *wavenumbers*.

5. Complex forms. It will be extremely useful to generalize problem 1 and obtain a sneak preview of the DFT. There are some advantages to formulating the above interpolating conditions in terms of the complex exponential. Show that with $N = 4$ the real form of the interpolating conditions

$$y_n = a_0 + a_1 \cos\left(\frac{2\pi n}{N}\right) + b_1 \sin\left(\frac{2\pi n}{N}\right) + a_2 \cos\left(\frac{4\pi n}{N}\right)$$

for $n = 0, 1, 2, 3$ may also be written in the form

$$y_n = \sum_{k=0}^{3} c_k e^{i2\pi nk/N}$$

for $n = 0, 1, 2, 3$, where the coefficients c_0, c_1, c_2, and c_3 may be complex-valued. Assuming that all of the data values are real, show that c_0 and c_2 are real-valued, and that $\mathrm{Re}\{c_1\} = \mathrm{Re}\{c_3\}$ and $\mathrm{Im}\{c_1\} = -\mathrm{Im}\{c_3\}$. Then the following relations between the coefficients $\{a_0, a_1, b_1, a_2\}$ and the coefficients $\{c_0, c_1, c_2, c_3\}$ result:

$$a_0 = c_0, \quad a_1 = 2\mathrm{Re}\{c_1\}, \quad b_1 = -2\mathrm{Im}\{c_1\}, \quad a_2 = c_2.$$

6. More general complex forms. To sharpen your skills in working with complex exponentials, carry out the calculation of the previous problem for a data set with N points in the following two cases.

(a) Assume that the real data pairs (x_n, y_n) are given for $n = 0, \ldots, N-1$ where $x_n = 2\pi n/N$. Show that the interpolating conditions take the real form

$$y_n = a_0 + \sum_{k=1}^{\frac{N}{2}-1} \left[a_k \cos\left(\frac{2\pi nk}{N}\right) + b_k \sin\left(\frac{2\pi nk}{N}\right)\right] + a_{\frac{N}{2}} \cos(\pi n),$$

for $n = 0 : N-1$. Show that these conditions are equivalent to the conditions in complex form

$$y_n = \sum_{k=0}^{N-1} c_k e^{i2\pi nk/N}$$

for $n = 0 : N-1$, where c_0 and $c_{\frac{N}{2}}$ are real, $\mathrm{Re}\{c_k\} = \mathrm{Re}\{c_{N-k}\}$, and $\mathrm{Im}\{c_k\} = -\mathrm{Im}\{c_{N-k}\}$. The relations between the real coefficients $\{a_k, b_k\}$ and the complex coefficients c_k are $a_0 = c_0, a_{\frac{N}{2}} = c_{\frac{N}{2}}$, and

$$a_k = 2\mathrm{Re}\{c_k\}, \quad b_k = -2\mathrm{Im}\{c_k\}$$

for $k = 1 : N/2 - 1$.

(b) Now assume the N real data pairs (x_n, y_n) are given for $n = -N/2 + 1, \ldots, N/2$, where $x_n = 2\pi n/N$. Notice that the sampling interval is now $[-\pi, \pi]$. The data fitting conditions take the real form

$$y_n = a_0 + \sum_{k=1}^{\frac{N}{2}-1} \left[a_k \cos\left(\frac{2\pi nk}{N}\right) + b_k \sin\left(\frac{2\pi nk}{N}\right)\right] + a_{\frac{N}{2}} \cos(\pi n),$$

where $n = -N/2 + 1 : N/2$. Show that these conditions are equivalent to the conditions in complex form

$$y_n = \sum_{k=-\frac{N}{2}+1}^{\frac{N}{2}} c_k e^{i2\pi nk/N}$$

for $n = -N/2+1 : N/2$, where c_0 and $c_{\frac{N}{2}}$ are real, $\mathrm{Re}\{c_k\} = \mathrm{Re}\{c_{-k}\}$, and $\mathrm{Im}\{c_k\} = -\mathrm{Im}\{c_{-k}\}$. The relations between the real coefficients $\{a_k, b_k\}$ and the complex coefficients c_k are $a_0 = c_0, a_{\frac{N}{2}} = c_{\frac{N}{2}}$, and

$$a_k = 2\mathrm{Re}\{c_k\}, \quad b_k = -2\mathrm{Im}\{c_k\}$$

for $k = 1 : N/2 - 1$.

We will see shortly that all of the above procedures for determining the coefficients c_k from the data values y_n (or vice-versa) are versions of the DFT.

7. Explain those figures. The cover and frontmatter of this book display several polygonal, mandala-shaped figures which were generated using the DFT. With just a small preview of things to come, we can explain these figures. As will be shown early in the next chapter, given a sequence $\{f_n\}_{n=0}^{N-1}$, one form of the DFT of this sequence is the sequence $\{F_k\}_{k=0}^{N-1}$, where

$$F_k = \frac{1}{N} \sum_{n=0}^{N-1} f_n e^{-i2\pi nk/N}.$$

(a) Let $N = 4$ for the moment and compute the DFT of the unit vectors $(1, 0, 0, 0)$, $(0, 1, 0, 0)$, $(0, 0, 1, 0)$, and $(0, 0, 0, 1)$. Verify that the corresponding DFTs are $\frac{1}{4}(1, 1, 1, 1)$, $\frac{1}{4}(1, -i, 1, i)$, $\frac{1}{4}(1, -1, 1, -1)$, and $\frac{1}{4}(1, i, -1, -i)$.

(b) Plot each of these DFT vectors by regarding each component of a given vector as a point in the complex plane and draw a connecting segment from each component to its predecessor, cyclically (i.e., $F_0 \to F_1 \to F_2 \to F_3 \to F_1$).

(c) Repeat this process with the eight unit vectors $\hat{e}_j = (0, 0, \ldots, 1, 0, 0, \ldots, 0)$, where the 1 is in the jth component. Show that the 8-point DFT of $\hat{e}_j$ is

$$F_k = \frac{1}{8}\left(\cos\left(\frac{\pi jk}{4}\right) - i \sin\left(\frac{\pi jk}{4}\right)\right),$$

which holds for $k = 0 : 7$.

(d) Verify that plotting each of these eight DFTs in the complex plane (connecting end-to-end as described above) produces a figure similar to the ones shown on the cover and elsewhere.