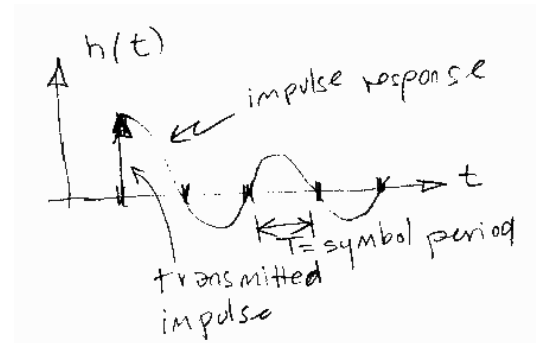


Data Transmission over Bandlimited Channels

After this lecture you should be able to: determine if a channel meets the Nyquist no-ISI criteria and, if so, the maximum signalling rate without ISI; determine the maximum error-free information rate over an AWGN channel; determine the specific conditions under which these two limits apply; explain the principles of operation for partial-response signalling, decision-feedback equalization and sequence estimation. You should be able to perform computations involving the OFDM symbol rate, sampling rate, block size and guard interval.

Introduction

All practical channels are band-limited (either low-pass or band-pass) and the channel bandwidth is often what limits the maximum data rate. We will study two theorems, the Nyquist no-ISI criteria and Shannon's capacity theorem, that provide some guidance about maximum data rate that can be achieved over a bandlimited channel.



One impulse response that meet this criteria is the sinc() function:

$$h(t) = \frac{\sin(\pi t/T)}{\pi t/T}$$

which has value 1 at $t = 0$ and 0 at multiples of T .

Inter-Symbol Interference

Waveforms transmitted over a bandwidth-limited channel tend to get "smeared" in time. The channel extends the duration of each transmitted symbol in time and into subsequently-transmitted symbols. This means there is a possibility that each symbol will interfere with subsequent symbols. This interference is called inter-symbol interference (ISI).

Nyquist no-ISI Criteria in Frequency

It is possible to derive the corresponding characteristics of the channel's frequency-domain transfer function that result in no ISI. This condition is that the channel frequency response have odd symmetry around half of the symbol frequency:


$$H\left(\frac{1}{2T} + f\right) + H\left(\frac{1}{2T} - f\right) = 1 \text{ for } 0 < |f| < \frac{1}{2T}$$

Nyquist no-ISI Criteria in Time


Consider a system that transmits symbols as (infinitely-)short pulses ("impulses"). A low-pass channel will limit the rise time of the pulses and cause the impulses to be smeared out in time. However, if the response to the impulses is zero at multiples of the symbol period the impulses will not interfere with each other. This is called the Nyquist no-ISI criteria.

Exercise 1: What is the transfer function of a channel with infinite bandwidth? Does this channel meet the Nyquist no-ISI condition?

Just as there could be many impulses responses that are zero at multiples of the symbol period, there are many no-ISI transfer functions. For example, the following two straight-line transfer functions meet the no-ISI condition:



The dashed line is a “brick-wall” filter whose response is 1 below half of the symbol rate and zero above that. Although such a filter would have the minimum overall bandwidth, it is not physically realizable and has other problems as described below. The filter with the linear transfer function is also difficult to implement. A more practical transfer function is the so-called raised-cosine function which is a half-cycle of a cosine function centered around half of the symbol rate.



Note that it is symmetry about $1/2T$ that ensures there will be no ISI rather than the exact filter shape. Thus we are free to implement other transfer functions, possibly arbitrary ones, if they make the implementation easier.



Pulse-Shaping Filter

Note that the no-ISI criteria applies for a channel that produces no ISI for *impulses*, not the square pulses typically used. Since practical systems don't transmit impulses, the Nyquist criteria cannot be applied directly to the physical channel itself. Instead, we consider that the transmitter includes a hypothetical filter that converts impulses to pulses before transmitting them over the channel. The response of this (im)pulse-shaping filter has to be included when evaluating the channel ISI. It is the combination of this impulse-shaping filter and the channel that has to meet the Nyquist criteria.

Excess Bandwidth

The width of the transfer function in the transition between passband and stopband has an impact on the shape of the impulse response and on the sensitivity of the receiver to errors in the timing of the sampling point.

This parameter, α , is called the “excess bandwidth”. The following diagram¹ shows how the excess bandwidth parameter for a raised-cosine transfer function affects the impulse response.



Larger values of excess bandwidth result in smaller values of the impulse response which in turn reduces the amount of ISI near the sampling point. This makes the receiver less sensitive to errors in the timing of the sampling point.

Exercise 2: What is the possible range of values of α ?

Equalization

To avoid ISI, the total channel response including the pulse-shaping filters, transmit filters, the channel and

¹From Simon Haykin, “An Introduction to Analog and Digital Communication”, 1989.

the receiver filter(s) have to meet the Nyquist no-ISI condition.

When the channel by itself is unlikely to meet the no-ISI conditions, the transmitter and receiver typically use filters, known as “equalizers”, to ensure the no-ISI condition is met.

Transmitter and receiver filters typically have other functions beside equalization. For example, the transmit filter may limit the bandwidth of the transmitted signal to limit interference to adjacent channels. The receiver filter may filter out noise and interference from adjacent channels and thus improve the SIR and SNR. The communication system designer would design the transmitter and receiver filters to meet both the filtering and equalization requirements.

A common situation is a flat channel where interference is not an issue. In this case a reasonable approach is to put half of the filtering at the transmitter and half at the receiver. In order to achieve an overall raised cosine transfer function, each side has to use a “root raised cosine” (RRC) transfer function. The product of the two filters is thus the desired raised-cosine function which meets the no-ISI condition.

Exercise 3: Could equalization be done at the transmitter only? Why or why not?

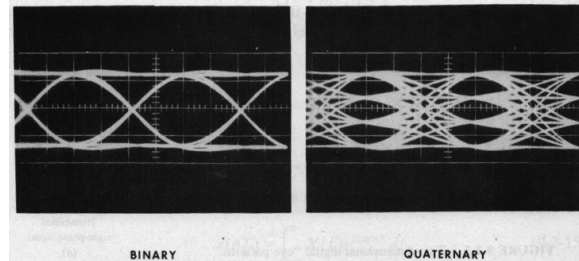
Adaptive Equalizers

In many communication systems the transfer function of the channel cannot be predicted ahead of time. One example is a modem used over the public switched telephone network (PSTN). Each phone call will result in a channel that includes different “loops” and thus different frequency responses. Another example is multipath propagation in wireless networks. The channel impulse response changes as the receiver, transmitter or objects in the environment move around.

To compensate for the time-varying channel impulse response the receiver can be designed to adjust the receiver equalizer filter response using various algorithms.


Eye Diagrams

An eye diagram is superimposed plots of duration T (the symbol period) of the received waveform (for random data). The eye diagram graphically demonstrates the effect of ISI. Some examples of eye diagrams produced by an appropriately-triggered oscilloscope²:



The vertical opening at the sampling time, called the “eye opening”, represents the amount of ISI at the sampling point.

The horizontal opening indicates how sensitive the receiver would be to errors in sampling point timing³:



Alternative Equalization Methods

Partial-Response Signalling

Instead of completely eliminating ISI, we can equalize the channel so that the effect of ISI is zero only for delays greater than one symbol period. For example, assume the inputs to the channel are 1 and -1. If the impulse response of the channel after one symbol period is 0.5 then the output of the channel will go up or down by 0.5. The output at symbol periods could thus be -1, 0 or +1 (assuming limiting to values of ± 1). Based on the change in channel output we

²From John G Proakis, “Digital Communications”, 3rd Ed., 1983.

³Proakis, op. cit.

can decide if the transmitted value was +1 or -1. This allows the transmitter to recover the transmitted data even though the received signal is corrupted by ISI.


Since the equalized channel has partial response to the input, this technique is called “Partial Response Signalling.”

Exercise 4: Draw the impulse response of a partial-response channel. Label the time axis.

Decision-Feedback Equalization

Another approach to equalizing the channel that is often simpler to implement is to predict the ISI that would result from previously-transmitted bits and subtract that ISI from the symbol currently being received.

The feed-back equalizer looks at the decoded bits that have already been received and subtracts out the ISI they would have caused to the symbol currently being received:



Decision-feedback equalizers suffer from a phenomenon called error propagation. If an error is made in decoding one symbol, an incorrect amount of ISI will be predicted and the result will be incomplete ISI cancellation and, potentially, additional error(s) being caused.

Sequence Estimation

Instead of trying to correct the ISI, the receiver can predict the ISI-corrupted waveforms that would result from all the possible sequences of transmitted symbols. The receiver would then pick the sequence that is most similar to the waveform that was actually received.

Exercise 5: How many possible transmitted waveforms could be received if a sequence of 10 symbols is transmitted, each symbol being chosen from 4 possible symbols?

Although this may seem like an impractical approach, various methods can be used to avoid the exponential increase in receiver complexity with mes-


sage length. The most common is to consider only the effect of a few of the most recent symbols. Another method is to use an algorithm, the Viterbi algorithm, that can significantly simplify the search for the “best match” (or “Maximum Likelihood”) waveform. Sequence estimation receivers are thus often said to use “Viterbi equalizers.”

OFDM

And alternative to adaptive equalization is to use a technique called Orthogonal Frequency Division Multiplexing (OFDM). OFDM transmits many (typically 64 or more) symbols in parallel at different “sub-carrier” frequencies. This reduces the symbol period by an amount equal to the number of parallel streams without reducing the overall bit rate.

An OFDM system converts blocks of N samples to/from the frequency domain thus reducing the symbol rate by a factor N . The frequency-domain signal can have up to N subcarriers. The value of N is typically a power of 2 for for efficient implementation using Fast Fourier Transforms (FFTs).

OFDM systems typically insert a “guard time” between symbols that is longer than the duration of the impulse response of the channel. This minimizes interference between symbols. The last samples of each block of N samples are prepended to the symbol and transmitted during the guard time to create a “cyclic” or “periodic” extension.



OFDM has become more popular than adaptive equalization because it is simpler to implement and more robust. OFDM is used by most ADSL, WLAN and 4G cellular standards.

Shannon Capacity

The Shannon Capacity of a channel is the information rate above which it is not possible to transmit data

with an arbitrarily low error rate. For the Additive White Gaussian Noise (AWGN) channel the capacity is:

$$C = B \log_2 \left(1 + \frac{S}{N} \right)$$

where C is the capacity (b/s), B is the bandwidth (Hz) and $\frac{S}{N}$ is the signal to noise ratio.

The Shannon limit does not say that you can't transmit data faster than this limit, only that if you do, you can't reduce the error rate to an arbitrarily low value.

Shannon's work also does not specify how to achieve capacity, for example, what modulation and coding should be used. However, Shannon's work does hint that using error-correcting codes should allow us to achieve arbitrarily-low error rates as long as we limit the data (actually, information) rate to less than the channel capacity.

Exercise 6: What is the channel capacity of a 3 kHz channel with an SNR of 20dB?

To implement systems that operate at close to channel capacity requires coding. Some systems using modern codes such as "Turbo" codes or Low Density Parity Check (LDPC) codes can operate within a fraction of a dB of channel capacity.

Note that the symbol rate limitations resulting from ISI do not by themselves limit the achievable bit rate or the capacity of the channel. For example, once the ISI is controlled we can use an arbitrarily large symbol set (any value of M) and transmit any number of bits per symbol.

The limitation on symbol rate is also different than the constraint on channel capacity defined by the Shannon bound which takes into account the signal to noise ratio as well as the bandwidth.

Also note that the Shannon capacity refers to *information* rate (the bit rate after maximum possible compression), not to the bit rate.