

Lecture 6B: Pulse Shaping

John M Pauly


February 7, 2026

Digital Communications: Pulse Shaping


- We have a long sequence of bits we want to transmit
- What analog waveform should I use?
- There are lots of options!
- For now we'll look at on-off keying (OOK).
- We send a pulse for a "1", and nothing for a "0"

On-Off Keying

The sequence of bits might look like this:



One option would be to transmit a square pulses for a "1", and nothing for a zero. This would look like:



This works, but the square edges use a lot of bandwidth, and will interfere with adjacent signals in the spectrum.

What would be a better choice?


Sinc Pulses

From sampling and reconstruction, a sinc pulse might seem to be a good idea. It has a narrow, sharp spectral profile.

We can choose the sinc to be "1" at the sample where it is centered, and zero at all of the other samples. There is no intersymbol interference or ISI.

This is the first criteria Nyquist proposed for dealing with ISI

$$p(t) = \begin{cases} 1 & t = 0 \\ 0 & t = \pm kT_b, k = \pm 1, \pm 2, \dots \end{cases}$$




Pulse Shaping: Sinc Pulse

Let $f_b = 1/T_b$. The sinc pulse $\text{sinc}(f_b t)$ satisfies Nyquist's first criterion for zero ISI:

$$\text{sinc}(f_b t) = \begin{cases} 1 & t = 0 \\ 0 & t = \pm k T_b, \quad k = \pm 1, \pm 2, \dots \end{cases}$$

This pulse is bandlimited. Its Fourier transform is


$$P(f) = \frac{1}{f_b} \Pi\left(\frac{f}{f_b}\right)$$



Unfortunately, this pulse has infinite width and decays slowly, as well as other issues as we will see shortly.

Sinc Pulse Filtered Waveform

This is an example using a sinc pulse to generate the analog waveform




The waveform goes exactly through the samples, but there is a large overshoot.

Nyquist Pulse

Nyquist increased spectral width to make the pulse fall off more rapidly.

The Nyquist pulse has spectrum width $\frac{1}{2}(1 + \beta)f_b$, where $0 < \beta < 1$.



β is called the "roll off factor".

- $\beta = 0$ is a square spectrum, and a sinc $p(t)$
- $\beta = 1$ is a raised cosine

In general β is on the order of 0.25 to 0.5. This gives a somewhat broader spectrum, but not double as $\beta = 1$ would.

Nyquist Pulse Spectrum

We want to choose $P(f)$, and hence $p(t)$, so that

- $p(t) = 0$ for $t = nT_b$, so there is no intersymbol interference.
- $p(t) = 1$ for $t = 0$

So that it basically functions like a sinc. Can we do this?

If we sample the pulse $p(t)$ at rate $f_b = 1/T_b$, then

$$\bar{p}(t) = p(t) \text{III}_{T_b}(t) = p(t)\delta(t) = \delta(t).$$

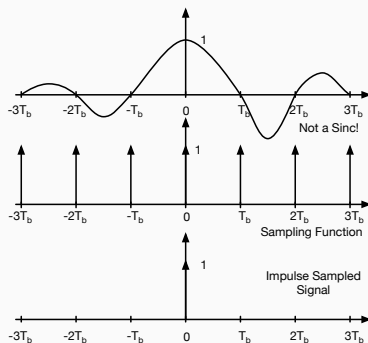
since $p(t)$ is zero at multiples of T_b , except at zero where it is 1.

Hence, if a pulse obeys these conditions, it will behave like a sinc once we sample!

This is a little hard to see!

First Nyquist Criterion

If we sample some non-sinc waveform that meets the criterion



No matter what the pulse looks like, if it is 1 at 0 , and zero at the other samples, it's sampled signal is an impulse at the origin.

It's spectrum is then flat.

First Nyquist Criterion (cont.)

The sinc does that because its spectrum is a rect.

However, we can get a flat spectrum without this!

What we want is that the Fourier transform of the sampled signal is

$$\overline{P}(f) = \sum_{k=-\infty}^{\infty} P(f - kf_b) = 1$$


The first term is the Fourier transform of the sampled signal $p(t) \text{III}_{T_b}(t)$, which equals the Fourier transform of $\delta(t)$, which is 1.

This is an interesting result! to see why this works, we need to look at how the spectrum of the pulses overlap.

Nyquist Pulse Spectrum (cont.)

Since we are sampling below the Nyquist rate $(1 + \beta)f_b$, the shifted transforms overlap.

Nyquist's criterion require the spectrum overlaps add to 1 for all f .



The result is the same as if we had used a sinc for $p(t)$, and summed them up. From the sampling perspective these do the same thing.

For parameter β with $0 < \beta < 1$, the resulting pulse has a half bandwidth


$$B_\beta = \frac{1}{2}(f_b + \beta f_b)$$

Nyquist Pulse Spectrum (cont.)

Nyquist chose a pulse with a “vestigial” raised cosine transform.

This transform is smoother than a trapezoid, so the pulse decays more rapidly.

The Nyquist pulse is parameterized by β . Let $f_x = \beta f_b/2$.



There are many other options that would work (linear, for example), but this is what is typically used.

Nyquist Pulse Spectrum (cont.)


Nyquist pulse spectrum is raised cosine pulse with flat porch.

$$P(f) = \begin{cases} 1 & |f| < \frac{1}{2}f_b - f_x \\ \frac{1}{2} \left(1 - \sin \pi \left(\frac{f - \frac{1}{2}f_b}{2f_x} \right) \right) & |f| - \frac{1}{2}f_b < f_x \\ 0 & |f| > \frac{1}{2}f_b + f_x \end{cases}$$


The transform $P(f)$ decays roughly as $1/t^3$, depending on the value of β .

Example Nyquist Pulse Spectra

Nyquist pulse spectra for roll-off factors β of 0, 0.25, 0.5, and 1.




Example Nyquist Pulses



Note that the pulse falls off dramatically as β increases. This means less computation for filtering on the transmit and receive side. But it costs bandwidth!


Overshoot

As β increases (and the bandwidth), the analog waveform is better behaved. Here are plots for $\beta = 0$, which is a sinc, and $\beta = 0.25$



Overshoot (cont.)


Here are plots for $\beta = 0.5$ and $\beta = 1.0$



These are much better behaved! Typically β is chosen to be somewhere between 0.25 and 0.5.

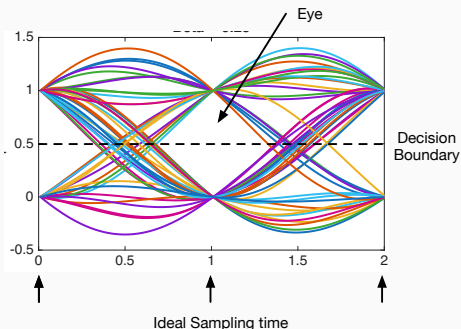
Decoding the Signals: Eye Diagram

At the receiver we have to find the symbols, and decide whether each is a "1" or a "0". The difficulty of deciding is captured in an eye diagram.



This is a plot of each of the symbol waveforms superimposed, here over a two symbol window.


Eye Diagram



The open area in the middle is called the "eye". The larger it is (the more open), the more reliable the decoding will be.

At each sampling time, we have to decide whether we've received a "1" or a "0". In this case we decide based on whether the sample is larger or smaller than $1/2$.

Eye Diagram (cont.)

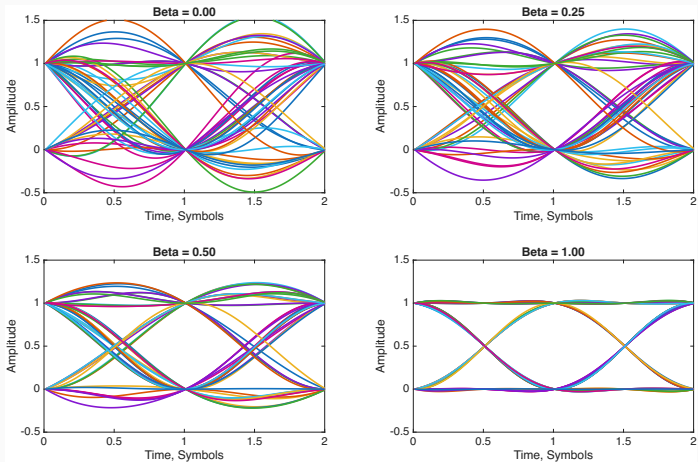


The height of the eye determines the sensitivity to noise and interference.

The width of the eye determine the sensitivity to timing and symbol synchronization.

Sampling too early or late will produce errors.

Eye Diagrams for Different Nyquist Pulse β 's



As β increases, the eye opens, and detection becomes easier. However, bandwidth also increases.

Square Root Nyquist Pulses

So far the filtering with the Nyquist pulse happens before transmission. We have to filter in order to generate the analog waveform.




We'll want to filter at the end also, to exclude other signals and optimally suppress noise (matched filter).

We can do this by using a "square root" Nyquist pulse at both ends of the system. This has a spectrum that is the square root of the normal pulse.




The two together give the same effect on the eye diagram as a single normal Nyquist pulse.

Square Root Nyquist Pulse for $\beta = 0.5$



The square root pulse is a little narrower, and hence wider bandwidth. It also is not zero at the other samples, so it will have ISI!


Square Root Nyquist Pulse Eye Diagram



The two pulses together give the same eye diagram as the single normal pulse.

Square Root Pulse Spectrum

The width of the square root pulse is wider than the normal pulse.



These are very commonly used. Most RF systems use square root Nyquist pulses. You'll need to include the square root filter to make the eye diagram look right.

In Matlab you can generate a normal Nyquist pulse with

```
>> p = rcosdesign(beta, symbol_span, samples_per_symbol, 'normal');
```

where `symbol_span` is the filter length in symbols, and `samples_per_symbol` is the oversampling factor. The total number of samples in the filter is the product of the two.


A square root Nyquist pulse is generated with

```
>> p = rcosdesign(beta, symbol_span, samples_per_symbol, 'sqrt');
```

This is in the signal processing toolbox.

Controlled ISI (Partial Response Signaling)

The second Nyquist approach embraces ISI. The interference is known if we know the adjacent bits! We can achieve a lower bandwidth by using an even wider pulse. Then ISI is canceled using knowledge of the pulse shape.



$p(t)$ is equal to 1 for $t = 0$ and $t = T_b$. The signal at time T_b depends on both $x(0)$ and $x(T_b)$.

- If both are 1, then the output $y(t) = 2$
- if one is 1, and the other is 0, the output is 1
- if both are 0, the output is 0.

With an initial a_0 we can subtract term by term, and recover $\{a_k\}$

Partial Response Signaling (cont.)

This is called a duobinary pulse. The ideal pulse is the sum two shifted sinc's


$$p(t) = \frac{\sin(\pi f_b t)}{\pi f_b t} + \frac{\sin(\pi f_b(t - T_b))}{\pi f_b(t - T_b)t} = \frac{\sin \pi f_b t}{\pi f_b t(1 - f_b t)}$$

This makes sense, it will be 1 at 0 and T_b and zero elsewhere.

The Fourier transform of $p(t)$ is


$$P(f) = \frac{2}{f_b} \cos\left(\frac{\pi f}{f_b}\right) \Pi\left(\frac{f}{f_b}\right) e^{-j\pi f/f_b}$$

The spectrum is confined to the theoretical minimum of $f_b/2$.



Duobinary Pulses

Suppose $p_a(t)$ satisfies Nyquist's first criterion (zero ISI).



Then


$$p_b(t) = p_a(t) + p_a(t - T_b)$$

is a duobinary pulse with controlled ISI.

We can use any of the Nyquist pulses from the previous slides as a basis for a duobinary pulse.

Sinc Based Duobinary Pulses

For a sinc pulse, the duobinary pulse looks like



By shift theorem,

$$P_b(f) = P_a(f)(1 + e^{-j2\pi T_b f})$$

Since $P_b(f_b/2) = 0$, most (or all) of the pulse energy is below $f_b/2$.


One issue with a duobinary pulse is that it has a non-zero DC value.

Modified Duobinary Pulses


We can eliminate unwanted DC component using modified duobinary, where $p_c(-T_b) = 1$, $p_c(T_b) = -1$, and $p_c(nT_B) = 0$ for other integers n .

$$p_c(t) = p_a(t + T_b) - p_a(t - T_b) \implies P_c(f) = 2jP_a(f) \sin 2\pi T_b f$$

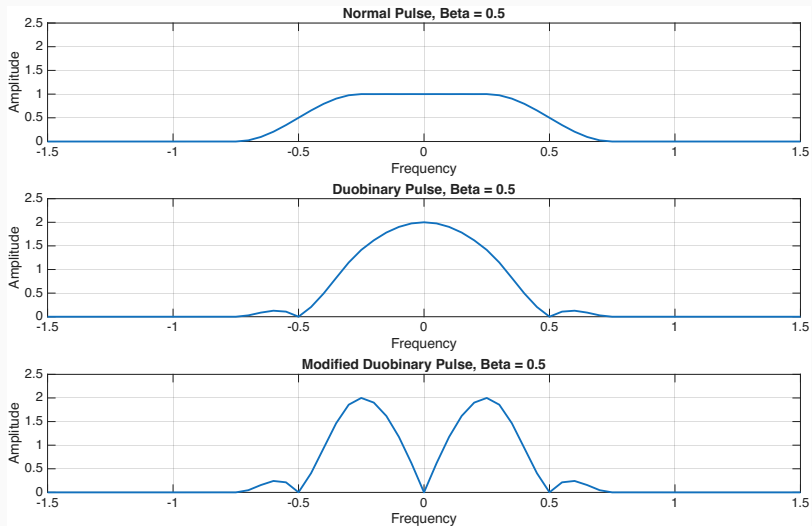
The transform of $p_c(t)$ has nulls at 0 and $\pm f_b/2$.



Example Duobinary and Modified Duobinary Pulses



Duobinary and Modified Duobinary Pulses Spectra



Partial Response Signaling Detection, Duobinary Pulses

Suppose that sequence 0010110 is transmitted (first bit is startup digit).

Digit x_k	0	0	1	0	1	1	0
Combined amplitude	0	1	1	1	2	1	
Decode sequence	0	0	1	0	1	1	0

For the combined amplitude, recall that for the previous two inputs,

- If both are 1, then the output is 2
- if one is 1, and the other is 0, the output is 1
- if both are 0, the output is 0.

Note that the signal now has multiple levels (0, 1, and 2). We could also do that directly, as we will see next week.

Partial Response Signaling Detection, Duobinary Pulses

Partial response signaling is susceptible to *error propagation*.

If a nonzero value is misdetectd, zeros will be misdetectd until the next nonzero value. The predicted next value will be the opposite of the current value.

This stops when we get a 0 or 2, which can only occur one way.

The need for partial response detection can be eliminated by differential encoding

- The input stream is preprocessed by subtracting adjacent samples,
- These sample differences are transmitted using duobinary pulses
- The receiver can directly read out the data bits.

We will see return to the shortly.

Pulse Shaping: Conclusion

- To transmit bits, we need to generate an analog waveform
- We would like to minimize intersymbol symbol interference (ISI)
- A sinc pulse would do this, but requires more computation, and doesn't produce a nice eye diagram
- Nyquist pulses trade off bandwidth, computation, and eye diagram characteristics
- These can be extended to duobinary pulses to further reduce bandwidth, and improve performance