1 Signal Space and Channel Models

1.1 The Signal Space

For a continuous-time channel $Y(t) = X(t) + N(t)$, consider the vector-space of *finite-energy* signals. Let $\mathcal{L}_2$ be the set of complex-valued signals (functions) $x(t)$ with finite energy:

$$\int_{-\infty}^{\infty} |x(t)|^2 dt < \infty$$

$\mathcal{L}_2$ is a *vector space* which is a set of elements (called *vectors*) that is closed under addition and scalar multiplication.

The inner product of $x(.), y(.) \in \mathcal{L}_2$ can be defined as follows:

$$\langle x, y\rangle = \int_{-\infty}^{\infty} x(t) y^*(t) dt$$

The norm of a signal is the square-root of its energy:

$$\|x\| = \sqrt{\langle x, x\rangle} = \left[\int_{-\infty}^{\infty} |x(t)|^2 dt\right]^{1/2}$$

1.2 Orthonormal Basis

For any vector space $\mathcal{L} \subset \mathcal{L}_2$, the set of functions $\{f_i(.), i = 1, 2, \ldots\}$ is called an orthonormal basis for $\mathcal{L}$ if:

1. Every $x(.) \in \mathcal{L}$ can be expressed as:

   $$x(t) = \sum_i x_i f_i(t)$$

   The coefficients $x_1, x_2, \ldots$ are called the *projection coefficients*, and $x_1 f_1(t), x_2 f_2(t) \ldots$ are the projections of the signal $x(t)$ along $f_1(t), f_2(t), \ldots$, respectively.

2. The functions $\{f_i(.), i = 1, 2, \ldots\}$ are orthonormal:

   $$\langle f_\ell, f_m\rangle = \begin{cases} 1 & \text{if } \ell = m \\ 0 & \text{if } \ell \neq m \end{cases}$$

The orthonormal basis f_i *span* the vector space $\mathcal{L}$ and the number of elements (functions) in the basis is called the *dimension* of $\mathcal{L}$.

The inner product between $x(t)$ and $y(t)$ is:

$$\langle x(t), y(t)\rangle = \sum_i x_i y_i^*$$

The energy of $x(t)$ can therefore be written as:

$$\int |x(t)|^2 dt = \langle x(t), x(t)\rangle = \sum_i |x_i|^2$$

Gram-Schmidt procedure: Given functions $\{x_1(t), x_2(t), \ldots, x_m(t)\}$, we find an orthonormal basis $\{f_1(t), f_2(t) \ldots\}$ as follows:

1. Let $f_1(t) = \frac{x_1(t)}{\|x_1(t)\|}$.
2. Find the part of x_2 orthogonal to f_1, and normalise. Let

   $$g_2(t) = x_2(t) - \langle x_2, f_1\rangle f_1(t)$$

   Then, $f_2(t) = \frac{g_2(t)}{\|g_2(t)\|}$.
3. Find the part of x_3 orthogonal to f_1, f_2, and normalise. Let

   $$g_3(t) = x_3(t) - \langle x_3, f_1\rangle f_1(t) - \langle x_3, f_2\rangle f_2(t)$$

   Then, $f_3(t) = \frac{g_3(t)}{\|g_3(t)\|}$.

If the dimension of the space is k, only $f_1, \ldots f_k$ will be non-zero.

1.3 Modelling a Channel

Channels are often modelled as *linear time-invariant* systems with additive noise $y(t) = h(t) \star x(t) + n(t)$. In frequency domain:

$$Y(f) = H(f)X(f) + N(f)$$

Suppose the input signal $x(t)$ is band1imited to $[-W_0, W_0]$, where $|H(f)|$ is constant, then we can compensate for the constant channel gain and the constant delay at the receiver. The channel is effectively:

$$y(t) = x(t) + n(t)$$

Passband channel models: The signal is restricted to have frequency components in the band $[f_c + W, f_c - W]$, where f_c is a carrier frequency (typically $f_c \gg W$). If $|H(f)|$ is constant throughout the band and delay is also a constant τ, then an additive noise model $y(t) = x(t) + n(t)$ can be used.

Noise model:

$$y(t) = h(t) \star x(t) + n(t)$$

We model $n(t)$ as a Gaussian noise process and the effective channel is called an *additive Gaussian noise* channel.

1.4 Modelling the Noise

Gaussian white noise process: For each t, $n(t)$ is Gaussian with zero mean and autocorrelation function:

$$\mathbb{E}[n(t)n(t+\tau)] = \frac{N_0}{2}\delta(\tau)$$

The power spectral density (PSD) is:

$$S_n(f) = \frac{N_0}{2}$$

In practice, transmitted signal $x(t)$ is bandlimited. $n(t)$ has PSD $S_n(f) = \frac{N_0}{2}$ for all f where the signal has non-zero spectrum.

1.5 Signal Detection

Suppose that $\{\phi_1(t), \ldots, \phi_K(t)\}$ is an orthonormal basis for the signal set consisting of M waveforms $\{s_1(t), \ldots, s_M(t)\}$. (Note that $K \leq M$)

Hence, each signal $s_i(t)$ can be expressed as:

$$s_i(t) = s_{i,1}\phi_1(t) + \ldots + s_{i,K}\phi_K(t)$$

The projection coefficients for $s_i(t)$ are:

$$s_{i,j} = \int s_i(t)\phi_j^*(t) dt$$

Thus each signal is equivalent to a K - dimensional vector. For $i = 1, \ldots, M$:

$$s_i(t) \leftrightarrow \underline{s}_i = \left[s_{i,1}, \ldots, s_{i,K}\right]$$

The transimitted signal $x(t)$ lies in the space spanned by $\{\phi_1(t), \ldots, \phi_K(t)\}$. Projecting $y(t)$ onto the space spanned by $\{\phi_1(t), \ldots, \phi_K(t))\}$, we obtain the coefficient vector $\underline{r} = [r_1, \ldots, r_K]$:

$$r_j = \left\langle x(t), \phi_j(t)\right\rangle + \left\langle n(t), \phi_j(t)\right\rangle$$

We write $\underline{r} = \underline{x} + \underline{n}$, where $\underline{x} = [x_1, \ldots, x_K]$ and $\underline{n} = [n_1, \ldots, n_K]$ are the projection coefficient vectors of the signal $x(t)$ and noise $n(t)$, respectively.

Let $\{\phi_m(t)\}_{m\in\mathbb{Z}}$ be any orthonormal set of functions, and $n(t)$ be a Gaussian noise white noise process with zero mean and spectral density $N_0/2$. For $m \in \mathbb{Z}$:

$$n_m = \int_{-\infty}^{\infty} n(t)\phi_m(t) dt$$

Then $\{n_m\}_{m\in\mathbb{Z}}$ are i.i.d. Gaussian with zero mean and variance $\frac{N_0}{2}$.

1.6 Optimal Detection

If $\hat{x}$ represents the vector decoded by the receiver, we wish to minimise the probability of detection error $P(\hat{\underline{x}} \neq \underline{x})$.

Given transmitted vector $\underline{x} \in \mathcal{S}$, suppose $\underline{r}$ is generated according to the conditional distribution $P(\underline{r} \mid \underline{x})$. The optimal detection rule that minimizes the probability of detection error is the *Maximum a posteriori probability* (MAP) rule:

$$\begin{aligned}\hat{\underline{x}} &= \arg\max_{\underline{s}_i\in\mathcal{S}} P\left(\underline{x} = \underline{s}_i \mid \underline{r}\right) \\ &= \arg\max_{\underline{s}_i\in\mathcal{S}} P\left(\underline{x} = \underline{s}_i\right) f\left(\underline{r} \mid \underline{x} = \underline{s}_i\right)\end{aligned}$$

If prior distribution on the signal vectors is uniform, the MAP rule becomes the *maximum-likelihood* (ML) decoding rule:

$$\hat{X} = \arg\max_{\underline{s}_i\in\mathcal{S}} f\left(\underline{r} \mid \underline{x} = \underline{s}_i\right)$$

For additive white Gaussian noise (AWGN) $n(t)$:

$$f\left(\underline{r} \mid \underline{x} = \underline{s}_i\right) = \frac{1}{(\pi N_0)^{K/2}} e^{-\|\underline{r} - \underline{s}_i\|^2/N_0}$$

If the prior distribution over the signal vectors is uniform, the optimal detection rule is *minimum distance* decoding:

$$\hat{\underline{x}}^{\text{ML}} = \arg\min_{\underline{s}_i\in\mathcal{S}} \left\|\underline{r} - \underline{s}_i\right\|^2$$

Suppose the prior probabilities of the signal vectors are $P\left(\underline{x} = \underline{s}_i\right) = p_i$:

$$\hat{\underline{x}}^{\text{MAP}} = \arg\min_{\underline{s}_i\in\mathcal{S}} \ln\frac{1}{p_i} + \frac{\left\|\underline{r} - \underline{s}_i\right\|^2}{N_0}$$

Probability of detection error:

$$P_e = \sum_{\underline{s}_i\in\mathcal{S}} P\left(\underline{x} = \underline{s}_i\right) P\left(\hat{\underline{x}} \neq \underline{s}_i \mid \underline{x} = \underline{s}_i\right)$$

2 Baseband Transmission

2.1 Pulse Amplitude Modulation

The most common modulation scheme for a baseband channel with additive Gaussian noise is *Pulse Amplitude Modulation* (PAM).

The set of values the bits are mapped to is called the *constellation* $\mathcal{C}$. In a constellation with M symbols, each symbol represents $\log_2 M$ bits.

The *pulse* waveform is a unit-energy baseband waveform denoted $p(t)$. A sequence of constellation symbols $X_0, X_1, X_2, \ldots$ is used to generate a baseband signal as follows:

$$x(t) = \sum_m X_m p(t - mT)$$

T is called the symbol time of the pulse.

The transmission rate is $\frac{1}{T}$ symbols / sec or $\frac{\log_2 M}{T}$ bits/second.

Time Decay vs. Bandwidth Trade-off: We want $p(t)$ to decay quickly in time but also be approximately band-limited.

Orthonormality: We choose $p(t)$ so that the shifted pulses $\{p(t - mT)\}_{m\in\mathbb{Z}}$ form an *orthonormal basis*.

Denoting $\phi_m(t) = p(t - mT)$, the PAM signal is:

$$x(t) = \sum_m X_m \phi_m(t)$$

Projecting $y(t)$ onto the space spanned by $\{\phi_1(t), \phi_2(t)), \ldots\}$, we obtain the coefficients, we have $Y_k = X_k + N_k$, for $k \in \mathbb{Z}$, where n_k is the projection coefficient of the noise with the basis function ϕ_m:

$$N_k = \langle n(t), \phi_k(t)\rangle = \langle n(t), p(t - kT)\rangle$$

2.2 Matched Filter

Let the filter impulse response be $q(t) = p(-t)$. Since $y(t) = \sum_m X_m p(t - mT) + n(t)$, the filter output is $y(t) \star q(t)$:

$$r(t) = \sum_m X_m \int_{-\infty}^{\infty} p(\tau - mT)p(\tau - t) d\tau + \int_{-\infty}^{\infty} n(\tau)p(\tau - t) d\tau$$

Sampling at $t = kT$, we get $r(kT) = X_k + N_k$.

2.3 Optimal Detection

Optimal (MAP) detection rule to recover X from $Y = X + N$ when $Y = y$:

$$\hat{X}(y) = \arg\max_{c\in\mathcal{C}} P(X = c) f(y \mid X = c)$$

If all the symbols in the constellation are equally likely, i.e. $P(X = c)$ is the same for all symbols $c \in \mathcal{C}$, then the MAP rule becomes:

$$\hat{X} = \arg\max_{c\in\mathcal{C}} f(y \mid X = c)$$

This is the *maximum-likelihood* (ML) decoding rule.

For an M -point constellation with equally likely symbols, the *average energy per symbol* $E_s = E_b \log_2 M$.

The *signal-to-noise ratio* $\frac{E_b}{N_0}$ is a key parameter of a transmission scheme. P_e is often plotted as a function of $\frac{E_b}{N_0}$.

2.4 The Nyquist Pulse Criterion

When there is no noise, the filter output is:

$$r(t) = \sum_k X_k \int_{-\infty}^{\infty} q(u)p(t-kT-u)du = \sum_k X_k g(t-kT)$$

$p(t)$ is sometimes called transmit filter, $q(t)$ the receive filter, and $g(t)$ the overall filter:

$$g(t) = q(t) \star p(t) = \int_{-\infty}^{\infty} q(u)p(t-u)du$$

If we want $r(mT) = X_m$ for all integers m, then $g(t)$ should satisfy:

$$g(mT) = \begin{cases} 1, & m = 0 \\ 0, & m = \ldots,-2,-1,1,2,1,\ldots \end{cases}$$

Otherwise, we have *inter-symbol interference* (ISI).

Nyquist Pulse Criterion: Let $G(f)$ denote the Fourier transform of the effective pulse $g(t)$. Then the time-domain condition for no ISI is equivalent to:

$$\sum_{n=-\infty}^{\infty} G\left(f - \frac{n}{T}\right) = T$$

The Nyquist pulse criterion implies that in order to have no ISI, $G(f)$ must have bandwidth at least $1/(2T)$.

If the pulse bandwidth B lies in $\left(\frac{1}{2T}, \frac{1}{T}\right]$ and $G(f)$ is real and even, then for all $\Delta \in \left[0, \frac{1}{2T}\right)$ we need:

$$G\left(\frac{1}{2T} - \Delta\right) + G\left(\frac{1}{2T} + \Delta\right) = T$$

This condition is called *band-edge symmetry*.

With the matched filter $q(t) = p(-t)$, the orthonormal shifts property of $p(t)$ (in time domain) is equivalent to $G(f) = |P(f)|^2$ satisfying the Nyquist pulse criterion.

2.5 Power Spectral Density

The *power spectral density* (PSD) of the PAM signal describes the average power in any frequency band. To calculate the PSD, we will consider a slightly modified PAM signal:

$$x(t) = \sum_{k=-\infty}^{\infty} X_k p(t-kT-\Theta)$$

The random dither (delay) Θ make $x(t)$ a wide-sense stationary (WSS) process.

We assume that the X_k 's are drawn from a constellation with zero mean and the random process $\{X_k\}_{k=-\infty}^{k=\infty}$ is a WSS discrete-time process.

The *autocovariance function* of $x(t)$ is:

$$R_x(t+\tau, t) = \frac{1}{T} \sum_{m=-\infty}^{\infty} R_x[m] R_p(\tau - mT)$$

$R_p(\tau)$ is defined as $\int_{-\infty}^{\infty} p(v)p(v-\tau)dv$. For real $p(t)$:

$$\mathcal{F}\left[R_p(\tau)\right] = P(f)P(-f) = |P(f)|^2$$

Therefore the PSD of the transmitted PAM signal $x(t)$ is:

$$S_x(f) = \frac{|P(f)|^2}{T} \sum_{m=-\infty}^{\infty} R_X[m] e^{-j2\pi m f T}$$

When the symbols $\{X_k\}$ are independent:

$$R_X[m] = \begin{cases} \mathbb{E}\left[X_k^2\right] = \mathcal{E}_s, & m = 0 \\ 0, & m \neq 0 \end{cases}$$

$\mathcal{E}_s$ denotes the average energy per constellation symbol. Then the formulas for autocovariance and PSD of $x(t)$ simplify to:

$$R_x(\tau) = \frac{\mathcal{E}_s}{T} R_p(\tau), \quad S_x(f) = \frac{\mathcal{E}_s}{T} |P(f)|^2$$

Parseval's Theorem: The average power of the PAM waveform is then calculated as:

$$\frac{\mathcal{E}_s}{T} \int_{-\infty}^{\infty} |P(f)|^2 df = \frac{\mathcal{E}_s}{T} \int_{-\infty}^{\infty} |p(t)|^2 dt$$

If the pulse $p(t)$ has unit energy, then power of PAM signal is $\mathcal{E}_s/T$. If the matched receive filter is chosen as $q(t) = p(-t)$, then the overall filter frequency response is:

$$G(f) = P(f)P(-f) = |P(f)|^2$$

The PSD can be written as:

$$S_x(f) = \frac{\mathcal{E}_s}{T} G(f)$$

3 Passband Modulation

3.1 Upconverted PAM

Up-convert a PAM signal to passband by multiplying with a carrier:

$$x(t) = x_b(t)\cos(2\pi f_c t) = \left[\sum_k X_k p(t-kT)\right] \cos(2\pi f_c t)$$

The passband signal $x(t)$ has spectrum:

$$X(f) = \frac{1}{2}\left[X_b(f-f_c) + X_b(f+f_c)\right]$$

Due to $X_b(-f) = X_b^*(f)$, the lower sideband in $X(f)$ will be completely determined by the upper one (and vice versa). If the bandwidth of the PAM signal is W, then the passband signal has bandwidth $2W$.

3.2 Quadrature Amplitude Modulation

For *Quadrature Amplitude Modulation* (QAM), the constellation from which the symbols X_k are drawn can be *complex-valued*.

The QAM signal is generated as $x(t) = \mathrm{Re}\left[\sqrt{2} x_b(t) e^{j2\pi f_c t}\right]$. Define the following functions for $k \in \mathbb{Z}$:

$$f_k^r(t) = p(t-kT)\sqrt{2}\cos(2\pi f_c t)$$
$$f_k^i(t) = -p(t-kT)\sqrt{2}\sin(2\pi f_c t)$$

The transmitted QAM waveform can be expressed as:

$$x(t) = \sum_k \left[X_k^r f_k^r(t) + X_k^i f_k^i(t)\right]$$

The set of functions $\left\{f_k^r(t), f_k^i(t)\right\}, k \in \mathbb{Z}$ is an orthonormal set.

3.3 Demodulation

At the receiver, we have:

$$y(t) = \sum_k \left[X_k^r f_k^r(t) + X_k^i f_k^i(t)\right] + n(t)$$

As $p(t)$ is a baseband pulse bandlimited to $[-W, W]$, we can reject the high-frequency components using low-pass filters. At the output of the low-pass filter after carrier multiplication, we get:

$$y^r(t) = \sum_k X_k^r p(t-kT) + n^r(t)$$
$$y^i(t) = \sum_k X_k^i p(t-kT) + n^i(t)$$

If we choose receive filter $q(t) = p(-t)$ so that the overall filter $g(t) = p(t) \star p(-t)$ satisfies Nyquist pulse criterion, then:

$$Y_m^r = X_m^r + N_m^r$$
$$Y_m^i = X_m^i + N_m^i$$

N_m^r, N_m^i are samples of $n^r(t) \star q(t)$ and $n^i(t) \star q(t)$ at time mT.

Since $\left\{f_1^r(t), f_2^r(t), \ldots, f_1^i(t), f_2^i(t), \ldots\right\}$ form an orthonormal basis, the random variables $N_1^r, N_2^r, \ldots, N_1^i, N_2^i, \ldots$ are i.i.d. $\sim \mathcal{N}\left(0, \frac{N_0}{2}\right)$.

3.4 Optimal Detection

The optimal rule for detecting a QAM symbol X from $Y = X + N$ is the MAP rule:

$$\hat{X} = \arg\max_{x \in \mathcal{C}} P(X = x) \cdot f(Y = y \mid X = x)$$

If the constellation symbols are equally likely MAP rule reduces to maximum-likelihood (ML):

$$\hat{X}_{ML} = \arg\max_x f(Y = y \mid X = x) = \arg\max_{x=(x^r, x^i)} \frac{1}{\pi N_0} e^{-\frac{(y^r - x^r)^2 + (y^i - x^i)^2}{N_0}}$$

For i.i.d. Gaussian noise, ML is therefore equivalent to minimum-distance decoding:

$$\hat{X} = \arg\min_x \|y - x\|^2 = \arg\min_x \|x\|^2 - 2x^T y$$

The ML decoding rule for Phase Shift Keying (PSK) is $\hat{X} = \arg\min_x \theta(x, y)$.

3.5 Frequency Shift Keying

In *Frequency Shift Keying* (FSK), the information modulates the frequency of the carrier. To transmit message $i \in \{1, \ldots, M\}$ in any symbol period $[\ell T, (\ell+1)T)$, the basis function f_i is:

$$\sqrt{\frac{2}{T}} \cos\left(2\pi\left(f_c + (2i - (M+1))\frac{\Delta_f}{2}\right)t\right)$$

$\Delta_f = \frac{1}{2T} = \frac{f_c}{K}$ for some large integer K.

The set $\{f_1(t), \ldots, f_M(t)\}$ forms an orthonormal basis. The M symbols are represented by M frequencies, with adjacent frequencies separated by Δ_f.

We can express the FSK signal in terms of the basis functions as $x(t) = \sum_{k=1}^{M} x_k f_k(t)$.

For message i the projection coefficients are:

$$[x_1, \ldots, x_i, \ldots x_M] = \left[0, \ldots, \sqrt{E_s}, \ldots, 0\right]$$

The demodulator computes:

$$Y_i = \langle y(t), f_i(t) \rangle = X_i + N_i$$

The noise variables $N_j \sim \mathcal{N}\left(0, \frac{N_0}{2}\right)$ are i.i.d.. Hence $[Y_1, \ldots, Y_i, \ldots, Y_M] = \left[N_1, \ldots, \sqrt{E_s} + N_i, \ldots, N_M\right]$.

The optimal detection rule for M-ary FSK is $\hat{m} = \arg\max_{1 \leq i \leq M} y_i$.

The total bandwidth required for M-ary FSK is $(M-1)\Delta_f = \frac{(M-1)}{2T}$. The *bandwidth efficiency* η is defined as rate/bandwidth:

$$\eta_{QAM} \approx \log_2 M \text{ bits /s/Hz}$$
$$\eta_{MFSK} = \frac{2\log_2 M}{M-1} \text{ bits /s/Hz}$$

4 Modulation Concepts

4.1 Channel Equalisation

A channel frequency response $H(f)$ that is not flat in the transmission band is called *frequency-selective* or *dispersive*. In general:

$$y(t) = \int h(u)x(t-u)du + n(t) = \sum_k X_k f(t-kT) + n(t)$$

$f(t) = p(t) \star h(t)$. At the receiver: In the absence of noise, $r(t) = \sum_k X_k g(t-kT)$, where the overall filter is:

$$g(t) = f(t) \star q(t) = p(t) \star h(t) \star q(t)$$

Defining $r_m = r(mT)$, and $g_m = g(mT)$:

$$r_m = \sum_{\ell} g_\ell X_{m-\ell}$$

Taking z-transforms on both sides, we obtain $R(z) = X(z)G(z)$.

4.2 The Zero-forcing Equaliser

This equalising filter $H_E(z) = 1/G(z)$ is called the *zero-forcing* filter as it completely eliminates the ISI. The output of the filter (in the noiseless case) has z-transform:

$$Y(z) = R(z)H_E(z) = X(z)$$

In the presence of noise, the output of the zero-forcing filter is:

$$y_m = x_m + \tilde{n}_m$$

The sequence $\{\tilde{n}_m\}$ is the inverse z - transform of $N(z)/G(z)$.

$$\begin{aligned} H_E(z) &= \frac{1}{\sum_{\ell=0}^{L} g_\ell z^{-\ell}} \\ &= h_0 + h_1 z^{-1} + h_2 z^{-2} + h_3 z^{-3} + \ldots \end{aligned}$$

The above zero forcing filter is IIR, i.e., impulse response $\underline{h}$ has infinitely many non-zero coefficients.

The IIR filter can be implemented using $\underline{r} = \underline{y} \star \underline{g}$:

$$y_n = \frac{1}{g_0}\left[r_n - g_1 y_{n-1} - \ldots - g_L y_{n-L}\right]$$

We want to design an equalising filter with $K+1$ taps, with impulse response $\underline{h} = [h_0, \ldots, h_K]$. If we pass $\underline{r}$ through the filter, the output is:

$$y_m = X_m f_0 + \sum_{j=1}^{L+K} X_{m-j} f_j + \sum_{i=0}^{K} h_i n_{m-i}$$

$\underline{f} = [f_0, \ldots, f_{L+K}] = \underline{h} \star \underline{g}$:

$$f_j = \sum_{i=0}^{K} h_i g_{j-i}$$

For an ideal zero-forcing equaliser, we would like $f_0 = 1$, $f_j = 0$ for $j = 1, \ldots, L+K$. Determine $\underline{h} = [h_0, \ldots, h_K]$ by solving:

$$\begin{aligned} f_0 &= h_0 g_0 = 1 \\ f_1 &= h_0 g_1 + g_0 h_1 = 0 \\ &\vdots \\ f_K &= h_0 g_K + h_1 g_{K-1} + \ldots + h_K g_0 = 0 \end{aligned}$$

The output of the FIR ZF equaliser with $K+1$ taps is:

$$y_m = X_m + \sum_{j=K+1}^{L+K} X_{m-j} f_j + \sum_{i=0}^{K} h_i n_{m-i}$$

Assuming the noise variables $\{n_m\}$ are i.i.d. $\mathcal{N}\left(0, \sigma^2\right)$ the output noise variance is:

$$\mathbb{E}\left(\sum_{i=0}^{K} h_i n_{m-i}\right)^2 = \sigma^2\left(\sum_{i=0}^{K} h_i^2\right)$$

Thus there is a trade-off between minimising residual interference and increasing the noise variance.

4.3 MMSE Equalisation

The minimum mean squared error (MMSE) equaliser explicitly tries to minimise the expected squared error between X_m and its estimate $\hat{X}_m$.

For a $(K+1)$ -tap MMSE equaliser, the estimate $\hat{X}_m$ is generated as:

$$\hat{X}_m = c_0 r_m + c_1 r_{m+1} + \ldots + c_K r_{m+K}$$

We want to choose $\underline{c} = [c_0, \ldots, c_K]$ in order to minimise:

$$\mathbb{E}\left[\left(X_m - \hat{X}_m\right)^2\right] = \mathbb{E}\left[\left(X_m - \underline{r}^T \underline{c}\right)^2\right]$$

The optimal $\underline{c} = [c_0, \ldots, c_k]$ satisfies:

$$\frac{\partial f(\underline{c})}{\partial c_0} = \mathbb{E}\left[r_m\left(X_m - \underline{r}^T \underline{c}\right)\right] = 0$$

$$\vdots$$

$$\frac{\partial f(\underline{c})}{\partial c_K} = \mathbb{E}\left[r_{m+K}\left(X_m - \underline{r}^T \underline{c}\right)\right] = 0$$

Therefore $\mathbb{E}[\underline{r} X_m] = \mathbb{E}\left[\underline{r}\underline{r}^T\right]\underline{c}$, the MMSE equaliser is: $\underline{c} = \mathbf{R}^{-1}\mathbf{p}$, where

$$\mathbf{R} = \mathbb{E}\left[\underline{r}\underline{r}^T\right], \quad \mathbf{p} = \mathbb{E}[\underline{r} X_m]$$

4.4 Orthogonal Frequency Division Multiplexing

The term-by-term multiplication of DFT corresponds to *circular* convolution in the discrete-time domain, but the matched filter output is a *linear* convolution.

OFDM transmitter:

1. Encode the information in DFT domain by choosing $X[0], \ldots, X[N-1]$ from a QAM constellation
2. Take inverse DFT to produce the time-domain sequence $x[0], \ldots, x[N-1]$.
3. Insert *cyclic prefix*:
$$\{x[-L], \ldots, x[-1]\} = \{x[N-L], \ldots, x[N-1]\}$$
4. Using time-domain sequence, form:
$$x_b(t) = \sum_{m=-L}^{N-1} x[m] p(t - mT_s)$$
5. Transmitted passband signal:
$$x(t) = \mathrm{Re}\left[x_b(t) e^{j2\pi f_c t}\right]$$

Each OFDM block has total duration $(L+N)T_s$ seconds. N QAM symbols are transmitted per block.

OFDM receiver: Received signal: $y(t) = x(t) \star h(t) + n(t)$.

1. Carrier multiplication + low-pass filter + matched filter
2. Sampled matched filter output at $t = kT_s$:
$$r[k] = \sum_{\ell=0}^{L} g[\ell] \times [k-\ell] + n[k]$$
The outputs $r[-L], \ldots, r[-1]$ are discarded.
3. Compute the N-point DFT of $r[0], \ldots, r[N-1]$ to get:
$$R[n] = G[n]X[n] + N[n]$$
4. Use MAP rule to recover the QAM symbols $X[n]$.

The period of duration LT_s to send the cyclic prefix is called the *guard interval* or *guard period*. The guard period occupies $L/(L+N)$ of the OFDM block.

The DFT coefficients $\{G[n]\}$ correspond to samples of the spectrum $G(f)$ at N equally spaced freqs between $\frac{-1}{2T_s}$ and $\frac{1}{2T_s}$.

$$G(f) = P(f)H_b(f)P(-f) = |P(f)|^2 H_b(f)$$

For $0 \leq n \leq (N-1)$, we have $G[n] = G(f_n) = |P(f_n)|^2 H_b(f_n)$:

$$f_n = \begin{cases} \frac{1}{T_s}\frac{n}{N}, & 0 \leq n < \frac{N}{2} \\ \frac{1}{T_s}\left(\frac{n-N}{N}\right), & \frac{N}{2} \leq n \leq (N-1) \end{cases}$$

If pulse such that $P(f)$ is roughly constant (say c) for $\frac{-1}{2T_s} \leq f \leq \frac{1}{2T_s}$, then:

$$R[n] = cH_b(f_n)X[n] + N[n]$$

Can interpret OFDM as having N sub-carriers of frequencies $f_0, \ldots, f_{N-1}$. Spacing between adjacent sub-carriers is $\frac{1}{NT_s}$. In passband, the sub-carrier frequencies are $\{f_c + f_0, f_c + f_1, \ldots, f_c + f_{N-1}\}$.

Bandwidth of OFDM signal $= \frac{1}{T_s}$.

Rate of transmission: $\frac{N}{(L+N)T_s}$ QAM symbols per second

5 Channel Coding

5.1 Convolutional Codes

In *convolutional codes*, a stream of input bits is transformed into a stream of code bits using a shift register (filter).

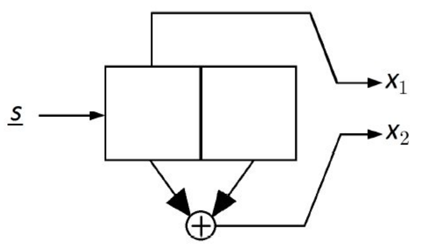

In a general convolutional code, at each time instant: k input bits are fed into a shift register with S stages, and its contents are used to produce n code bits, for a rate of $R = k/n$.

The generation of the code bits can be described by n *generators*, which indicate which shift register bits are added to generate each code bit.

$$g_1 = (10), \quad g_2 = (11)$$

We obtain the code bits $\underline{x} = \left(x_1^{(1)}, x_2^{(1)}, x_3^{(1)}, \ldots, x_1^{(m)}, x_2^{(m)}, x_3^{(m)}\right)$ simply by interleaving the convolutions of $\underline{s}$ with each of the generators: $\underline{x}_i = \underline{s} * g_i$. We can express this in terms of the Z-transform (with $Z = z^{-1}$:

$$g_1(Z) = 1, \quad g_2(Z) = 1 + Z$$

Hence $x(Z) = s\left(Z^3\right) g_1\left(Z^3\right) + Zs\left(Z^3\right) g_2\left(Z^3\right)$.

Convolutional codes can also be represented via *state diagrams* corresponding to *finite-state machines*.

5.2 Decoding Convolutional Codes

Optimal decoding is minimum distance decoding: If we receive $\underline{y}$, we need to find a path in the trellis which gives a code sequence $\hat{\underline{x}}$ minimising $d(\underline{y}, \hat{\underline{x}})$.

There is a simple *dynamic programming* algorithm to find the minimum-distance path called the *Viterbi algorithm*.

Consider the *free distance* d_{free} defined as the minimum weight among all codewords generated with the code starting and ending in the all-zero state. If $d_{\text{free}} = d$, then any collection of $\leq \left\lfloor \frac{d-1}{2} \right\rfloor$ errors can be corrected, as long as these *error bursts* are not *too close* to one another.

5.3 Transfer Function

Modified state diagram:

1. Label every branch by D raised to the Hamming weight of the sequence it produces.
2. Add a J to every branch, counting total path length.
3. Add an N to each branch corresponding to an input bit 1.

From the diagram we can obtain the *extended state equations* which we can solve to obtain the *extended transfer function* $T(J,N,D)$.

If the greatest common divisor of the polynomials $\{g_1(Z), g_2(Z), \ldots, g_n(Z)\}$ is of the form Z^ℓ for some integer $\ell \geq 0$, then the encoder defined by the generators $\{g_1, g_2, \ldots, g_n\}$ is not *catastrophic*.

6 Routing and Congestion Control

6.1 Transmission Control Protocol (TCP)

Congestion occurs when too many sources sending too much data too fast for the network to handle. It manifests as *packet loss* (buffer overflow at routers) or *large delay* (long queues at router buffers).

Let T denote the RTT: T = *round-trip time* between packet being sent and `ack` received. At time t, the rate $R(t) = \frac{W(t)}{T}$ packets/second.

TCP-Reno:

1. Variable *window size* W regulates transmission speed.
2. *Slow start:* Multiplicative (i.e., exponential) rate increase until `ssthresh`/T packets/second
3. *Congestion avoidance:* Additive (i.e., linear) rate increase until delay/loss detected.
4. *Fast retransmit* when delay detected.
5. *Fast recovery* when `timeout`/loss detected.

Loss (`timeout`) is treated as a more serious issue than delay (3 `dupack`'s).

Consider a continuous-time approximate model and let $q(t)$ = loss rate at time t.

At equilibrium, $R'(t) = W'(t)/T = 0$.

TCP square-root law:

$$R(t) = \frac{1}{T\sqrt{\beta}}\sqrt{\frac{1-q(t)}{q(t)}}$$

6.2 Dijkstra's Routing Algorithm

6.3 The Bellman-Ford Routing Algorithm

(The End)