


Baseband Transmitters and Receivers

This chapter covers the electrical aspects of the physical layer of baseband communication systems.

After this chapter you should be able to: compute noise margins and error rates for Gaussian noise; explain two advantages of current loop signalling; compute common-mode and differential voltages; define, calculate and explain the purpose of slew-rate limiting; specify source and load impedances that avoid reflections; select and design the most appropriate bus driver technology for a bus with multiple drivers; list some functions of line drivers and receivers; justify a choice between polled and contention-based buses; explain how transformers can be used to: interface between balanced and unbalanced transmission lines, separate common-mode and differential signals, provide protection from DC or low-frequency AC; explain the purpose for, and design an optoisolator circuit.

Decision Thresholds

A receiver converts the received waveform into bits. The simplest receivers operate on multi-level baseband signals by comparing the received voltage to one or more “decision thresholds”. For example, if the possible received levels were $+V$ and $-V$ and the noise was symmetrical and zero-mean (equally likely to be positive or negative) the optimum¹ threshold would be set at zero. If there were four levels we would need three thresholds.




A circuit called a comparator could be used to convert the analog input to a digital (logic-level) signal. If the symbols are at more than two levels we would need more than one comparator and some logic (e.g. a priority encoder) to convert the input level to a multi-bit digital output.

In most cases the received signal is integrated (added up) over the duration of the symbol to reduce the effect of noise. For example, the signal could be sampled several times over one symbol and these sample values added up. The signal components would all add in the same direction but the (zero-mean) noise components would tend to cancel. The effect of integration is similar to low-pass filtering. It is the integrated result that is compared to thresholds at the end of the symbol.

¹Optimum in the sense of minimizing the error rate.

Noise Margin

Noise margin is the noise level required to cause an error. For example, a bipolar signal received with levels of $\pm 1V$ would have a noise margin of $1V$ because a noise voltage of $1V$ (in the wrong direction) would cause the received signal to cross the decision threshold and cause an error.



Error Rates

How often does noise cause an error? This depends on the transmitted levels, the noise margins and probability distribution of the noise. An error will happen whenever the noise voltage causes the received level to cross a decision threshold. For Gaussian noise, we can compute the likelihood of an error (the error rate) using the Gaussian CDF.


In the general case we need to consider all possible error events (threshold crossings) and their likelihoods.

Exercise 1: Gaussian noise with a mean of $0.5V$ and a variance of $0.25V^2$ is added to a bipolar signal with levels of $\pm 1V$. Assuming a decision threshold equally spaced between the two levels, what is the likelihood of error if $+1$ is transmitted? If -1 is transmitted? What is the average error rate if $+1$ is transmitted 25% of the time?

Current Loops

Data can be transmitted by using different currents rather than different voltages.

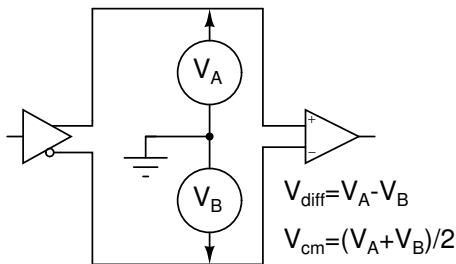
Current signalling uses a current source at the transmitter instead of a voltage source. The current swing stays constant regardless of the resistance of the line. The voltage across the terminating resistor at the receiving end always has the same voltage swing regardless of the length of the cable. The current loop's noise margin is thus independent of the length (resistance) of the circuit.



The use of current loops also makes it easier for the transmitter to determine when the loop has failed – when the loop becomes an open circuit the driver output voltage goes to the maximum possible value.

Differential Signalling

The average of the voltages on two conductors, with each voltage referenced to ground, is called the *common-mode* voltage. The voltage difference between these two conductors is the *differential* voltage.



A differential transmitter encodes data as the differential voltage – the voltage difference between its two outputs. For example, $V_A = +5\text{ V}$ and $V_B = 0\text{ V}$ for a logical '1' and $V_A = 0\text{ V}$ and $V_B = +5\text{ V}$ for a logical '0'.

Exercise 2: What are the differential and common-mode voltages for this example?

A differential receiver measures the differential voltage by subtracting the two voltages (each measured relative to ground). Note that the differential voltage can be negative even though neither of the two voltages is negative relative to ground.


Data transmitted using the differential voltage (or current) is called *differential signalling*². Differential signalling is commonly used at higher speeds and longer distances where the use of a shared (“common”) ground path can add significant amounts of noise as described below.

Noise From Shared Grounds

Some communication channels such as RS-232 transmit over one conductor and use a return path (a “ground” connection) that is shared with other circuits. These other circuits could be other communication links or could be used for power distribution.

A problem with a common ground is that there will be a voltage shift in the apparent “ground” voltage level equal to the sum of the return currents multiplied by the resistance of the return path. This results in an offset voltage at the receiver that is proportional to the product of the sum of the currents on the return path and the resistance of the return path. This voltage will be superimposed on the signal. Therefore the common-ground approach is only practical when this offset voltage is small relative to the signal levels.

For example, in the figure below the current on one circuit (i_1) affects the “ground” voltage (V_g) and thus the voltage seen on the second circuit (V_2):



These effects can be particularly severe when the shared ground circuit is used as the return path for a power supply.

Differential signalling, such as used with POTS³, RS-485, or twisted-pair Ethernet, avoids this problem because the receiver only measures the voltage difference between the two conductors. Since offsets in the ground voltage affect the voltage on both conductors

²Not to be confused with differential line codes which encode data as the difference between successive symbols.

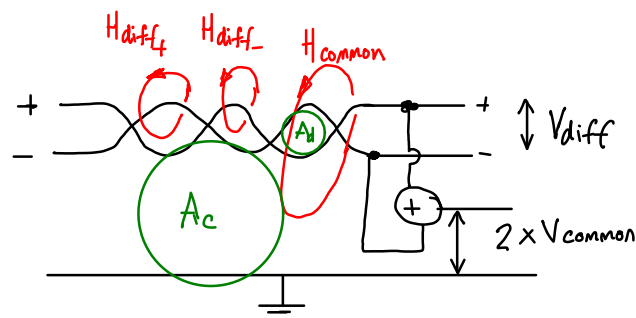
³Plain Old Telephone Service – “landlines”.

equally, this common-mode voltage does not affect the differential voltage.

Of course, the drawback is that each communication link requires two conductors instead of one.

Inductively-Coupled Noise

The following diagram shows a twisted-pair transmission line⁴ and a ground return path. The differential and common-mode voltages are shown. Also shown are the two magnetic field loops (H_{diff+} and H_{diff-} and H_{common} which could cause differential and common-mode noise to be induced on the conductors (imagine the magnetic field direction perpendicular to the page):



The two conductors of the twisted pair are next to each other and the area between the wires (A_d) is much smaller than the area between the wires and ground (A_c). This will result in a smaller voltage being induced on the differential signal. In addition, the voltages induced on the twisted pair by the magnetic field (the parts of the field H_{diff+} and H_{diff-}) will cancel out due to the twisting.

Thus the use of differential signalling over twisted pair reduces inductively-coupled noise both compared to common-mode signalling and compared to differential signalling without twisted pair.

Shielding and Grounding


Some differential communication cables also use a shield⁵) to limit radiation from the signals carried by the cable. To avoid radiation or inducing noise on

⁴A twisted-pair transmission line consists of two conductors twisted around each other.

⁵Although the construction is similar to co-axial cables, the differential signals do not use TEM propagation.

other circuits, it's desirable to limit currents flowing along the shield.

The diagram below shows some approaches that can be used to limit current flow along the shield. In some cases the shield can be left disconnected at one end. It can also be connected through a capacitor to block low-frequency currents (if the main concern is low-frequencies, typical from 50/60 Hz line voltage) or placing a ferrite cylinder around the shield (if the concern is RF radiation).



Slew Rate


Slew rate is the rate at which a voltage changes as it switches between levels. It is typically specified in units of $V/\mu s$.

Limiting the slew rate makes transitions between voltage levels more gradual. Since the transmission line has capacitance that must be charged up, by reducing the slew rate we reduce the current that must be supplied to the transmission line.

Exercise 3: What is the current flowing into a 1nF capacitor if it is being charged at a rate of $10V/\mu s$?

Reduced current flow also reduces the high-frequency components that are more likely to radiate from transmission lines and cause interference to other devices.

However, limiting the slew rate also limits the data rate because of the time it takes the signal to transition between voltage levels. Ideally, the slew rate is matched to the data rate so that the transition between voltage levels is a small fraction of the symbol duration.




Exercise 4: The RS-232 standard specifies a maximum slew rate of $30V/\mu s$. Assuming a voltage swing of 30 volts, what is the maximum

data rate for which two signal level transitions occupy 10 % of the bit period?

Impedance Matching and Terminations

The source (transmitter) and load (receiver) impedances should be the complex conjugates of the transmission line characteristic impedance to maximize power transfer and avoid reflections. Typically the characteristic impedance of a transmission line is resistive (real).

Often the receiver has a high impedance and a resistor is used to provide the appropriate termination. Terminations (often called “bus terminations” for multi-drop or multi-conductor topologies) are used at the ends of transmission lines to prevent reflections that would cause inter-symbol interference.



Passive terminations are resistors of a resistance equal to the characteristic impedance of the transmission line.

Active terminations connect the terminating resistor(s) to a low-impedance voltage source. The voltage is chosen to minimize the average current⁶.


Tri-State and Open-Collector Outputs

Bus-type network topologies consist of multiple devices connected in parallel to the same transmission line. There are two ways that bus drivers avoid having two outputs connected together.


The first method is to use open-collector (or open drain) outputs. The open-collector outputs can pull the line low but cannot source current or voltage. An external pull-up resistor sets the bus voltage high when none of the open-collector outputs are active. Open-collector outputs are simple to implement and provide all devices an indication that the bus is connected and active. However, the pull-up resistor results in constant power consumption and the RC time constant limits data rates.

⁶Power consumption in terminations is more of an issue for multi-conductor buses.

Exercise 5: If the capacitance of the transmission line joining several OC drivers is 1 nF and the pull-up resistor is 1 k Ω , how long will it take for the pull-up to pull the line from 0V to 63% of the logic high voltage?



The second method is to use bus drivers with tri-state outputs. In addition to high- and low-voltage outputs, tri-state outputs also have a high-impedance output state that effectively disconnects the output from the bus. The design of a bus with tri-state drivers is more complex because some mechanism must be provided to switch each driver on and off at the appropriate times. This mechanism is usually part of a multiple-access protocol (e.g. CSMA or TDMA) that allocates the bus to different devices at different times



RS-485 is a standard similar to RS-422 that allows multiple drivers to be connected to a multi-drop bus. Only one device can have an active output at any given time – the others must be in a high-impedance mode.

Bus Contention Protocols

One way to ensure that only one device at a time drives the bus is to select one device to act as the bus master. The other devices are called slaves. The master can then “poll” a slave to give it permission to drive the bus. This is typically done by sending a control message addressed to a specific slave.

Exercise 6: What are the consequences of increasing the delay between polls? What other factor might determine the maximum delay before slave gets access to the bus in a system using polling?

Although polling simplifies the design of slave-only devices, it can be wasteful of bus bandwidth and/or can result in long delays before a slave is able

to transmit. Because of this, other bus arbitration techniques have been developed to allow different devices to contend (fight) for access to the channel. This is usually done by having devices detect the effect of a collision (when multiple devices try to drive the line at the same time).

For example, in the case of an open-collector output, the high (“recessive”) state is when the line is high. If the bus is at the low (“dominant”) voltage but a specific device is not pulling the line low itself then this implies that some other device must be using the bus. Different protocols use this fact to handle contention in different ways.

Some contention-based bus arbitration protocols also implement the concept of priority. This enables some devices to obtain access to the bus ahead of others when more than one device wants to access the bus at the same time. This is important in real-time systems where it is necessary to guarantee a maximum delay for transmission of a message between devices.

Exercise 7: Consider a communication bus in a car that connects an airbag activation controller with a collision detector, a passenger-seat occupancy sensor and an airbag-disabling switch. Would it be more appropriate to use a polling- or contention-based bus arbitration protocol? Would it be appropriate for the arbitration protocol to allow different priorities for bus access? If so, what priorities might be assigned the different sensors?

Many (dozens of) bus contention mechanism have been developed that provide different tradeoffs between complexity and performance.

Line Drivers and Receivers

Line drivers and receivers are ICs that interface digital logic circuits to transmission lines. The primary functions of these ICs are:

- convert to/from the voltage or current levels on the transmission line to logic-level signals (TTL, CMOS, etc).
- match the impedance of the transmission line to minimize reflections and maximize power transfer
- limit slew rate to limit interference to other devices


- provide protection from accidental short-circuits and overvoltage or overcurrent conditions
- provide switchable tri-state outputs or open-collector outputs if used on bus-type connections
- detect invalid voltage/current levels or disconnected lines when possible
- detect disconnected or floating transmission lines
- convert unbalanced logic level signals to/from balanced transmission lines (if appropriate)

We will study a typical line driver and receiver in the lab.

Transformers

Transformers are widely used in baseband communications systems for various reasons.

Transformers can convert between balanced and unbalanced circuits. In this configuration the transformer is called a “balun”. If the balanced side of the transformer has a center-tap then this tap will be at the common-mode voltage. The common-mode circuit can be grounded, used to carry power (as in “power over ethernet”) or as a secondary (typically low-speed) signaling circuit.



Exercise 8: If the common-mode circuit is used to carry 500mA, how much current flows through each half of the transformer secondary? What is the net effect on the flux in the transformer core?

Transformers can also convert impedances. The voltage ratio is the same as the turns ratio while the impedance ratio is the square of the turns ratio.

Transformers also protect logic circuits against high DC (or low frequency AC) voltages that might be accidentally applied to the transmission line. Each


transformer will have a rated maximum isolation between primary and secondary. Typically this is on the order of 1kV or more. The term used for the transformers used for data interfaces such as 100BT Ethernet is “magnetics”.

The main disadvantages of transformers are that (1) they do not pass low frequency or DC signals and that (2) their cost and size remains relatively high because they cannot be integrated into ICs.

Optoisolators

Optoisolators provide a way of isolating two circuits by coupling via light transmission.

An optoisolator (or optocoupler) consists of an LED and a photodiode in the same package. The input of the optoisolator is the LED. The LED generates light when current flows through it. The light from the LED shines on a photodiode or phototransistor. The photodiode conducts when light from the LED strikes it. The photodiode is connected to the output circuit.



Since the connection between the LED and photodiode is by light alone, the input and output can be completely isolated electrically. The two sides do not even have to share a ground connection.

Although optocouplers can provide good isolation and are inexpensive, they have the disadvantage that they cannot supply power to the output. For many applications this makes an optocoupler interface more complex than one using a transformer.

Exercise 9: When the input to the optocoupler is high, will the output be high or low? Assume a pull-up is connected to the output.