ENGR 76: Information Science and Engineering — Winter 2026 — Lectures 11–12

Digital Communication, Modulation, and Demodulation

Siddharth Chandak

Over the first half of the course, we studied how to *represent information using bits*. We introduced entropy and lossless compression to quantify and remove redundancy in discrete data, then studied sampling and interpolation to understand when continuous-time signals can be represented by discrete samples without losing information. We introduced quantization, which maps real-valued samples to a finite alphabet and is the fundamental source of unavoidable distortion. In parallel, we studied frequency-domain representations using Fourier series and spectra to reason about bandwidth and degrees of freedom. Taken together, the first half of the course established how signals, discrete or continuous, can be converted into bit sequences.

The second half of the course shifts from representation to *communication*. Once information has been converted into a bit sequence, those bits must be transmitted over a physical channel. Communication introduces new challenges, including noise, interference, and physical constraints such as bandwidth and spectrum allocation. We will study how digital systems address these challenges by mapping bits to continuous-time waveforms for transmission and by introducing redundancy through error-correcting codes to achieve reliable communication.

1 Communication: Analog and Digital

Communication is the problem of sending information reliably over noisy channels. A channel may be a wire, fiber, or wireless medium, but in all cases the transmitted waveform is corrupted by noise and distortion before reaching the receiver.

In analog communication, the message itself is represented directly as a waveform. The transmitter sends a continuous-time signal through the channel, and the receiver attempts to reproduce the original waveform from a version that has been corrupted by noise and distortion.

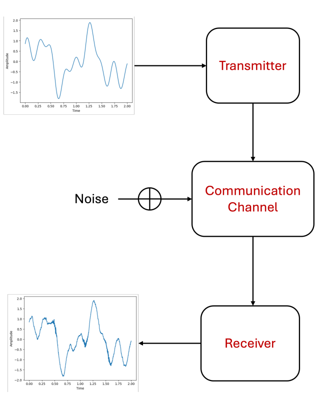

(a) Analog communication

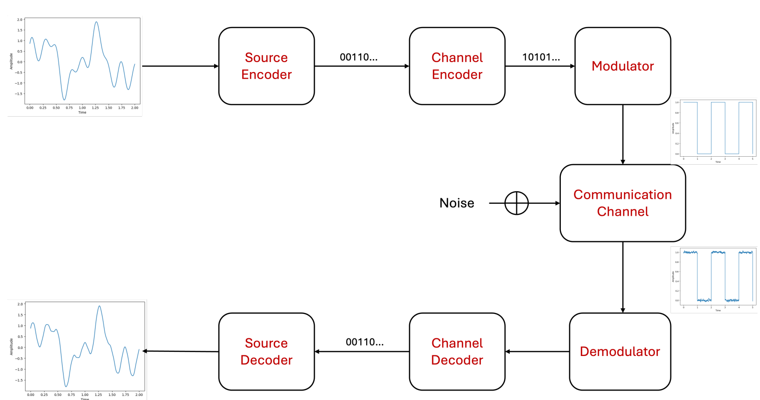

(b) Digital communication

In digital communication, the message is first converted into a bit sequence. A *source encoder* represents the signal using bits, and a *channel encoder* then adds controlled redundancy so that errors introduced by the channel can later be detected or corrected. The resulting bits are mapped to a continuous-time waveform by a *modulator*, transmitted through the physical channel where noise and distortion are introduced, and processed at the receiver by a *demodulator*, which converts the noisy waveform back into bit decisions. A *channel decoder* then corrects errors when possible, and finally a *source decoder* reconstructs the original message.

In analog communication, the message is represented by a continuous waveform, and noise perturbs this waveform continuously. Since the signal takes values in a continuum, even small amounts of noise change the received signal, and it is generally not possible to recover the original waveform exactly. In digital communication, by contrast, the receiver makes discrete decisions (for example, deciding between 0 and 1). If the noise is small enough that it does not change this decision, the bit is recovered exactly. Thus small distortions in the waveform do not necessarily cause errors. This discrete representation also enables modular system design, where compression, coding, modulation, and transmission can be designed and analyzed separately.

Modulation and Demodulation. In the digital communication pipeline, the first blocks we study are modulation and demodulation. Given a bit sequence $\{b_m\}$, modulation is the process of constructing a continuous-time signal $s(t)$ that encodes these bits in a form suitable for transmission over a physical channel. Demodulation is the reverse process: given a received waveform $y(t)$, the receiver processes it to produce estimates $\{\hat{b}_m\}$ of the transmitted bits.

2 On-Off Keying (OOK) Baseband Signal

We now introduce the simplest possible modulation scheme. The basic idea is extremely simple: during each bit interval, we either transmit a constant signal or transmit nothing at all. A bit 1 is represented by turning the signal "on" at a fixed amplitude, and a bit 0 is represented by turning the signal "off". The waveform therefore consists of rectangular pulses whose presence or absence encodes the bit sequence.

Definition 1 (On-Off Keying: Baseband Signal)**.** *Let $\{b_m\}_{m\geq 1}$ be a bit sequence with $b_m \in \{0, 1\}$. Fix amplitude $V > 0$ and bit duration $T_b > 0$. The OOK baseband waveform is*

$$x(t) = \begin{cases} V, & t \in [(m-1)T_b, mT_b) \text{ and } b_m = 1, \\ 0, & t \in [(m-1)T_b, mT_b) \text{ and } b_m = 0. \end{cases}$$

The bit duration (or symbol duration) T_b is the amount of time allocated to transmit a single bit and the bit rate R_b is defined as the number of bits transmitted per unit time:

$$R_b = \frac{1}{T_b} \quad \text{bits per second.}$$

In other words, the time axis is divided into intervals of length T_b, and within each interval the signal takes one of two possible amplitudes: either 0 or V. Information is therefore encoded purely through amplitude levels.

Example 1. *Consider the bit sequence* $1, 0, 1, 0, 1, 1, 0$. *Then the resulting OOK baseband waveform with bit duration T_b and amplitude V is as follows:*

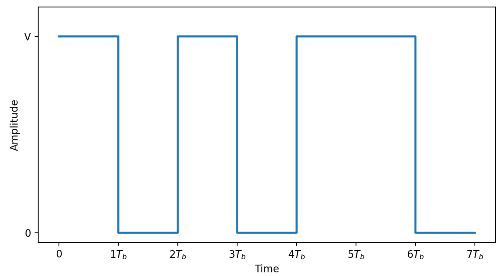

2.1 Demodulation

We now turn to the receiver. The transmitter maps bits $\{b_m\}$ to a waveform $x(t)$. The channel introduces noise and distortion, and the receiver observes a waveform $y(t)$. The goal of demodulation is to process $y(t)$ and produce estimates $\{\hat{b}_m\}$ of the transmitted bits. In other words, we must decide, for each bit interval $[(m-1)T_b, mT_b)$, whether a 0 or a 1 was transmitted.

2.1.1 Method I: Midpoint Sampling

The simplest idea is to sample the received signal once per bit interval. Since $x(t)$ is constant within each interval in OOK, a natural choice is to sample at the midpoint:

$$t_m = (m-1)T_b + \frac{T_b}{2}.$$

We then declare

$$\hat{b}_m = \begin{cases} 1, & y(t_m) > \alpha, \\ 0, & y(t_m) \le \alpha, \end{cases}$$

where α is a threshold, typically chosen to be $V/2$. If the channel is noiseless, so that $y(t) = x(t)$, this method works perfectly. Indeed, when $b_m = 1$, we have $y(t_m) = V$, and when $b_m = 0$, we have $y(t_m) = 0$. Choosing any threshold α in the interval $(0, V)$ guarantees correct decoding.

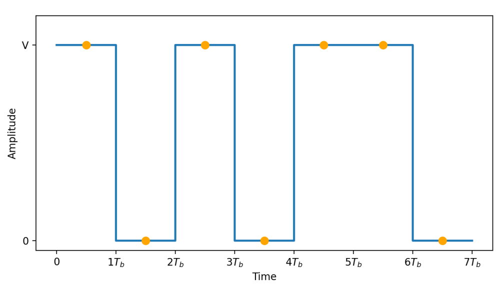

(a) Sampling in a noiseless scenario

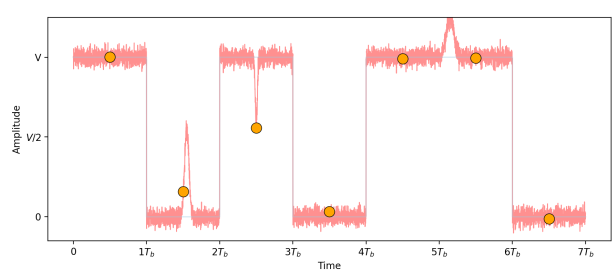

(b) Sampling in presence of noise

However, this strategy is very sensitive to noise. Since the decision depends on the value of the signal at a single time instant, a brief noise spike near the sampling time can flip the decision even if the rest of the interval is relatively clean. In the figure above, most of the third bit interval contains a signal close to V, but a noise spike near the sampling time causes $y(t_m)$ to fall below the threshold, leading to an incorrect decision. This illustrates a fundamental weakness of single-sample detection: it does not use all the information available in the interval.

2.1.2 Method II: Power Detection and Thresholding

A more robust strategy is to use the entire bit interval rather than a single sample. Since OOK encodes information through amplitude, it is natural to measure the average received power in each interval:

$$p_m = \frac{1}{T_b} \int_{(m-1)T_b}^{mT_b} y^2(t)\, dt.$$

This quantity captures the overall energy of the signal during the bit interval. The decision rule becomes

$$\hat{b}_m = \begin{cases} 1, & p_m > p_{\text{thresh}}, \\ 0, & p_m \le p_{\text{thresh}}. \end{cases}$$

Let us first understand what happens in the noiseless case, where $y(t) = x(t)$. If $b_m = 0$, then $y(t) = 0$ throughout the interval and hence

$$p_m = \frac{1}{T_b} \int_{(m-1)T_b}^{mT_b} 0\, dt = 0.$$

If $b_m = 1$, then $y(t) = V$ throughout the interval and therefore

$$p_m = \frac{1}{T_b} \int_{(m-1)T_b}^{mT_b} V^2\, dt = \frac{1}{T_b}(V^2 T_b) = V^2.$$

Thus, in the absence of noise, the average power takes one of two values: 0 when a 0 is transmitted and V^2 when a 1 is transmitted. Any threshold chosen strictly between 0 and V^2 guarantees perfect decoding.

The key advantage of power detection is that it averages the signal over the entire bit interval. Random fluctuations or short bursts of noise have a reduced effect on the average compared to their effect on a single sample. This makes power detection significantly more robust than midpoint sampling, especially in the presence of bursty or impulsive noise. In practice, noise may still push the average power across the threshold, so errors cannot be completely eliminated. However, by using the full interval rather than a single point, we make much better use of the available information and substantially reduce the probability of error.

2.2 Frequency-Domain Properties of OOK

We now analyze OOK in the frequency domain. Since modulation produces a continuous-time waveform, its frequency content determines how much bandwidth is required for transmission. Recall that a continuous-time signal of duration L seconds can be written as

$$x(t) = b_0 + \sum_{j=1}^{\infty} a_j \sin\left(2\pi \frac{j}{L} t\right) + b_j \cos\left(2\pi \frac{j}{L} t\right).$$

Its spectrum is visualized by plotting the magnitudes

$$A_j = \sqrt{a_j^2 + b_j^2}$$

as stems at frequencies j/L.

Effect of Increasing the Bit Rate. Suppose we double the bit rate. Since $R_b = 1/T_b$, doubling the rate means replacing T_b by $T_b/2$. In time, this squeezes the waveform by a factor of 2. If the original waveform is $x(t)$, the faster waveform is approximately $x(2t)$.

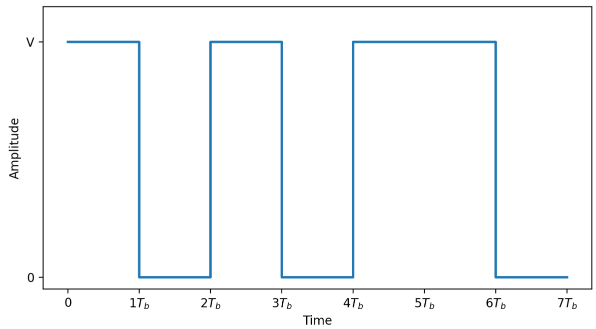

(a) $x(t)$: Bit rate $1/T_b$

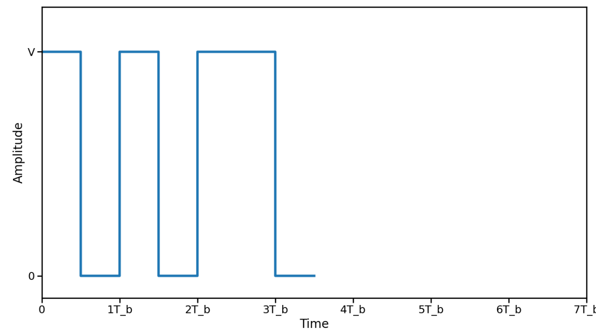

(b) $x(2t)$: Bit rate $2/T_b$

Now consider what happens to the frequency components. If $x(t)$ contains a term $\cos(2\pi ft)$, then $x(2t)$ contains $\cos(2\pi f2t) = \cos(2\pi(2f)t)$. Every frequency component is doubled. Thus, squeezing a signal in time stretches it in frequency. For example, if $z(t) = 1 + 3\cos(2\pi 5t)$, then its band is $[0, 5]$. Replacing t by $2t$ gives $z(2t) = 1 + 3\cos(2\pi 10t)$, whose band is $[0, 10]$. Squeezing the signal in time doubles its bandwidth. Since reducing T_b squeezes the OOK waveform in time, increasing the bit rate increases its bandwidth. We therefore expect the bandwidth to scale proportionally to $1/T_b$.

Actual Spectrum of OOK. Consider an OOK waveform with bit duration T_b and total length L, where typically $L \gg T_b$. The waveform consists of rectangular pulses of width T_b placed at multiples of T_b, with amplitudes either 0 or V. When we compute its Fourier series representation over the interval of length L, we obtain discrete coefficient magnitudes $A_j = \sqrt{a_j^2 + b_j^2}$ at frequencies j/L.

The magnitudes A_j follow a sinc-shaped envelope. In particular, the envelope has zero-crossings at frequencies that are integer multiples of $1/T_b$, and the first zero occurs at frequency $1/T_b$. Although the exact coefficients depend on the specific bit sequence, their overall decay pattern is governed by this envelope.

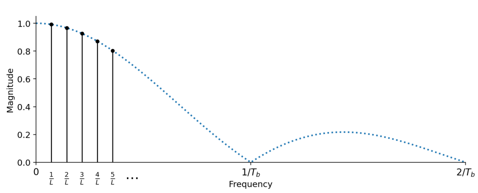

(a) Discrete spectrum with sinc envelope

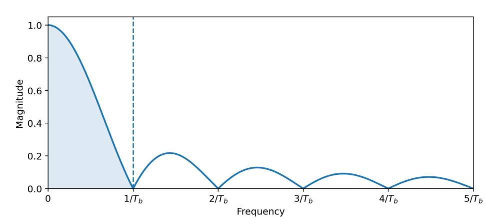

(b) Continuous view of sinc-shaped decay

Although the spectrum extends to arbitrarily large frequencies, the magnitudes decrease rapidly beyond the first zero at $1/T_b$. Most of the signal energy lies in the main lobe between 0 and $1/T_b$, and the side lobes beyond this frequency are significantly smaller. This gives us the following approximate bandlimited property of the OOK waveform.

Fact 1. *For OOK modulation with bit duration T_b (equivalently, bit rate $1/T_b$), the baseband waveform is approximately bandlimited to the frequency range $[0,\ 1/T_b]$.*

Remark 1. *This bandlimiting is only approximate. The spectrum of OOK has a sinc-shaped decay and extends to arbitrarily high frequencies, but most of the energy lies in the main lobe up to $1/T_b$. There exist other pulse-shaping schemes that produce exactly bandlimited signals. We study some of these, along with a more general modulation framework, in the appendix.*

2.3 Parameter Choice and the Role of Bandwidth

There are two primary design parameters in OOK modulation: the amplitude V and the bit duration T_b.

Choice of Amplitude V. The amplitude V determines the signal strength. Increasing V increases the separation between the two hypotheses (bit 0 and bit 1). In midpoint sampling, the received values are closer to either 0 or V. In power detection, the received average power is either close to 0 or close to V^2 (or $V^2/2$ in passband). A larger separation makes the system more robust to noise and reduces the probability of bit error. However, V cannot be increased arbitrarily. Transmitters are subject to power constraints due to hardware limitations, battery life, regulatory limits, and interference considerations. Thus, while increasing V improves reliability, it is fundamentally constrained by physical and regulatory considerations.

Choice of Bit Duration T_b. The bit duration determines the bit rate: $R_b = \frac{1}{T_b}$. To transmit information faster, we want a larger bit rate, which means making T_b smaller.

Making T_b smaller, however, reduces the amount of time available to average noise within each bit interval. In power detection, the integral

$$\frac{1}{T_b} \int_{(m-1)T_b}^{mT_b} y^2(t)\, dt$$

is computed over a shorter interval, so the averaging effect is weaker. As a result, the probability of bit error typically increases as T_b decreases. Nevertheless, this is not the most fundamental limitation. Error probability can be reduced by increasing V (within limits) or by using channel coding to introduce redundancy. The more fundamental constraint arises from bandwidth.

Bandwidth Constraint. We have seen that the OOK waveform is approximately bandlimited to $[0, 1/T_b]$. Therefore, reducing T_b to increase the bit rate necessarily increases the required bandwidth. Increasing bandwidth is costly for two major reasons.

First, physical channels are frequency-selective. Different frequency components may experience different attenuation or distortion. If the signal bandwidth exceeds the usable band of the channel, distortion and inter-symbol interference may occur. Reliable transmission therefore requires that the signal bandwidth fit within the favorable portion of the channel spectrum.

Second, bandwidth is a scarce and regulated resource, especially in wireless communication. The electromagnetic spectrum is divided among different services, and users are assigned non-overlapping frequency bands to avoid interference. A communication system must operate within its allocated band. Thus, the available bandwidth directly limits how small T_b can be, and therefore limits the achievable bit rate.

In summary, while increasing V improves reliability and decreasing T_b increases rate, the achievable bit rate is fundamentally constrained by the available bandwidth.

2.4 Why Baseband Signals Are Problematic

So far, the OOK waveform $x(t)$ we constructed is a baseband signal, meaning its frequency content lies approximately in $[0, B]$ with $B \approx 1/T_b$. In other words, most of its energy is concentrated near zero frequency. Transmitting such a baseband signal directly over the air is impractical for two fundamental reasons.

First, antenna size is inversely proportional to frequency. Roughly speaking, an efficient antenna must have a physical length on the order of the wavelength of the signal, and the wavelength is inversely proportional to frequency. Very low-frequency signals therefore require extremely large antennas. Since baseband signals concentrate energy near frequency 0, their effective wavelength is extremely large, making direct radiation physically infeasible.

Second, in practice many users transmit simultaneously over the same physical medium. If all transmitted signals occupied frequencies near zero, their spectra would overlap completely and interfere with one another. There would be no mechanism to separate different users in frequency.

Instead, communication systems allocate distinct frequency bands to different users or services. The electromagnetic spectrum is divided and regulated. For example, the FM broadcast band spans 88–108 MHz and is divided into many non-overlapping channels, each assigned to a different station. The full spectrum allocation can be seen in the frequency allocation chart here.

Because users are assigned specific frequency bands $[f_{\min}, f_{\max}]$, a baseband signal must be shifted from $[0, B]$ to lie inside its allocated band. This motivates passband modulation, where a baseband waveform is multiplied by a high-frequency carrier so that its spectrum is translated to a band centered around some carrier frequency f_c.

3 Passband Signals

So far, we constructed a baseband waveform $x(t)$ whose spectrum lies approximately in $[0, B]$ with $B \approx 1/T_b$. In other words, most of its energy is concentrated near frequency 0. However, in practice, communication systems are allocated frequency bands of the form $[f_{\min}, f_{\max}]$ with $f_{\min} > 0$. We therefore need a mechanism to move the spectrum of $x(t)$ from near zero frequency to a higher band.

The key idea is to multiply the signal by a high-frequency sinusoid. Before describing the general scheme, let us first look at a simple example to build intuition.

Example 2. *Let* $z_1(t) = 1 + \cos(2\pi 100t)$*, whose frequency band is* $[0, 100]$*. Choose* $f_c = 1000$ *and consider*

$$z_2(t) = z_1(t)\sin(2\pi 1000t).$$

Using the identity

$$\sin(a)\cos(b) = \frac{1}{2}\big(\sin(a+b) + \sin(a-b)\big),$$

we obtain

$$z_2(t) = \sin(2\pi 1000t) + \frac{1}{2}\sin(2\pi 1100t) + \frac{1}{2}\sin(2\pi 900t).$$

The original frequency band $[0, 100]$ *has been shifted to* $[900, 1100]$*, which equals* $[f_c - 100, f_c + 100]$.

This example illustrates the general phenomenon: multiplying by $\sin(2\pi f_c t)$ creates new frequency components at $f_c + f$ and $f_c - f$ for every baseband frequency f. Now consider a general signal $x(t)$ that contains a frequency component $\cos(2\pi f t)$ with $0 \le f \le B$. Multiplying by $\sin(2\pi f_c t)$ produces components at $f_c + f$ and $f_c - f$. Thus, each baseband frequency is translated symmetrically around f_c. If $x(t)$ occupies the band $[0, B]$, then after multiplication the signal occupies $[f_c - B, f_c + B]$. The width of the band becomes $2B$ and its location changes.

We now apply this frequency-translation idea directly to OOK modulation. Let $x(t)$ be the baseband OOK waveform constructed using amplitude V and bit duration T_b. To shift its spectrum from $[0, 1/T_b]$ to a band centered around f_c, we multiply by a high-frequency sinusoid and define

$$s(t) = x(t)\sin(2\pi f_c t).$$

The sinusoid $\sin(2\pi f_c t)$ is called the **carrier wave**, and f_c is the **carrier frequency**.

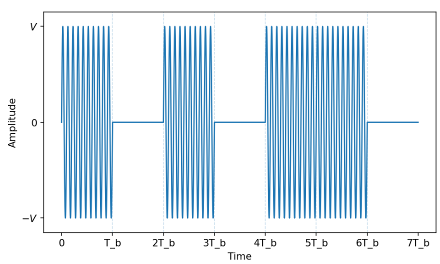

(a) Passband signal $s(t)$

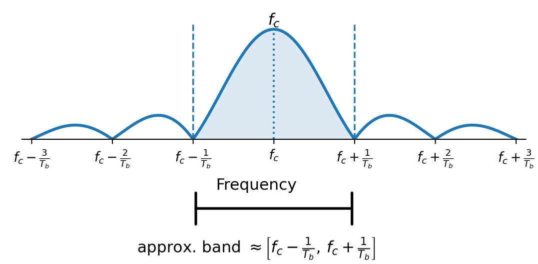

(b) Spectrum of $s(t)$

When $b_m = 1$, the transmitted signal over the interval $[(m-1)T_b, mT_b)$ is $V\sin(2\pi f_c t)$. When $b_m = 0$, the transmitted signal is 0. Thus, the carrier is turned on and off according to the bit sequence. The baseband waveform $x(t)$ acts as a slowly varying envelope that controls the amplitude of the rapidly oscillating carrier.

Fact 2 (Bandwidth of OOK Modulation)**.** *For OOK modulation with bit duration* T_b *(equivalently, bit rate* $1/T_b$*), the baseband waveform* $x(t)$ *is approximately bandlimited to the frequency range* $[0, 1/T_b]$*. The passband signal constructed using* $s(t) = x(t)\sin(2\pi f_c t)$*, where* $f_c > 1/T_b$*, is approximately bandlimited to the frequency range* $[f_c - 1/T_b, f_c + 1/T_b]$.

The bit rate remains $R_b = 1/T_b$. For baseband OOK, the approximate band is $[0, 1/T_b]$, so the bandwidth is $1/T_b$. For passband OOK, the approximate band is $[f_c - 1/T_b, f_c + 1/T_b]$, so the bandwidth is $2/T_b$. Thus, for passband transmission, the bandwidth equals $2R_b$. Increasing the bit rate necessarily increases the required bandwidth.

3.1 Choosing Carrier Frequency and Bit Duration Based on an Allocated Band

Suppose we are allocated a frequency band $[f_{\min}, f_{\max}]$ for transmission. We would like the passband OOK spectrum to lie entirely inside this band. Since passband OOK occupies approximately $[f_c - 1/T_b, f_c + 1/T_b]$, we require

$$f_c - \frac{1}{T_b} \geq f_{\min} \quad \text{and} \quad f_c + \frac{1}{T_b} \leq f_{\max}.$$

To utilize the band efficiently, we center the signal within the allocated range by choosing the carrier frequency as the midpoint,

$$f_c = \frac{f_{\min} + f_{\max}}{2}.$$

Let $B = f_{\max} - f_{\min}$ denote the bandwidth of the allocated channel. With the midpoint choice of f_c, the conditions above reduce to $2/T_b \leq B$. Thus, the maximum achievable bit rate satisfies

$$R_b = \frac{1}{T_b} = \frac{B}{2}.$$

This shows that the bit rate depends only on the width of the allocated band, not on its location. The carrier frequency determines where we transmit, while the available bandwidth determines how fast we can transmit.

3.2 Demodulation of Passband OOK

We now analyze power detection for passband OOK in the noiseless case, that is, assuming the received signal satisfies $y(t) = s(t)$.

Suppose $b_m = 1$. Over the interval $[(m-1)T_b, mT_b)$, the transmitted signal is $s(t) = V \sin(2\pi f_c t)$. The receiver computes the average received power

$$p_m = \frac{1}{T_b} \int_{(m-1)T_b}^{mT_b} s^2(t)\, dt = \frac{V^2}{T_b} \int_{(m-1)T_b}^{mT_b} \sin^2(2\pi f_c t)\, dt.$$

Why does this strategy work? The carrier $\sin(2\pi f_c t)$ oscillates rapidly and changes sign every half cycle. If we simply integrated $s(t)$ over the interval, the positive and negative portions would largely cancel. Squaring removes the sign changes and converts the rapidly oscillating sinusoid into a nonnegative signal whose average value over a full cycle is $1/2$. Thus, power detection measures whether carrier energy is present during the bit interval.

Assume that T_b contains an integer number of carrier cycles. Over one carrier period,

$$\int_0^{1/f_c} \sin^2(2\pi f_c t)\, dt = \frac{1}{2f_c}.$$

This follows from $\sin^2\theta = \frac{1-\cos(2\theta)}{2}$ and integrating over one full period. Since there are $T_b f_c$ complete cycles in one bit interval,

$$\int_{(m-1)T_b}^{mT_b} \sin^2(2\pi f_c t)\, dt = (T_b f_c) \cdot \frac{1}{2f_c} = \frac{T_b}{2}.$$

Substituting back gives

$$p_m = \frac{V^2}{T_b} \cdot \frac{T_b}{2} = \frac{V^2}{2}.$$

If $b_m = 0$, then $s(t) = 0$ and therefore $p_m = 0$. Thus, in the noiseless case, the average power takes only two possible values: 0 when the carrier is absent and $V^2/2$ when the carrier is present. Any threshold chosen in $(0, V^2/2)$ therefore ensures perfect decoding.

3.3 Synchronization

So far, we have assumed that the receiver knows exactly where each bit interval begins, that is, we assumed perfect alignment of the intervals $[(m-1)T_b, mT_b)$. In practice, this assumption is unrealistic. The received signal may begin with an unknown delay, and the receiver does not know where the first bit starts. If the bit boundaries are misaligned, power detection fails. The integral

$$p_m = \frac{1}{T_b} \int_{(m-1)T_b}^{mT_b} y^2(t)\, dt$$

would then average over portions of two adjacent bits. Even in the noiseless case, this mixing can produce intermediate power values that do not clearly correspond to either 0 or $V^2/2$, leading to incorrect decisions. Therefore, before demodulating the bits, the receiver must first determine the correct alignment of the bit intervals. This process is called synchronization.

A simple synchronization scheme is the following. We assume that the first transmitted bit is always a 1. The receiver continuously monitors the incoming signal and detects the first time the received power exceeds a threshold. This moment is taken as the start of the first bit interval. Once the starting point is determined, subsequent intervals are spaced at multiples of T_b.

This approach works reasonably well when the signal-to-noise ratio is sufficiently high and when the channel distortion is limited. But it is particularly vulnerable to bursty noise. A short, high-energy disturbance can exceed the detection threshold and be mistakenly interpreted as the beginning of transmission. Since synchronization errors affect all subsequent bit decisions, such false detections can corrupt an entire packet of data. In practice, more sophisticated synchronization schemes are often used to improve reliability and robustness. We briefly describe one such idea in the appendix.

4 Multiple Simultaneous Transmissions

Earlier, we noted that in practical communication systems many users transmit simultaneously. For example, in radio broadcasting, different stations operate at different carrier frequencies within the same overall spectrum. The signals coexist in the medium and are separated at the receiver using frequency-selective filtering.

To model this mathematically, suppose two independent transmitters send signals at carrier frequencies f_{c_1} and f_{c_2}. Let their respective baseband waveforms be $x_1(t)$ and $x_2(t)$, constructed using bit durations T_{b_1} and T_{b_2} (these may be equal or different). The corresponding passband signals are

$$s_1(t) = x_1(t)\sin(2\pi f_{c_1} t), \qquad s_2(t) = x_2(t)\sin(2\pi f_{c_2} t).$$

The total transmitted signal in the medium is the sum

$$s(t) = s_1(t) + s_2(t).$$

The receiver observes

$$y(t) = s(t) + n(t) = s_1(t) + s_2(t) + n(t),$$

where $n(t)$ represents noise. The key point is that $y(t)$ contains both transmissions simultaneously. Without additional processing, the receiver cannot directly apply the demodulation scheme described earlier, because the signal from one transmitter acts as interference for the other.

4.1 Frequency Domain

Each passband signal occupies a distinct frequency band. If the first transmission has bit duration T_{b_1} and carrier frequency f_{c_1}, then its spectrum lies approximately in the interval $[f_{c_1} - 1/T_{b_1},\, f_{c_1} + 1/T_{b_1}]$. Similarly, if the second transmission has bit duration T_{b_2} and carrier frequency f_{c_2}, then its spectrum lies approximately in $[f_{c_2} - 1/T_{b_2},\, f_{c_2} + 1/T_{b_2}]$.

To ensure that the two transmissions do not interfere with one another in frequency, these bands must not overlap. Without loss of generality, suppose $f_{c_2} > f_{c_1}$. A sufficient condition for non-overlap is $f_{c_1} + \frac{1}{T_{b_1}} < f_{c_2} - \frac{1}{T_{b_2}}$. Equivalently, the separation between the carrier frequencies must satisfy $f_{c_2} - f_{c_1} > \frac{1}{T_{b_1}} + \frac{1}{T_{b_2}}$. When this condition holds, the two occupied frequency intervals are disjoint, and frequency-selective filtering can isolate each transmission.

Example 3 (Two simultaneous transmissions and frequency separation)**.** *Consider two users transmitting simultaneously over the same medium. User 1 uses bit rate* $R_{b_1} = 100$ *so* $T_{b_1} = 0.01$*, carrier frequency* $f_{c_1} = 1000$ *Hz, and amplitude* $V_1 = 1$*. User 2 uses bit rate* $R_{b_2} = 150$ *so* $T_{b_2} = 1/150 \approx 0.0067$*, carrier frequency* $f_{c_2} = 3000$ *Hz, and amplitude* $V_2 = 0.8$*. Let the bit sequences be* $(1, 0, 1, 1)$*, and* $(0, 1, 1, 0, 1, 0)$ *for the first and the second user, respectively.*

The spectrum for user 1's signal occupies approximately

$$[f_{c_1} - 1/T_{b_1},\, f_{c_1} + 1/T_{b_1}] = [900, 1100],$$

while user 2 occupies approximately

$$[f_{c_2} - 1/T_{b_2},\, f_{c_2} + 1/T_{b_2}] = [2850, 3150].$$

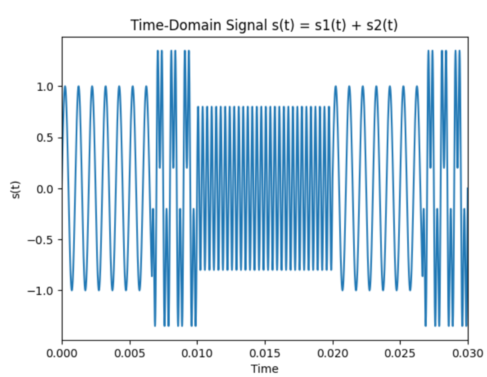

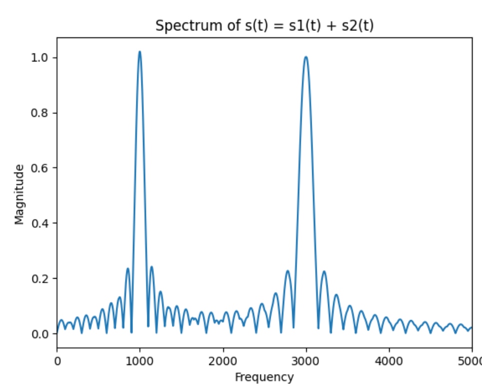

In the time domain, the received signal is the sum $s(t) = s_1(t) + s_2(t)$*. The two waveforms overlap and add pointwise, so there is no obvious way to separate the contributions of the two transmitters from the time-domain signal alone. In the frequency domain, however, the situation* *is different. The spectrum of* $s_1(t)$ *is concentrated approximately in* $[900, 1100]$*, while that of* $s_2(t)$ *lies in* $[2850, 3150]$*. These intervals are disjoint, so although the signals overlap in time, they are cleanly separated in frequency.*

4.2 Bandpass Filters

Definition 2 (Ideal Bandpass Filter)**.** *An ideal bandpass filter with passband* $[f_1, f_2]$ *is a system that, in the frequency domain, leaves the spectrum unchanged for* $f \in [f_1, f_2]$ *and sets the spectrum to* 0 *for all other frequencies.*

Thus, if a signal has spectrum $X(f)$, the output spectrum equals $X(f)$ for $f \in [f_1, f_2]$ and equals 0 elsewhere. In other words, frequencies inside the band remain as they are, while frequencies outside the band are completely removed. An example of the operation of the bandpass filter is shown in the figure below.

(a) Spectrum of signal (b) Spectrum after bandpass filter with band $[4, 7]$ Hz

4.3 Demodulation

Return now to the received signal $y(t) = s_1(t) + s_2(t) + n(t)$. Assume that the occupied bands of the two transmissions are disjoint. To recover the first transmission, the receiver applies an ideal bandpass filter with passband $[f_{c_1} - 1/T_{b_1},\ f_{c_1} + 1/T_{b_1}]$. Since $s_1(t)$ lies in this band while $s_2(t)$ lies outside it, the resulting signal $y_1(t)$ contains $s_1(t)$ (together with noise components inside the same band) and removes $s_2(t)$. Once $y_1(t)$ has been isolated, demodulation proceeds exactly as in the single-user case. For passband OOK, the receiver computes $p_m = \frac{1}{T_{b_1}} \int_{(m-1)T_{b_1}}^{mT_{b_1}} y_1^2(t)\, dt$ and compares p_m with a threshold to decide whether the transmitted bit was 0 or 1.

The same procedure is applied independently to recover the second transmission: first isolate its frequency band using a bandpass filter centered at f_{c_2}, then apply the standard power-detection rule with bit duration T_{b_2}.

Remark 2. *Even in single-user transmission, bandpass filtering is useful. Suppose a transmitted signal occupies only the band* $[f_c - 1/T_b,\ f_c + 1/T_b]$*, while the noise* $n(t)$ *is spread across a much wider range of frequencies. The received signal is* $y(t) = s(t) + n(t)$*, and although* $s(t)$ *is confined to its narrow band, the noise may contain components at many other frequencies.*

Applying a bandpass filter that keeps only $[f_c - 1/T_b,\ f_c + 1/T_b]$ *removes all noise components outside the signal band while leaving the desired signal essentially unchanged. As a result, the total noise power at the receiver is reduced, while the signal power remains approximately the same. This increases the signal-to-noise ratio before demodulation and improves the reliability of the subsequent thresholding step.*

A Appendices

The following appendices contain supplementary material that is not required for this course and is provided for additional reading.

A General Modulation and Pulse Shaping

In the main text, we constructed baseband OOK by assigning a constant amplitude over each bit interval. While this construction is simple, it is only one specific instance of a much more general framework for building communication signals. In this appendix, we describe that general framework and show how OOK fits into it.

A.1 General Pulse-Shaping Framework

Recall from the sampling and interpolation part of the course that a continuous-time signal can be reconstructed from discrete samples by forming a weighted sum of shifted interpolation kernels. Modulation can be understood in a very similar way. Let $\{b_m\}$ be the transmitted bit sequence with $b_m \in \{0, 1\}$. Instead of transmitting a constant signal over each interval, we choose a fixed waveform $F(\cdot)$ and construct the baseband signal as

$$x(t) = \sum_{m=1}^{\infty} b_m \, F\left(\frac{t - mT_b}{T_b}\right).$$

Here, $F(\cdot)$ is a pulse shape that determines how each bit influences the transmitted waveform. For each m, we place a shifted and time-scaled copy of F centered at $t = mT_b$. The bit b_m either activates that pulse (if $b_m = 1$) or removes it (if $b_m = 0$).

This construction is directly analogous to interpolation. The bits $\{b_m\}$ play the role of discrete samples, while the function F plays the role of an interpolation kernel. The choice of F determines the spectral properties of the transmitted signal.

A.2 OOK as a Special Case

In OOK, the pulse shape is rectangular. Define

$$F(u) = \begin{cases} V, & 0 \le u < 1, \\ 0, & \text{otherwise.} \end{cases}$$

Then the general formula $x(t) = \sum_{m=1}^{\infty} b_m \, F\left(\frac{t-mT_b}{T_b}\right)$ reduces exactly to the OOK waveform defined in the main text.

A.3 Sinc Pulse and Exact Bandlimiting

Instead of choosing a rectangular pulse, we can choose the pulse shape F to be a sinc function. In normalized form, consider

$$F(u) = V \operatorname{sinc}(u),$$

so that the baseband signal becomes

$$x(t) = \sum_{m=1}^{\infty} b_m \, V \operatorname{sinc}\left(\frac{t - mT_b}{T_b}\right).$$

With this choice, the resulting signal is **exactly** bandlimited to the frequency range $[0, 1/(2T_b)]$. In contrast to OOK, whose spectrum has a sinc-shaped decay and extends to arbitrarily high frequencies, this construction produces a signal whose spectrum is strictly confined to a finite band.

The intuition comes from the Fourier transform pair between rectangular functions and sinc functions. Earlier, we saw that a rectangular pulse in time produces a sinc-shaped spectrum that decays slowly and is not exactly bandlimited. The roles can be reversed: a sinc-shaped pulse in time corresponds to a rectangular spectrum in frequency. Thus, choosing a sinc pulse in time forces the spectrum to be perfectly confined to a finite frequency interval.

This is exactly the same principle as sinc interpolation. There, shifted sinc functions combine to reconstruct a signal whose spectrum is limited to a fixed band. Here, shifted sinc pulses combine to produce a communication signal that is exactly bandlimited.

The advantage of this construction is precise bandwidth control. The disadvantage is that the sinc pulse extends infinitely in time and decays slowly. Each transmitted bit therefore affects the waveform over an infinite time horizon, making exact implementation impossible and practical implementation challenging. In practice, one truncates or approximates the sinc pulse, which reintroduces slight spectral leakage.

A.4 Other Common Pulse Shapes

Among practical pulse shapes, the most important is the raised-cosine pulse. It is exactly bandlimited and is designed so that symbols transmitted at spacing T_b do not interfere with one another at the sampling instants. A parameter called the roll-off factor controls how much excess bandwidth beyond $1/(2T_b)$ is used. Larger roll-off leads to better time localization but increases bandwidth slightly.

In practice, many systems use a root-raised-cosine filter at both the transmitter and the receiver, so that their combined effect produces a raised-cosine response. This preserves the desired bandwidth properties while improving performance. Other pulse shapes, such as Gaussian pulses, are used in some specialized systems because of their smooth time and frequency behavior, but the raised-cosine family remains the most widely used in modern digital communication.

B Synchronization Schemes

The synchronization scheme described in the main text and implemented in the project is intentionally simple. We assumed that the transmission begins with a known bit (for example, a 1), and the receiver declares the first time the received power exceeds a threshold as the start of the first bit interval. Subsequent intervals are then spaced at multiples of T_b.

While this approach is easy to implement and works in controlled settings, it is not robust in realistic environments. First, noise may produce large fluctuations that momentarily exceed the threshold, causing the receiver to lock onto an incorrect starting point. This is especially problematic in the presence of bursty noise, where short but high-energy disturbances can mimic the presence of a transmitted bit. Second, if the first few bits happen to contain low energy (for example, several consecutive zeros), the receiver may fail to detect the true start of transmission.

For these reasons, practical systems employ more sophisticated synchronization mechanisms.

B.1 Preamble-Based Synchronization

A widely used approach is to prepend a known bit sequence, called a preamble, to each transmission. Let the preamble consist of K known bits $p_1, \ldots, p_K$. The transmitted baseband signal over the preamble interval is

$$x_p(t) = \sum_{m=1}^{K} p_m \, F\left(\frac{t - mT_b}{T_b}\right).$$

The receiver does not know the starting time, so it considers candidate offsets τ and computes a correlation metric

$$C(\tau) = \int y(t) \, x_p(t - \tau) \, dt.$$

When τ equals the true delay, the received waveform aligns with the expected preamble pattern and $C(\tau)$ becomes large. For incorrect offsets, the correlation is significantly smaller. The receiver declares synchronization at the value of τ that maximizes $C(\tau)$.

For example, suppose the preamble is $1, 0, 1, 1$. If the receiver integrates over four consecutive candidate intervals and observes high, low, high, high energy in that order, the pattern matches and synchronization is declared. Because several bits must match simultaneously, the probability that noise alone produces a false synchronization is much smaller than in the single-bit threshold scheme.

In practice, the choice of the preamble sequence matters. Some bit patterns produce clearer correlation peaks than others. Ideally, we would like a preamble that matches very strongly when perfectly aligned and matches poorly for all incorrect shifts. This reduces the chance that noise or partial overlap produces a false synchronization. Certain short sequences, known as Barker codes, are designed to have this property. Informally, when such a sequence is correlated with shifted versions of itself, the perfectly aligned position produces a much larger value than any misaligned position. For example, the length-7 Barker sequence $1, 1, 1, 0, 0, 1, 0$ produces a sharp correlation peak at the correct alignment and much smaller values elsewhere. Using such structured preambles makes synchronization significantly more robust than using an arbitrary bit pattern, while keeping the implementation simple.