Fundamentals of Signal Processing

A Tutorial on Waves, the Frequency Domain, and Their Use in Contactless Sensing

CS60055: Ubiquitous Computing
Department of Computer Science and Engineering
Indian Institute of Technology Kharagpur

Note on preparation. This tutorial was compiled from the CS60055 lecture material on contactless sensing and radio frequency sensing principles. The text was drafted and polished with the assistance of Claude Opus 4.8, a large language model developed by Anthropic. The mathematical statements follow the lecture material; readers are nevertheless encouraged to verify derivations and numerical values against a standard signal processing reference before relying on them.

Contents

1 Introduction

Ubiquitous computing depends on the ability to observe human activity and the state of the physical environment continuously and unobtrusively. Wearable sensing, built on inertial measurement units embedded in smartwatches, earables, smart glasses, and smart rings, offers one route to this goal, but it requires the subject to carry or wear an instrumented device at all times. For many settings, unobtrusive sleep monitoring being a familiar example, this requirement is inconvenient or impractical. Contactless sensing removes it. By observing how a wave that already pervades the environment is altered as it travels, reflects, and scatters, we can infer the presence, motion, and even the material composition of objects without attaching anything to them. The same idea applies when the target is not a person but a material, a liquid, or a structure, which makes contactless techniques attractive for continuous and passive monitoring across a wide range of applications.

Every contactless sensing system, whether it uses radio, sound, light, or heat, rests on a shared foundation: the physics of waves and the mathematics used to analyse them. This tutorial develops that foundation. It begins with the wave modalities available for sensing and the way waves interact with matter, proceeds to the representation of a signal and the distinction between its time and frequency descriptions, and closes with the discrete transforms and filters that make signal analysis practical on sampled data. The treatment is deliberately general. Specific sensing pipelines built on these principles are taken up in later material.

2 Sensing Modalities and the Wave Spectrum

2.1 Electromagnetic and Mechanical Waves

Sensing systems are built almost exclusively on two categories of wave. The first is the electromagnetic wave, on which the majority of sensing devices operate, spanning radio, microwave, infrared, visible, and higher frequency radiation. The second is the mechanical wave, whose propagation requires a medium and which underlies acoustic and ultrasonic sensing. The operating principle of a given system depends strongly on the type of wave it uses and on the frequency at which it operates, because these determine how the wave couples to the objects of interest.

The existence of electromagnetic waves was predicted by James Clerk Maxwell in 1864 and confirmed experimentally by Heinrich Hertz in 1887. Hertz further demonstrated that such waves are affected and reflected by solid objects, which is precisely the property that makes them useful for sensing. A wave is characterised by its frequency and, equivalently through the propagation speed, by its wavelength.

2.2 The Electromagnetic Spectrum

Across the electromagnetic spectrum, wavelength and frequency vary over many orders of magnitude, from gamma rays and X-rays at the shortest wavelengths and highest energies, through ultraviolet, the narrow visible band, and infrared, out to microwave and radio at the longest wavelengths and lowest energies (Figure 1). It is useful to remember that short wavelength corresponds to high energy and high frequency, while long wavelength corresponds to low energy and low frequency. The objects that everyday sensing systems care about, roughly from the scale of a pinpoint or an insect up to the scale of a person, fall in the wavelength range spanned by the infrared, microwave, and radio portions of the spectrum, which is one reason these bands dominate contactless sensing.

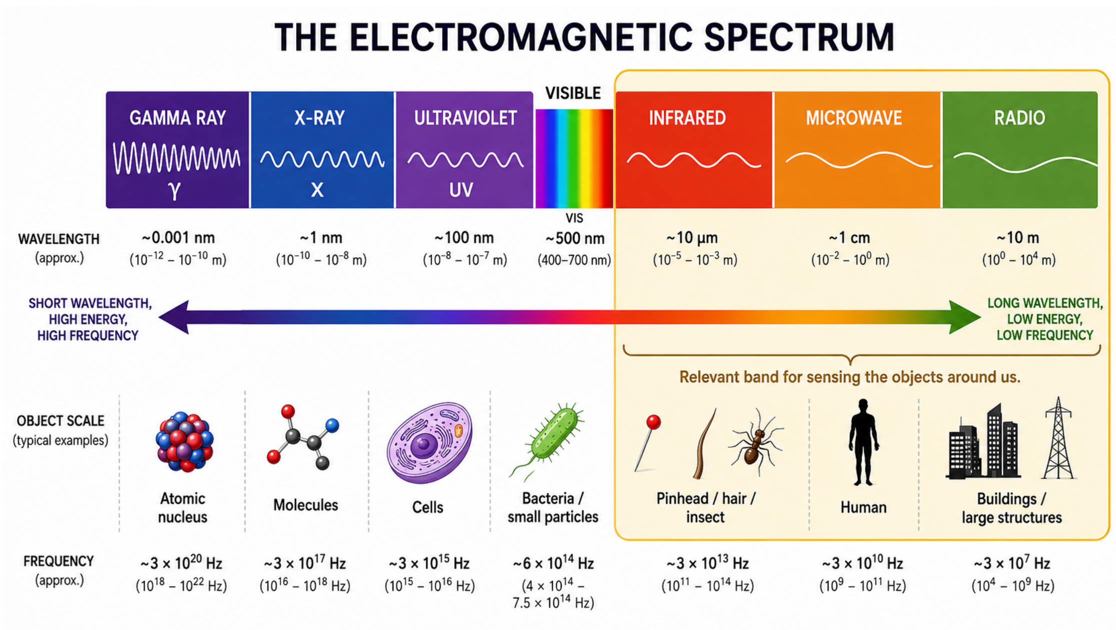

Figure 1: The electromagnetic spectrum arranged by radiation type (gamma ray, X-ray, ultraviolet, visible, infrared, microwave, radio) with representative wavelengths decreasing from radio toward gamma, and an indication of the object scale associated with each region. An arrow marks the trend from short wavelength, high energy, high frequency on one side to long wavelength, low energy, low frequency on the other. The infrared to radio region is highlighted as the band relevant to sensing the objects around us.

2.3 Penetration and Reflection

The way a wave interacts with an object depends on the relationship between its wavelength and the object. As a general rule, long wavelength, low frequency radiation tends to penetrate through physical objects; radiation of intermediate wavelength and frequency tends to reflect from them; and very short wavelength, high frequency radiation again penetrates, now passing through the object at a finer scale. This picture is only a first approximation. The actual balance between penetration and reflection also depends on the bandwidth of the signal and on other channel parameters, a point we return to in Section 5.2.

Two regions of this tradeoff are especially relevant to passive contactless sensing. The band spanning microwave and radio frequencies is attractive because reflected energy returns to a receiver and carries information about the reflecting object, so for these longer wavelength signals we analyse the reflection properties. The infrared and visible region, the vision domain, is also used widely for sensing, but it raises privacy concerns because it can reveal identifiable detail. At the shortest wavelengths, sensing exploits penetration rather than reflection, and here we analyse the penetration properties of the signal. The remainder of this tutorial concentrates on the reflective regime, since it is the one on which radio frequency sensing is built.

2.4 Radio Bands for Sensing

Within the radio portion of the spectrum, several sub-bands host the sensing techniques studied in this course (Figure 2). Radio frequency identification, together with toll tags, occupies the lower radio frequencies. Wireless local area networking, on which WiFi sensing is based, sits in the microwave region around and above a few gigahertz. Ultra wideband spans a wide swath of the microwave region and is used where fine range resolution is required. Millimeter wave, around and above 30 GHz, supports high resolution radar for applications such as advanced driver assistance. The techniques studied later in the course, in the RFID, WiFi, ultra wideband, and millimeter wave bands, all live in this part of the spectrum.

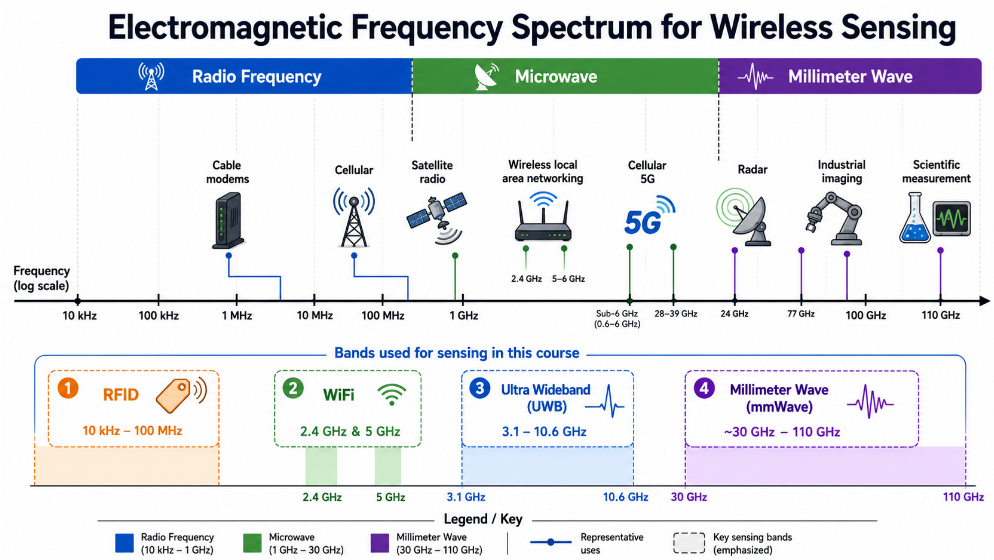

Figure 2: The electromagnetic frequency spectrum from roughly 10 kHz to 110 GHz, partitioned into radio frequency, microwave, and millimeter wave regions, with representative uses such as cable modems, cellular, satellite radio, wireless local area networking, industrial imaging, cellular 5G, radar, and scientific measurement. The four bands used for sensing in this course are highlighted: RFID at the low end, WiFi and ultra wideband in the microwave region, and millimeter wave near and above 30 GHz.

3 How Waves Interact with the Environment

3.1 From Radar to Reflection-based Sensing

The idea that a propagating wave can reveal a distant moving object is old. In 1897, Alexander Popov, a physicist of the Russian Imperial Navy, was testing early wireless communication between two ships in the Baltic Sea when he observed an interference pattern produced by a third ship passing between them, and he proposed that the effect might be used to detect moving objects. Radar as an engineered instrument was proposed by the military engineer Piotr Oshchepkov in 1932, and the first industrial radar, RUS-1, entered use in 1939 with a 4 m wavelength and a transmitter and receiver separated by 35 km. These developments established the two observations that underpin radio sensing: that objects reflect radio waves, and that motion imprints a detectable signature on the reflected signal. The rest of this section examines that signature.

3.2 The Doppler Effect

The signature of motion is the Doppler effect: a change in the observed frequency of a wave when the observer and the source move relative to one another. When the source approaches the observer, successive wave cycles are emitted from positions progressively closer to the observer, so the time between arriving cycles shortens and the observed frequency rises. When the source recedes, the time between cycles lengthens and the observed frequency falls.

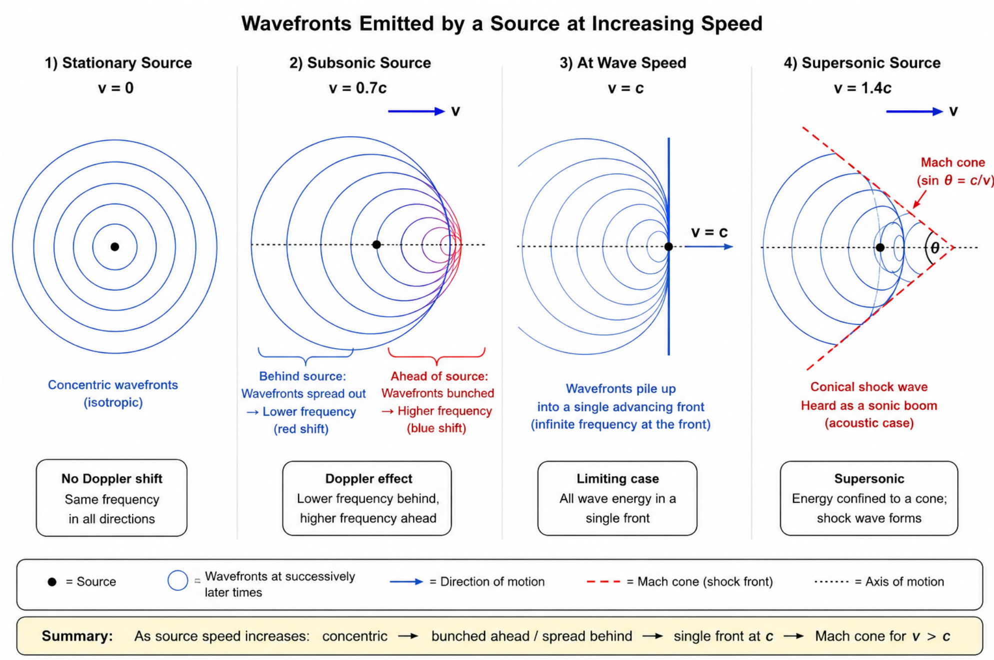

Figure 3: Wavefronts emitted by a source at increasing speed. For a stationary source the wavefronts are concentric; at $0.7c$ they bunch ahead of the source and spread behind it, producing an upward frequency shift in front and a downward shift behind; at the wave speed c they pile up into a single advancing front; and at $1.4c$ they form a conical shock wave, heard as a sonic boom in the acoustic case. Here c denotes the speed of the wave in the medium.

Let f_0 denote the emitted frequency, f the observed frequency, c the propagation speed of the wave in the medium, v_r the speed of the receiver, and v_s the speed of the source. The observed frequency is

$$f = \left(\frac{c \pm v_r}{c \mp v_s} \right) f_0. \tag{1}$$

Here v_r is added to c when the receiver is moving toward the source and subtracted when it is moving away, with the opposite convention applying to v_s in the denominator. Equivalently, when the source is directly approaching or receding from the observer,

$$\frac{f}{v_{wr}} = \frac{f_0}{v_{ws}} = \frac{1}{\lambda}, \tag{2}$$

where v_{wr} and v_{ws} are the wave speeds relative to the receiver and to the source respectively, and λ is the wavelength.

3.3 Wavefronts of a Moving Source

The geometry of the emitted wavefronts makes the effect vivid. For a stationary source the wavefronts are concentric circles, and the frequency is the same in every direction (Figure 3). As the source speed increases to a fraction of the wave speed, illustrated at $0.7c$ where c is here the speed of the wave in the medium rather than the speed of light, the wavefronts crowd together ahead of the source and spread apart behind it, raising the frequency for an observer in front and lowering it for one behind. When the source speed reaches the wave speed c, the wavefronts coincide at the leading edge and reinforce one another into a single advancing front. Beyond the wave speed, illustrated at

1.4c, the accumulated wavefronts form a conical shock wave, which in the acoustic case is heard as a sonic boom.

For contactless sensing the useful regime is the ordinary Doppler shift at speeds far below the wave speed, where the frequency change is small and is proportional to the relative velocity. Measuring this shift lets a receiver infer whether a target is moving, in which direction, and how fast, which is the basis of Doppler based motion sensing.

4 Representing a Signal

4.1 The Sinusoid

The elementary building block of signal analysis is the sinusoid. A single frequency signal may be written

$$y(t) = A\sin(\omega t + \varphi), \tag{3}$$

where A is the amplitude, ω is the angular frequency, and φ is the phase. Geometrically, the sinusoid is the projection onto one axis of a point moving uniformly around a circle (Figure 4). This picture is worth keeping in mind, because it is the same rotating point that connects the time waveform to its representation in the complex plane, developed in Section 5.3.

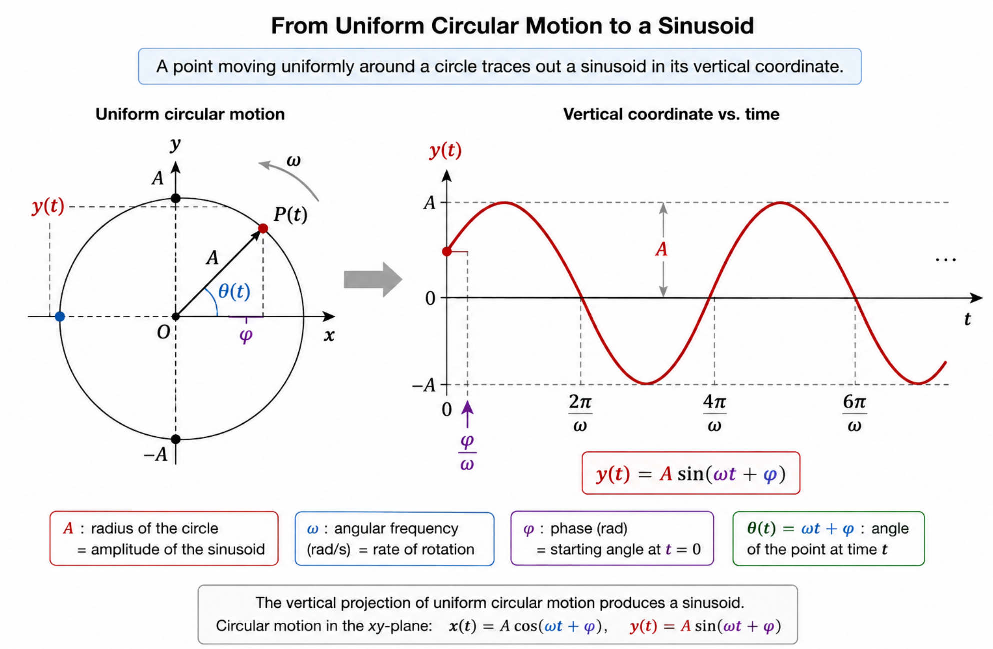

Figure 4: A point moving uniformly around a circle, with its vertical coordinate traced out over time as a sinusoid. The radius of the circle is the amplitude A, the rate of rotation is the angular frequency ω, and the starting angle is the phase φ. This construction links the circular motion on the left to the waveform on the right.

4.2 Angular and Ordinary Frequency

The angular frequency ω, measured in radians per second, relates to the ordinary frequency f, measured in cycles per second or hertz, through $\omega = 2\pi f$. Writing the sinusoid in terms of ordinary frequency,

$$y(t) = A\sin(2\pi ft + \varphi), \tag{4}$$

makes explicit the quantity that a Fourier analysis will recover. The ordinary frequency f identifies which sinusoid is present, and exposing this frequency is precisely what a frequency domain description is for.

4.3 Phase and Phase Shift

The phase φ is the fraction of a cycle that has elapsed up to a chosen time instant. Phase is an important sensing quantity in its own right, because the phase of a reflected or scattered signal depends on the material at which the transmitted signal was incident, and reflection introduces a shift in phase relative to the transmitted wave. The phase difference, or phase shift, between two otherwise identical sinusoids is the horizontal displacement between them expressed as an angle (Figure 5); a shift of ninety degrees, for instance, corresponds to a quarter of a cycle. Because different materials impose different phase shifts, a measured phase carries information about the composition of the reflecting surface, and because phase advances continuously with the distance the wave travels, it also encodes range.

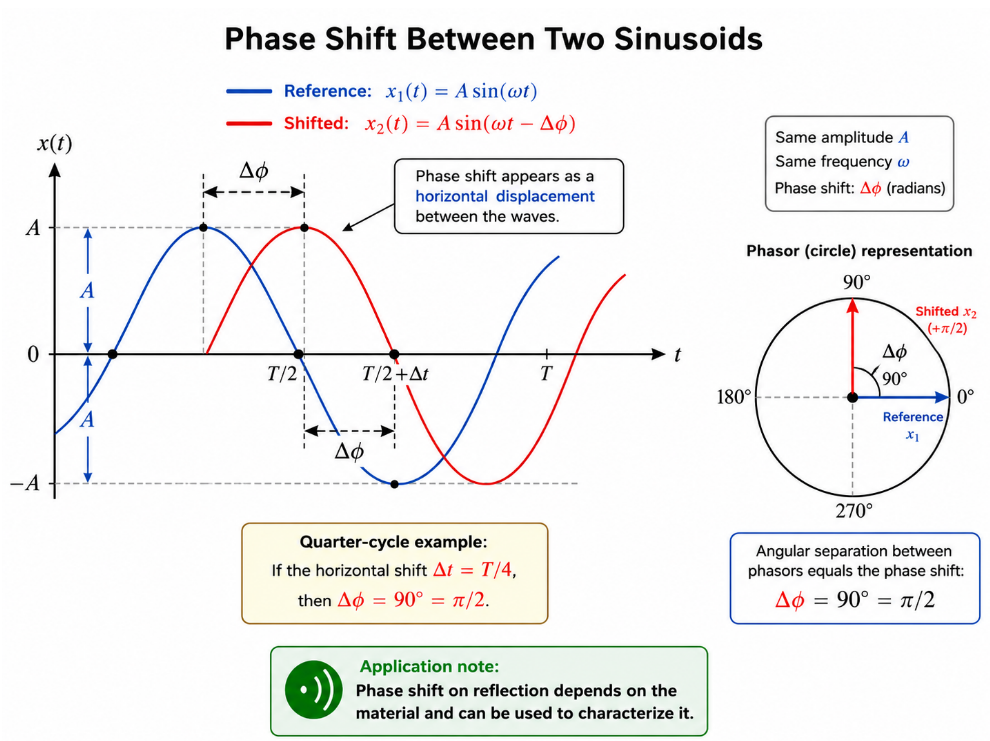

Figure 5: Two sinusoids of the same amplitude and frequency separated by a phase shift, shown as a horizontal displacement between the two curves and equivalently as an angular offset. A quarter cycle offset corresponds to a phase shift of ninety degrees. The phase shift introduced on reflection depends on the material and can therefore be used to characterise it.

5 From the Time Domain to the Frequency Domain

5.1 Frequency Components

Any arbitrary wave signal can be decomposed into a set of sinusoids of different frequencies, called the frequency components of the signal, as shown in Figure 6. This is a far reaching statement. A complicated waveform that appears structureless in time is, in the frequency domain, a specific combination of simple oscillations. Identifying those oscillations is the central task of Fourier analysis, and the value of a component at each frequency is what the analysis returns.

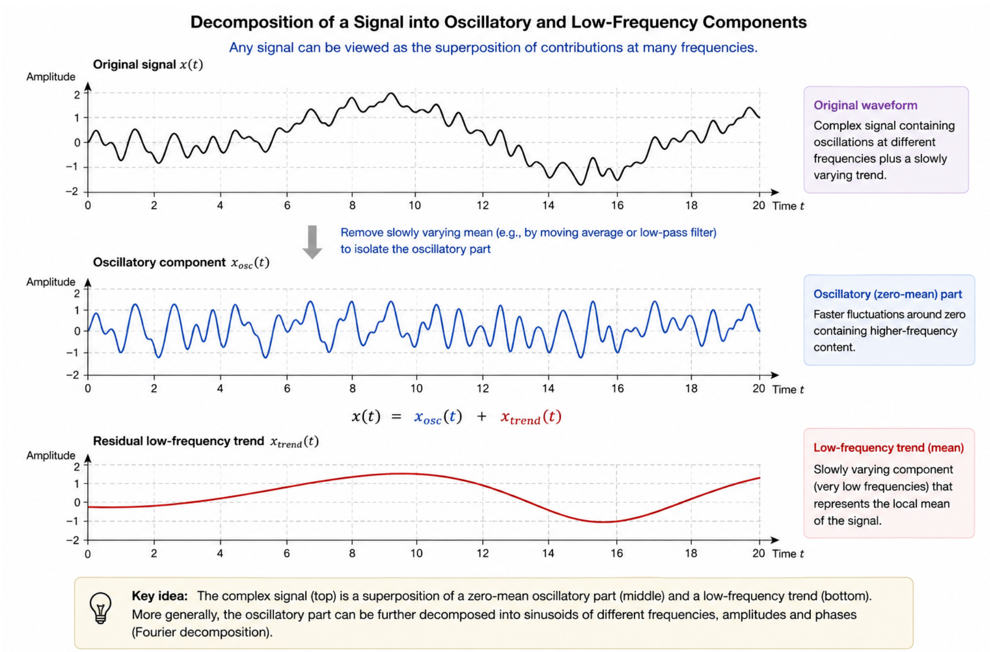

Figure 6: Decomposition of an arbitrary signal into simpler oscillatory components. The original waveform (top) is separated into an oscillatory part obtained after removing a slowly varying mean, together with a residual low frequency trend (bottom), illustrating that a complex signal is a superposition of contributions at different frequencies.

5.2 Bandwidth and Its Impact on Sensing

The bandwidth of a signal is the difference between its highest and its lowest frequency components, under the assumption that any component between these two extremes may be present. Bandwidth has a direct bearing on sensing. A higher bandwidth signal is likely to contain a larger number of frequency components, so it interacts with the environment in a richer way. Scattering, diffraction, and reflection all tend to be more pronounced for high bandwidth signals, because the many components respond differently to features of different sizes (Figure 7).

This richness is a nuisance for communication, where the additional interactions manifest as distortion and noise, but it is exactly what sensing wants, since each additional interaction is another observation of the environment. It also explains why the wavelength based rule of thumb for penetration and reflection introduced earlier is only approximate: the balance between penetration and reflection depends not on the centre frequency alone but also on the bandwidth of the signal.

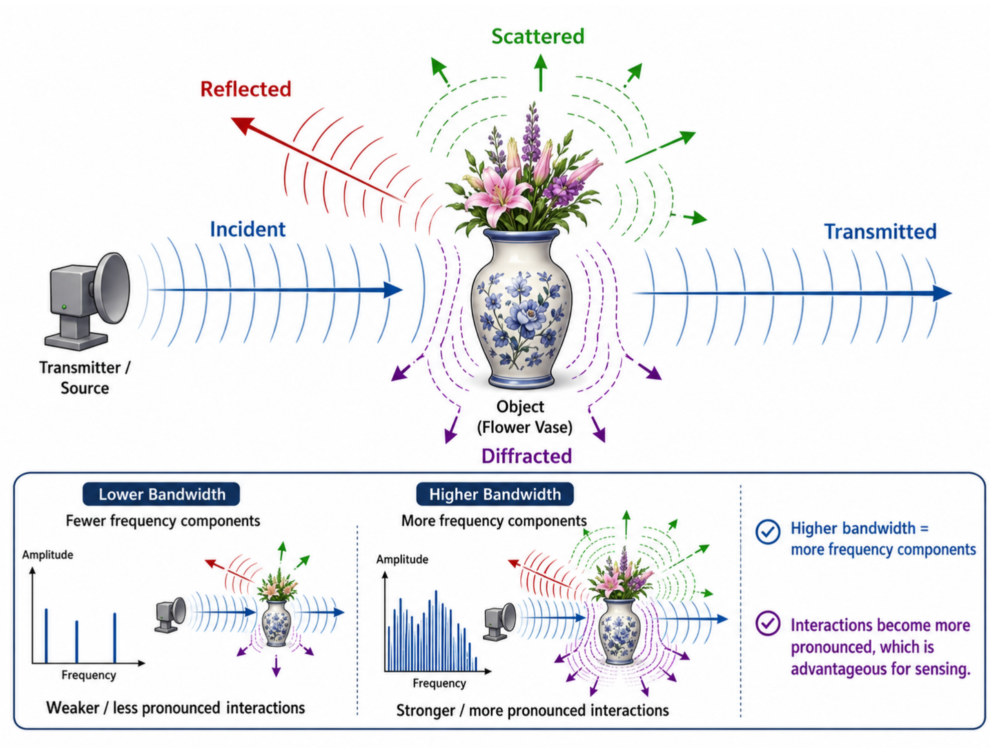

Figure 7: A signal incident on an object, showing the several ways the energy is redistributed: part is reflected back toward the source, part is scattered in other directions, part is diffracted around the object, and part passes through it. A higher bandwidth signal, having more frequency components, exhibits these effects more strongly, which is advantageous for sensing.

5.3 The Fourier Transform

To obtain the frequency components of a signal we use Fourier analysis. A pure sinusoid appears in the frequency domain as a single spike at its frequency. A sum of several sinusoids of different frequencies and amplitudes appears as several spikes at the corresponding frequencies, each with the corresponding height, even though their sum in the time domain is a complicated waveform. The frequency spectrum therefore separates a superposition back into its parts.

Even a discontinuous waveform has a frequency description. A square wave, for example (Figure 8), is the sum of a fundamental sinusoid and its odd harmonics at frequencies f, $3f$, $5f$, $7f$, and so on, with decreasing amplitude. Adding successive harmonics reconstructs the square wave with increasing fidelity, the approximation becoming sharper at the transitions as more terms are included.

The connection between the time waveform and its frequency content is made precise by the continuous Fourier transform. It rests on Euler's formula, $e^{i\varphi} = \cos\varphi + i\sin\varphi$, which represents a phase as a point on the unit circle in the complex plane, with real part $\cos\varphi$ and imaginary part $\sin\varphi$. Using this representation, the transform of a time signal $s(t)$ is

$$S(f) = \int_{-\infty}^{\infty} s(t)\, e^{-i2\pi ft}\, dt. \tag{5}$$

For each candidate frequency f, the integral correlates the signal against a complex sinusoid of that frequency. The result $S(f)$ is a complex number whose magnitude gives the strength of that frequency component and whose argument gives its phase. Carrying out the analysis in the complex

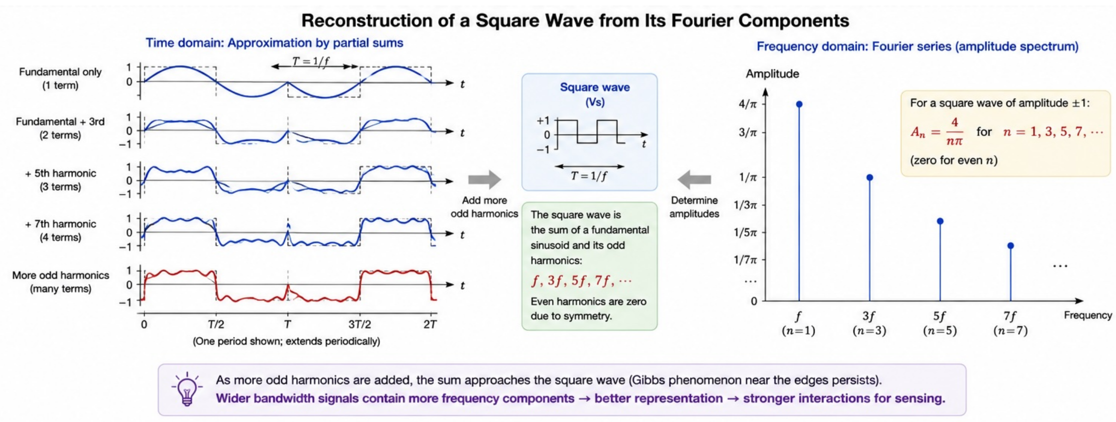

Figure 8: Reconstruction of a square wave from its Fourier components. The square wave in time (top) corresponds to a fundamental sinusoid together with odd harmonics at f, $3f$, $5f$, $7f$ (spectrum on the right), whose amplitudes decrease with frequency. Superposing these components (curve on the left) approximates the square wave, with the fit improving as more harmonics are added.

domain in this way keeps track of both amplitude and phase in a single quantity, which is essential because, as we saw in the discussion of phase, the phase is itself an informative sensing observable.

6 Processing Sampled Signals

6.1 Sampling and the Discrete Transforms

In practice we do not have access to a continuous signal but to its samples, recorded at discrete instants, which gives a discrete representation of the signal. The analysis must therefore be carried out on these samples. Applying the Fourier idea to a sampled signal gives the discrete time Fourier transform, or DTFT. Because, in addition, the number of samples we record and the number of frequency values we compute are both finite, the practical tool is the discrete Fourier transform, or DFT, which operates on a finite length time series and returns a finite set of frequency values.

The relationship between these transforms is best seen through periodicity (Figure 9). A continuous signal and its Fourier transform form one pair. Sampling the signal in time makes its spectrum periodic, which gives the DTFT, whose spectrum repeats. Requiring the computation to be finite corresponds to sampling the spectrum as well, so the DFT computes discrete samples of the continuous DTFT. The DFT is thus the fully discretised, finite counterpart of the Fourier transform, and it is what a real signal processing system actually evaluates.

Formally, let $x_0, \ldots, x_{n-1}$ be complex numbers representing the samples. The DFT is defined by

$$X_k = \sum_{m=0}^{n-1} x_m \, e^{-i2\pi km/n}, \qquad k = 0, \ldots, n-1, \tag{6}$$

where $e^{-i2\pi/n}$ is a primitive nth root of unity. Each output X_k is the correlation of the sampled signal with a discrete complex sinusoid whose frequency is indexed by k. Evaluated directly from this definition, the DFT requires on the order of $O(n^2)$ complex operations, since each of the n outputs is a sum over n inputs.

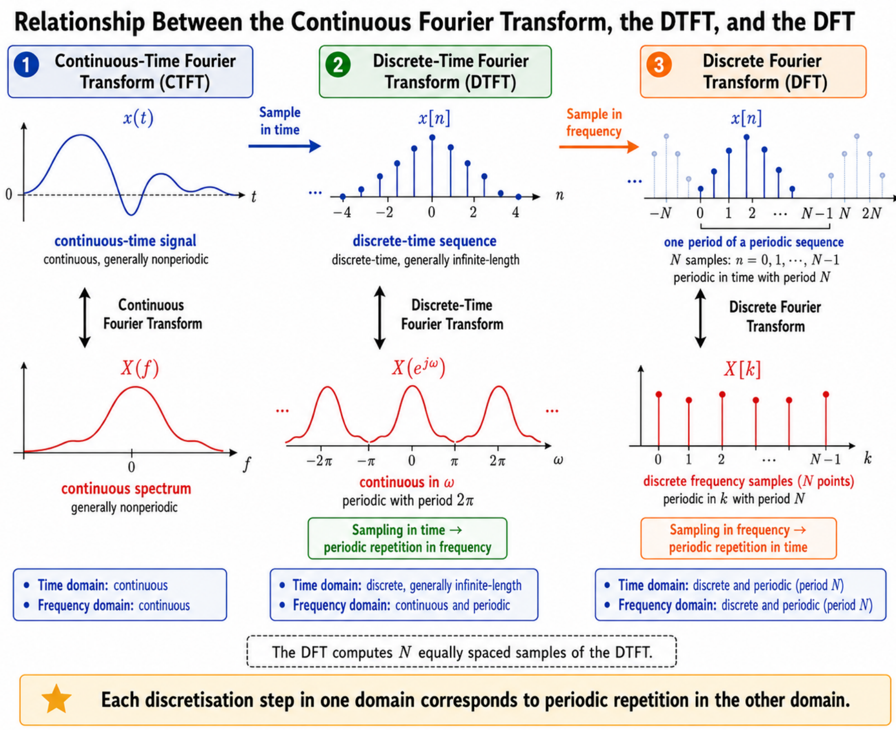

Figure 9: The relationship between the continuous Fourier transform, the DTFT, and the DFT. A continuous signal and its transform form the first pair. Sampling in time makes the spectrum periodic, giving the DTFT. Sampling the spectrum in turn produces the DFT, which computes discrete samples of the continuous DTFT. Each step of discretisation in one domain corresponds to periodic repetition in the other.

6.2 The DFT as a Matrix

Because the DFT is linear, an n point DFT can be written as a matrix multiplication $X = Wx$, where x is the vector of input samples and W is the $n \times n$ DFT matrix

$$W = \left(\frac{\omega^{jk}}{\sqrt{n}}\right)_{j,k=0,\ldots,n-1}, \qquad \omega = e^{-2\pi i/n}, \tag{7}$$

with ω a primitive nth root of unity and the factor $1/\sqrt{n}$ making the transform unitary. Each row j of W is a discrete complex sinusoid of a particular frequency, so multiplying by W simultaneously correlates the input against sinusoids of all the represented frequencies.

6.3 An Eight Point Example

The structure is clearest in a small case. For $n = 8$ the primitive root is

$$\omega = e^{-2\pi i/8} = \frac{1}{\sqrt{2}} - \frac{i}{\sqrt{2}}, \tag{8}$$

and substituting its powers turns the symbolic matrix of entries ω^{jk} into an explicit array whose elements are drawn from $\{1, -1, i, -i\}$ together with $(1 \pm i)/\sqrt{2}$ and $(-1 \pm i)/\sqrt{2}$.

Reading the rows of this matrix as waveforms is illuminating (Figure 10). The first row is constant and extracts the DC component, the average of the signal. Each subsequent row corresponds

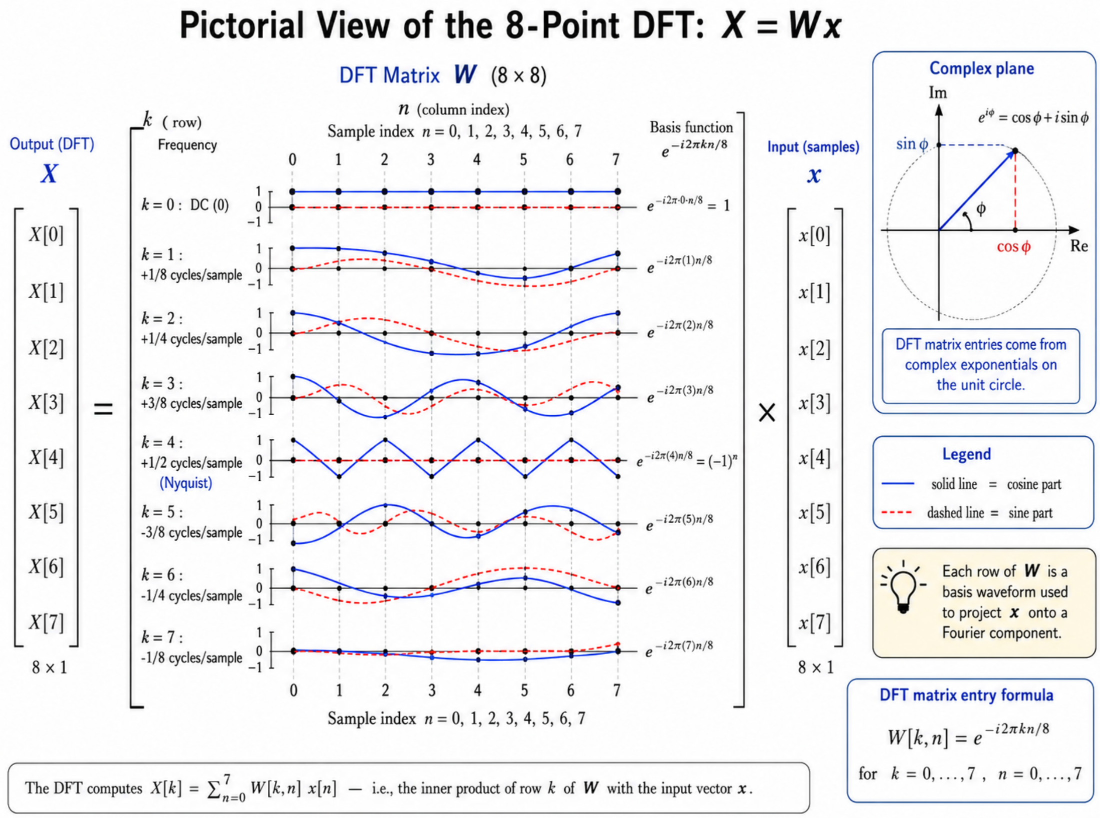

Figure 10: Pictorial view of the eight point DFT as $X = Wx$. Each row of the 8×8 matrix is a basis waveform, drawn with the cosine part as a solid line and the sine part as a dashed line. The top row is constant and gives the DC component; successive rows correspond to fractional frequencies $+1/8$, $+1/4$, $+3/8$, up to $+1/2$ at the middle row; and the lower rows correspond to the equivalent negative frequencies $-3/8$, $-1/4$, $-1/8$. The complex plane inset shows the phasor $e^{i\varphi} = \cos\varphi + i\sin\varphi$ from which the entries are obtained.

to a sinusoid of higher fractional frequency, namely $+1/8$, $+1/4$, $+3/8$ as one moves down, with the real part of the entries behaving as a cosine and the imaginary part as a sine. The middle row corresponds to the highest representable frequency of $+1/2$. The rows in the lower half of the matrix correspond to negative frequencies: a fractional frequency of $+5/8$ is equivalent to $-3/8$, $+3/4$ to $-1/4$, and $+7/8$ to $-1/8$. This equivalence follows from the identities $\cos(2\pi - \varphi) = \cos\varphi$ and $\sin(2\pi - \varphi) = -\sin\varphi$. The DFT therefore resolves the signal into a symmetric set of positive and negative frequency components.

6.4 The Fast Fourier Transform

The direct $O(n^2)$ cost of the DFT is prohibitive for long signals, and the fast Fourier transform, or FFT, reduces it by exploiting the structure of the DFT matrix. The FFT divides the DFT recursively into smaller DFTs and combines their results, using the multiplicative structure of the complex roots of unity. In the radix 2 decimation in time formulation, the sum defining X_k is split into contributions from the even indexed and the odd indexed samples,

$$X_k = \sum_{m=0}^{n/2-1} x_{2m}\, e^{-\frac{2\pi i}{n}(2m)k} + \sum_{m=0}^{n/2-1} x_{2m+1}\, e^{-\frac{2\pi i}{n}(2m+1)k}. \tag{9}$$

Rearranging, each half becomes a DFT of length $n/2$ over the even and odd subsequences respectively, coupled by a single complex factor,

$$X_k = E_k + e^{-\frac{2\pi i}{n}k}\, O_k, \qquad k = 0, \ldots, \tfrac{n}{2} - 1, \tag{10}$$

where E_k is the DFT of the even indexed part of the input and O_k that of the odd indexed part. Applying this splitting recursively yields the familiar butterfly structure, in which two half length DFTs feed a stage of combining operations with the twiddle factors W_N^k, check Figure 11. The recursion reduces the cost from $O(n^2)$ to $O(n \log n)$, which is what makes real time frequency analysis feasible.

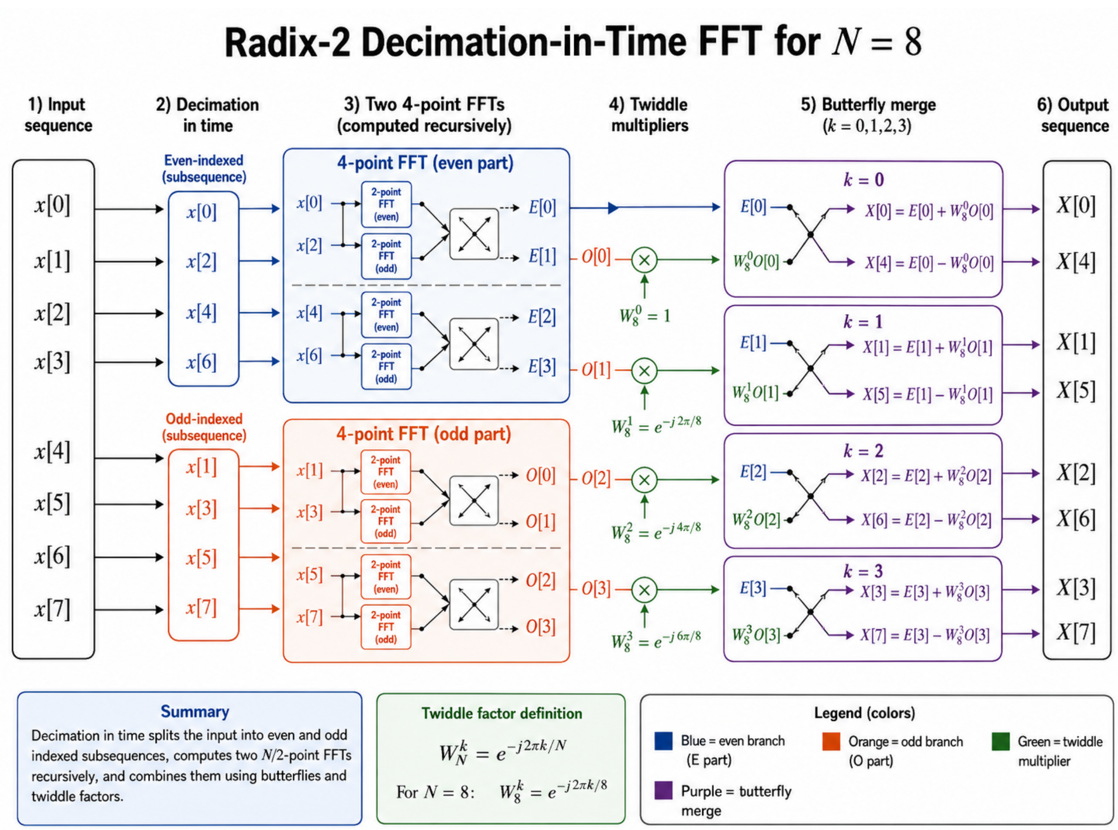

Figure 11: Radix 2 decimation in time FFT for $N = 8$. The even indexed inputs $x[0], x[2], x[4], x[6]$ feed one $N/2$ point DFT producing $E[0] \ldots E[3]$, and the odd indexed inputs $x[1], x[3], x[5], x[7]$ feed a second $N/2$ point DFT producing $O[0] \ldots O[3]$. A stage of butterfly combinations, weighted by the twiddle factors $W_N^0 \ldots W_N^7$, merges these into the outputs $X[0] \ldots X[7]$. The same construction is applied recursively within each half.

7 Filtering

Once a signal is described by its frequency content, that content can be manipulated deliberately. A filter passes some frequency components and suppresses others. The four canonical shapes (Figure 12) are the low-pass filter, which retains components below a cutoff; the high-pass filter, which retains components above a cutoff; the band-pass filter, which retains a contiguous band and rejects components on either side; and the band-stop filter, which rejects a band and passes the rest.

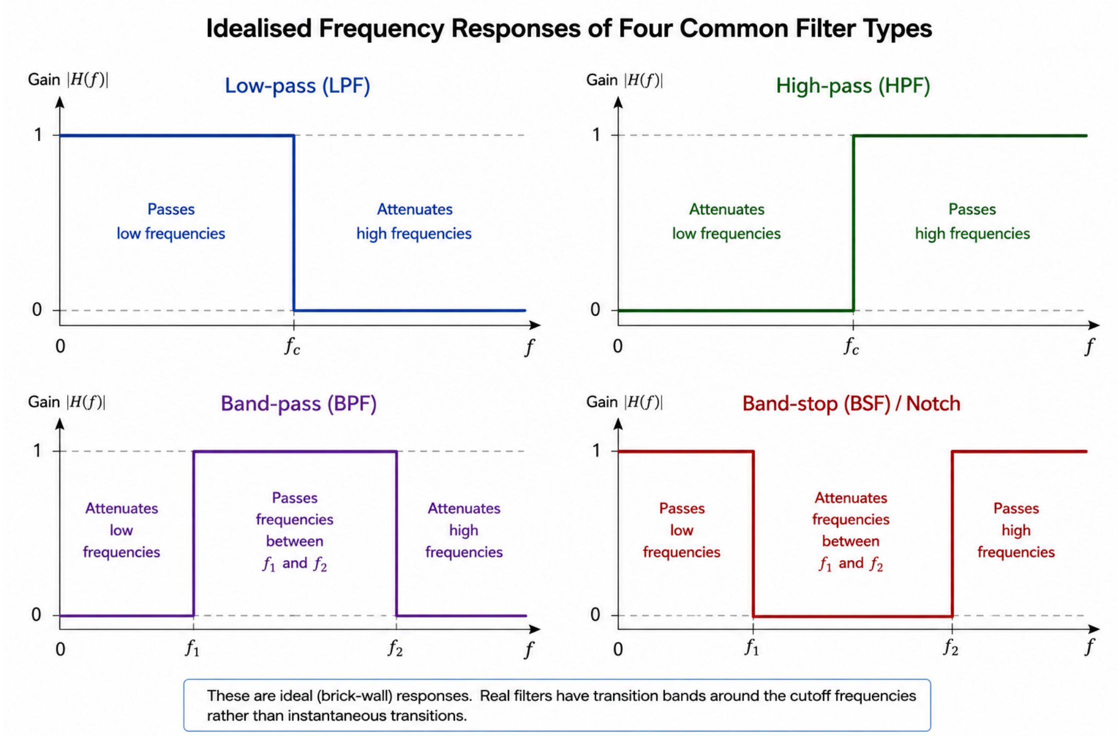

Figure 12: Idealised frequency responses of the four filter types. The low-pass response passes low frequencies and attenuates high ones; the high-pass response does the reverse; the band-pass response passes a contiguous middle band; and the band-stop response rejects a middle band while passing frequencies on either side. In each panel the horizontal axis is frequency and the vertical axis is the gain applied to each component.

Filtering is indispensable in sensing because a received signal is rarely clean. Consider an acoustic ranging system that emits a chirp from a smartphone and listens for its return. The emitted sound mixes with environmental noise across many frequencies, so the receiver passes the recorded signal through a band-pass filter tuned to the chirp's frequency band, discarding energy outside it and keeping only the components that carry information about the target. The same principle, applied with different filter shapes, underlies most of the preprocessing in a sensing pipeline.

8 Summary and Outlook

Signal propagation, distortion, and reflection allow us to sense the environment without contact. Doppler analysis identifies moving objects and movement patterns. The reflection pattern, through its frequency components and time of flight, yields the distance of an object from the transmitter and its angle of arrival. The phase shift on reflection reveals material properties. The Fourier transform is the instrument that exposes the frequency components of a signal, both for the analysis that extracts these quantities and for the preprocessing, through low-pass, high-pass, and band-pass filters, that removes unwanted components before analysis begins. The discrete and fast transforms make this analysis practical on sampled data, and the filters make it robust.

These are the common tools on which the specific sensing techniques of the course are built. Later material develops methods in the RFID, WiFi, ultra wideband, and millimeter wave bands, each of which applies the wave phenomena and the frequency domain analysis introduced here to a particular sensing problem.