

1 GSW... Multipath Channel Models

In the general case, the mobile radio channel is pretty unpleasant: there are a lot of echoes distorting the received signal, and the impulse response keeps changing. Fortunately, there are some simplifying assumptions that can often be used to make the mathematical description of the channel a bit easier, and which allow the determination of few simple parameters that provide a good indication the sort of problems any given system is likely to have.

This chapter is about these assumptions, and a couple of these simple parameters that can be derived using them that characterise the quality of a channel: most notably the delay spread and the coherence bandwidth. There's also a brief introduction to fading, but there's much more about time-varying channels in the next chapter on Time-Variant Multipath Channels.

1.1 The Mobile Radio Channel in General

In the most general case, the mobile radio channel can be characterised either by a time-dependent impulse response $h(\tau, t)$ or a time-dependent frequency response $H(\omega, t)$, where t is the time. This is a simple extension of the usual representation of linear time-invariant systems to the case where the system isn't time-invariant and the impulse response changes with time.

With time-variant systems, the output of the system for any given input signal $x(t)$ can be determined using the time-dependent impulse response:

$$y(t) = \int_{-\infty}^{\infty} h(\tau, t) x(t - \tau) d\tau \quad (0.1)$$

and in the general case, $h(\tau, t)$ can be anything.

The sort of information we'd like to know is over what range of time will the energy in a transmitted impulse arrive at the receiver, and how quickly the impulse response changes with time. Long impulse responses and fast-changing channels both cause problems for receivers.

1.1.1 The Problem of Long Impulse Responses

Long impulse responses (long relative to the length of a symbol) result in *intersymbol interference*. Consider a simple channel consisting of two rays, with a delay between them of just over one symbol period:

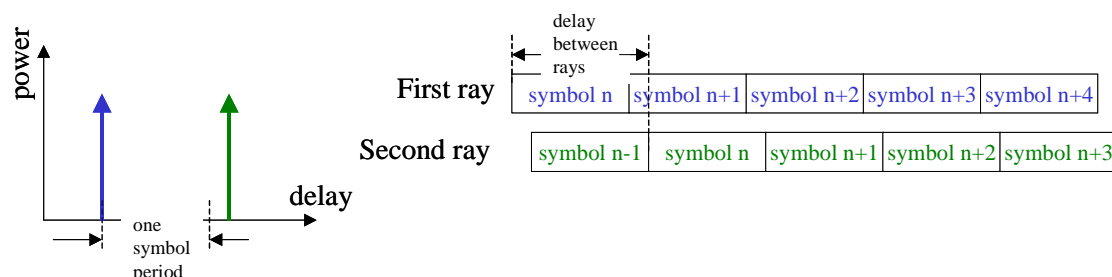


Figure 1-1 Intersymbol Interference Caused by Multipath

At all times, the receiver, which is receiving the sum of the signals from both rays in the impulse response, is receiving energy from two different symbols. The interference caused by the echo is known as intersymbol interference, and makes the receiver's task of working out

what symbol was transmitted much more difficult. (The part of the receiver with the task of removing the effects of multipath interference is called the *equaliser*. This is often the most difficult task a receiver has to perform.)

However, reduce the delay to one tenth of a symbol period (or alternatively increase the length of the symbol to ten times the delay between the rays), and the intersymbol interference almost disappears, as now the receiver is, for almost all of the time, only receiving energy from one symbol.

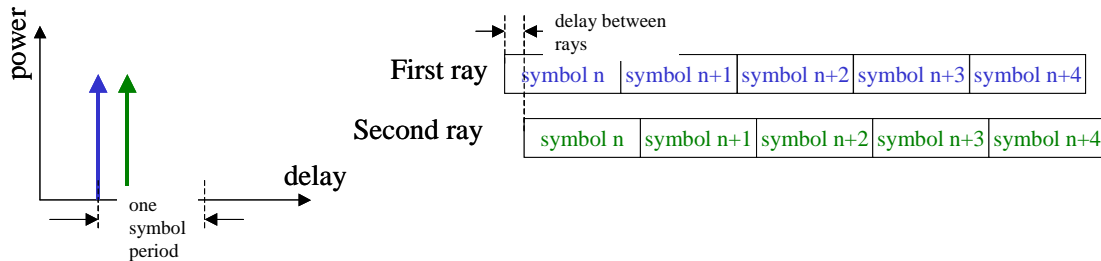


Figure 1-2 Multipath Causing Less Intersymbol Interference

The delay spread of a channel is a measure of the length of time over which the energy transmitted at one instant arrives at the receiver. It tells you how powerful an equaliser you will need when you're transmitting at a certain symbol rate; or alternatively, what symbol rate you can use without the need for an equaliser.

1.1.2 The Problem of Time-Varying Channels

Knowing how fast the channel is changing can be very useful for a number of reasons. An equaliser has to work out what the impulse response of the channel is, so it can undo the effects the multipath interference in the received signal, and this task is much harder if the channel keeps changing. Systems using equalisers typically work by transmitting a known series of bits (known as a *training sequence* or as *pilot symbols*) that allow the receiver to work out the impulse response of the channel. The receiver can then adapt the equaliser based on this channel impulse response.

The problem with using a training sequence is that the receiver works out what the channel impulse response was during the time when the training sequence is transmitted, not when the information is being transmitted; and the channel is changing all the time. Know how fast the impulse response of a channel is changing, and you know how often you need to transmit these training sequences so that the receiver can keep its equaliser up-to-date¹.

1.1.3 The Time-Variant Impulse Response

It's perhaps interesting to pause and consider what $h(\tau, t)$ is. It's the impulse response for an impulse leaving the transmitter at time t . It's not the impulse response of the channel at time t . If you took a photograph of the system at any time t_1 , and then tried to determine the impulse response $h(\tau, t_1)$ just by looking at the photograph, you couldn't do it.

¹ This is perhaps an over-simplification. For example, some designs of equaliser, once set-up using a training sequence, can track changes in the channel using the received data itself. However the general principle remains true: it is harder to design an accurate equaliser for a radio channel that is changing quickly.

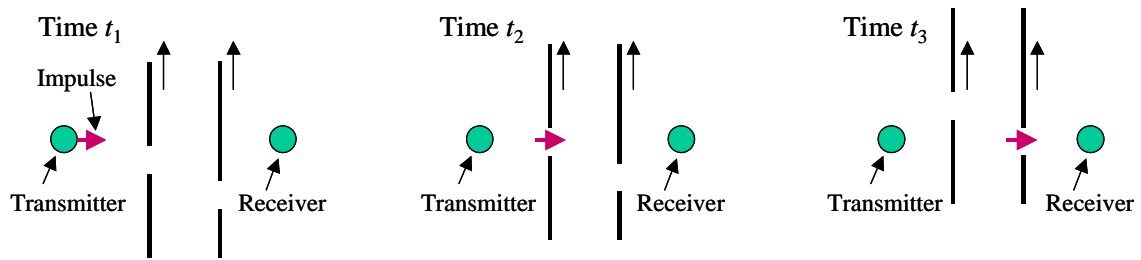


Figure 1-3 Path of an Impulse Through a Fast-Moving Channel

For example, consider the case shown in the figure above, where the direct path between the transmitter and the receiver is obstructed by two large sheets of copper, each with a hole in it. Both sheets are travelling upwards (very fast), and the holes are arranged so that an impulse leaving the transmitter at time t_1 would go straight through both holes and arrive at the receiver. However, at any given time, a photograph would show that the holes in the two sheets of copper are not in line, and that there is never a line-of-sight link between the transmitter and the receiver. The photograph does not contain enough information to work out $h(\tau, t)$, you need to know how fast things are moving in the environment as well².

In most cases, for the mobile radio channels that we're interested in, this point is largely academic, since the speed of just about anything in the environment is a very small fraction of the speed of light. This is the first of the simplifying assumption we can make: the impulse response does not change over the amount of time taken for an impulse to travel from the transmitter to a receiver.

This approximation is quite reasonable: consider a 6 km long radio channel (quite long for mobile radio), so that an impulse travelling at the speed of light would take 20 μs to get from the transmitter to the receiver. Anything travelling more than one millimetre in that time would be moving at more than 50 m/sec (180 km/hr), and anything moving less than one millimetre is unlikely to make a significant difference to the channel impulse response at the frequencies used by mobile phones (with wavelengths typically around 15 cm).

1.2 Clusters of Uncorrelated Scatterers

Another assumption that's often made is that the receiver receives a finite number of echoes of a transmitted impulse, and as the objects move in the environment of the transmitter and receiver, the amplitudes and relative phases of these echoes can change relatively quickly, but their delays change relatively slowly.

There are two physical models of scattering that predict exactly this. The first assumes that there are groups of *scatterers* (objects which reflect radio energy) spread all around the receiver, sort of like this:

² The same thing is true for the time-dependent frequency response as well. $H(\omega, t)$ is the gain of the component at frequency ω of an impulse transmitted at time t , and is the Fourier transform of the time-dependent impulse response.

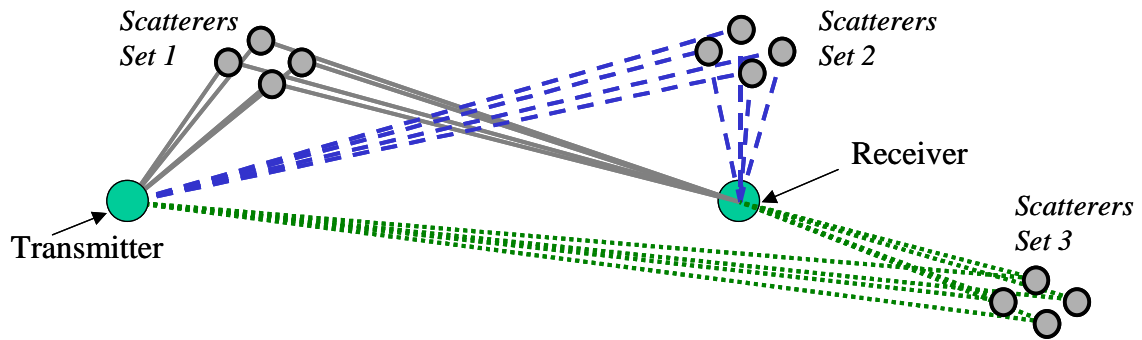


Figure 1-4 Model with Groups of Distributed Scatterers

The second, and in many cases rather more plausible explanation is that there are a lot of local scatterers around the receiver, and a smaller number of large reflectors spaced all around the receiver:

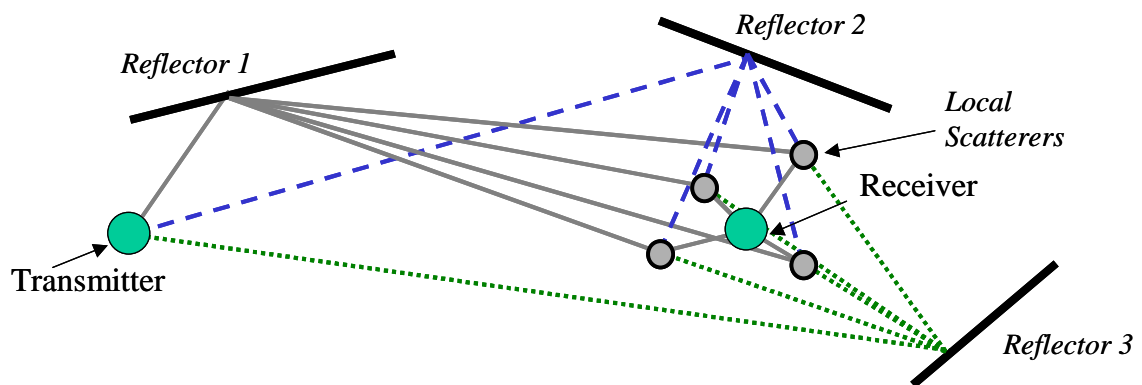


Figure 1-5 Model with Local Scatterers and Distant Reflectors

Both models predict that there will be groups of rays arriving at the receiver, with each group arriving at about the same time. If the difference in delays between two rays from the same group is too small, then the receiver cannot resolve these into two separate rays³, they effectively merge into one ray, with an amplitude and phase given by the sum of the vectors representing the individual rays in the group. In both cases, the impulse response then looks something like the following diagram, with the amplitude and relative phases of each of the groups of rays changing with time:

³ What ultimately determines whether two rays arriving at slightly different times are resolvable or not is the sampling rate and bandwidth of the receiver. If information about a change in the transmitted signal arrives via two different rays less than a sample period apart, then the receiver will not be able to distinguish the rays: in one sample neither has the new information, the next time the receiver 'looks', they both have. Receivers typically sample the incoming signal at between two and four times the transmitted symbol rate, so any rays arriving within a small fraction of the symbol rate will not be resolved. This means that the number of resolvable multipath rays in a radio channel is a function of the symbol rate: systems with higher transmitted symbol rates (and hence higher sampling rates) have more resolvable multipath components in the channel impulse response.

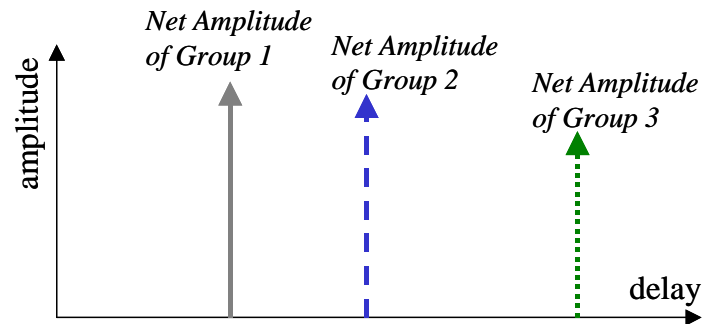


Figure 1-6 Impulse Response for Groups of Scatterers

In order to change the phase of the received signal from each group of scatterers, the receiver (or possibly the scatterers themselves) have to move a significant fraction of a wavelength. At typical mobile phone frequencies (about 2 GHz), this implies a movement of a few centimetres. That's all it takes to significantly change the relative phase of the component rays in any of these group of rays, and hence the amplitude and phase of the resultant of these rays.

However, to change the delay associated with this ray (the time between transmitting the signal and receiving the energy from this group of scattered rays), the receiver would have to move a significant distance towards (or away from) the transmitter. Even assuming a fast symbol rate of 5 Mbaud, to make a significant difference to the delay itself (say, one-tenth of a symbol period) would require the receiver to move 60 meters. The difference is over three orders of magnitude: the phase and amplitudes of the resolvable rays arriving at the receivers change much more rapidly than their delays.

It's often helpful to plot the impulse responses as a function of time. In the case of the three clusters of scatterers considered here, we might get something a bit like this:

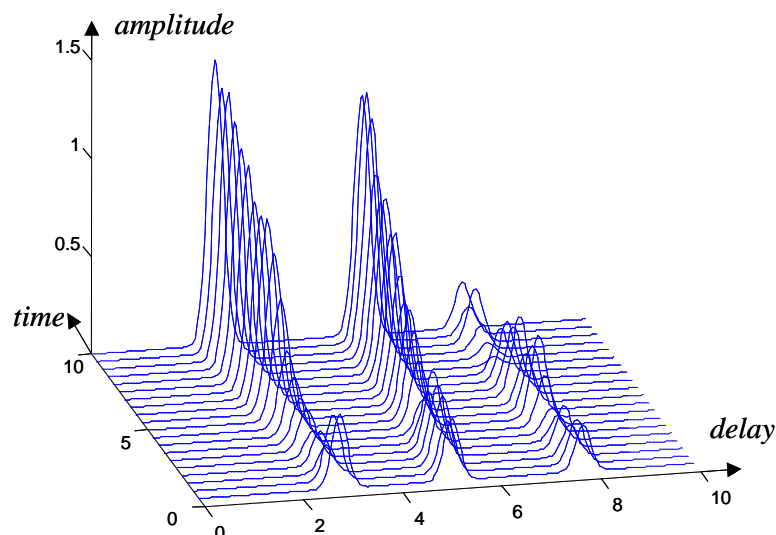


Figure 1-7 Example Time-Variant Three-Ray Impulse Response

Note that while the delays of the three rays don't change very fast with time, the amplitudes of the rays do change. This is typical of most mobile radio channels.

1.3 Multipath and the Power Delay Profile

Although the amplitude of each of the resolvable multipath rays in the channel impulse response is constantly changing, the power received in each ray can usually be averaged over a time during which the delays of the rays themselves are reasonably constant (as we've seen, the delays usually change much less rapidly than the amplitudes and phases of these rays).

Plotting this average power against delay gives the *power delay profile*. A typical power delay profile might look something like:

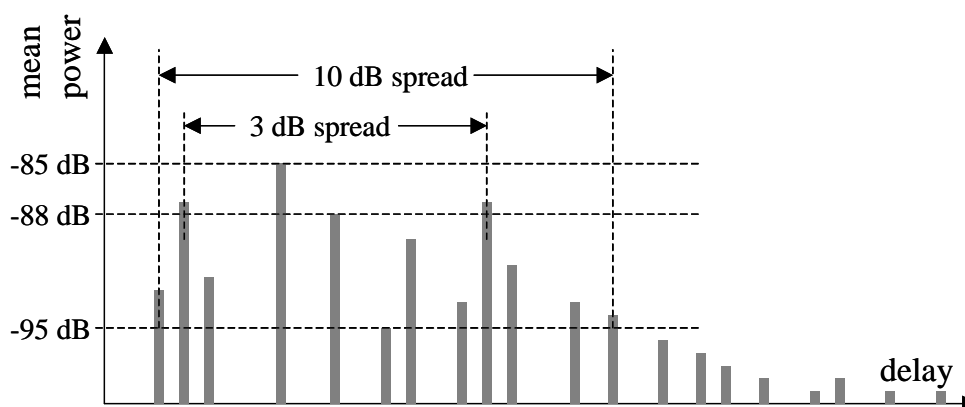


Figure 1-8 Sample Power-Delay Profile

(Remember that in this figure the short-term fading has been averaged out over time: this is a plot of the mean power received in each ray, and remains valid over the length of time that a receiver moves sufficiently far to change the delays of the rays.) The power-delay profile plot is useful, since it provides a simple visual image of the average range of times over which the energy transmitted in one symbol arrives at the receiver, and hence a measure of the intersymbol interference expected.

Whilst a picture is useful, what would be even more useful is a single number that can be used to describe this range of times, and tell us how much intersymbol interference to expect from the channel. The most obvious way to come up with a single number (the difference in time between the first multipath ray arriving and the last multipath ray arriving) isn't very helpful, since in theory at least the multipath rays decay away for ever⁴. At some point you have to say "there is so little energy arriving after this time I'm just going to ignore it".

Two other possibilities are to define the spread of the channel in terms of the range of delays over which rays arrive within 3 dB or 10 dB of the ray with the maximum power (the '3dB spread' and '10 dB spread', see the figure above). These have the advantage that they are very easy to measure from a plot of the power delay profile, but the disadvantage that they can give misleading answers. For example, consider the following two power delay profiles:

⁴ Imagine a receiver on a city street between two buildings. In theory, some of the energy arriving at the receiver will have bounced between the two buildings thousands of times before ending up at the receiver, and have a huge delay. Of course, the amount of energy that arrives after this length of time is negligible.

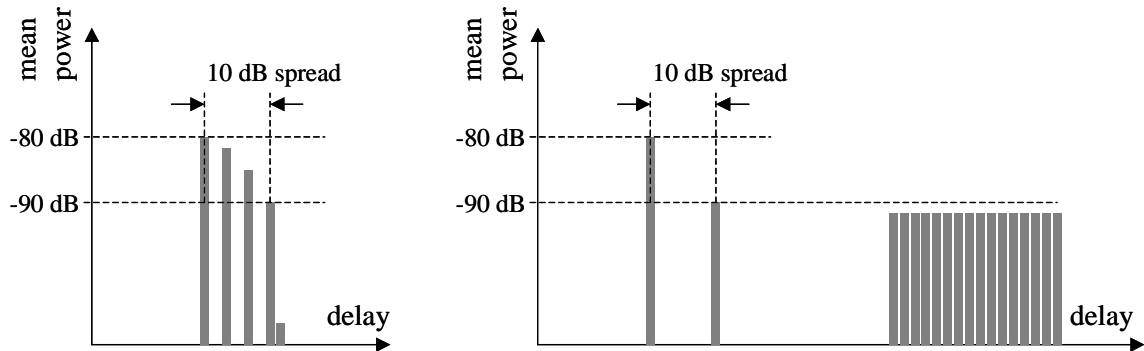


Figure 1-9 Two Power Delay Profiles with the Same 10 dB Delay Spread

Both of these delay profiles have the same 10 dB delay spread, but the second one actually has most of the received power arriving after a long delay; it's just that none of these individual rays quite manages to make it to within 10 dB of the largest beam. Effectively, most of the received power is being ignored.

It's a bit of a silly example perhaps, and totally unrealistic, but it does illustrate the point that the 3 dB or 10 dB delay spread figure completely neglects all of the power arriving outside the range specified, and there could be a considerable amount of this power, spread out over a wide range of delays, being ignored. That power could cause a lot of intersymbol interference.

The usual solution is to use the *rms delay spread* (sometimes just called the *delay spread* since it's so commonly used). This figure includes all of the received power, weighted by what effect it has on the amount of intersymbol interference it causes: exactly what we want to know.

1.3.1 The rms Delay Spread

If we normalise the power delay profile by ensuring that the area under the curve is one (or equivalently that the sum of all the powers in all the rays is one), then we have a probability density function, describing the probability that any random photon of energy arriving at the receiver has taken a given length of time to get from the transmitter to the receiver. The standard deviation of this probability density function is the *rms delay spread*. It's a good measure of the range of delays over which most of the power arrives.

In other words, the rms delay spread is the square root of the mean value of the square of the difference between the delay experienced by each photon of energy and the average delay. It's probably easier to understand in mathematical notation:

$$\Delta = \sqrt{\frac{\sum_i P_i (\tau_i - \bar{\tau})^2}{\sum_i P_i}} \quad (0.2)$$

where Δ is the rms delay spread, P_i is the received power in the i^{th} ray, τ_i is the delay of the i^{th} ray, and $\bar{\tau}$ is the mean delay given by:

$$\bar{\tau} = \frac{\sum_i P_i \tau_i}{\sum_i P_i} \quad (0.3)$$

Consider a simple example: the impulse response shown in the figure below:

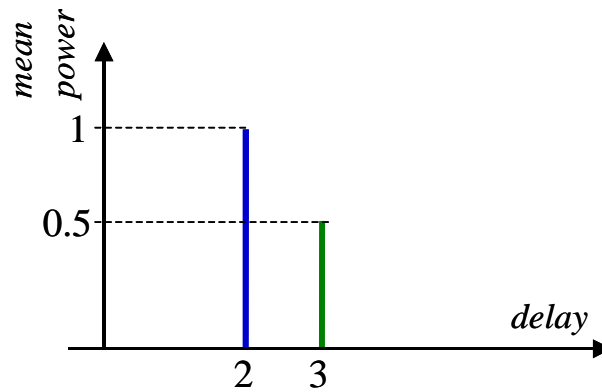


Figure 1-10 Simple Power-Delay Profile

This consists of just two rays: arriving at times of 2 and 3 units after the impulse was transmitted. The first ray arrives with a power of 1, and the second ray arrives with exactly half this power (it doesn't matter what the units of power are, they cancel out).

Then, the mean delay is:

$$\bar{\tau} = \frac{\sum_i P_i \tau_i}{\sum_i P_i} = \frac{1 \times 2 + 0.5 \times 3}{1 + 0.5} = \frac{3.5}{1.5} = \frac{7}{3} \quad (0.4)$$

and the rms delay spread is:

$$\Delta = \sqrt{\frac{\sum_i P_i (\tau_i - \bar{\tau})^2}{\sum_i P_i}} = \sqrt{\frac{1 \times (2 - 7/3)^2 + 0.5 \times (3 - 7/3)^2}{1 + 0.5}} = \sqrt{\frac{1/9 + 0.5 \times 4/9}{1 + 0.5}} = \frac{\sqrt{2}}{3} \quad (0.5)$$

and in both cases the units are whatever units the delays of the original rays were specified in⁵.

⁵ There's another way to do this sum which is often easier, and that's to use the fact that the variance of a distribution is the mean of the square of the values minus the square of the mean of the values. Here, this gives:

$$\Delta^2 = \frac{\sum_i P_i \tau_i^2}{\sum_i P_i} - \left(\frac{\sum_i P_i \tau_i}{\sum_i P_i} \right)^2$$

and:

$$\Delta = \sqrt{\frac{1 \times 2^2 + 0.5 \times 3^2}{1 + 0.5} - \left(\frac{1 \times 2 + 0.5 \times 3}{1 + 0.5} \right)^2} = \sqrt{\frac{8.5}{1.5} - \left(\frac{3.5}{1.5} \right)^2} = \sqrt{\frac{2}{9}} = \frac{\sqrt{2}}{3}$$

1.3.2 Delay Spread of Continuous Power-Delay Profiles

In some cases the power-delay profile doesn't consist of a finite number of discrete rays, arriving with different times, but it looks much more like a continuous curve, with some power occurring with any value of delay. We might get a power-delay profile that looks more like this:

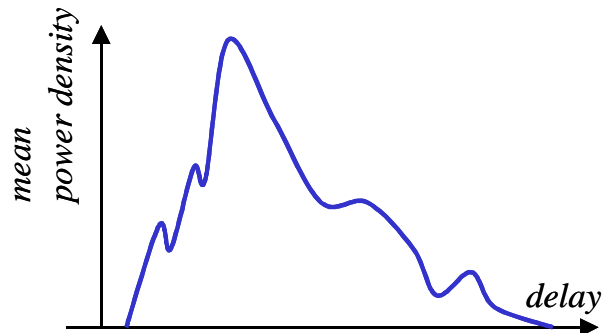


Figure 1-11 An Example Continuous Power Delay Profile

where the function plotted is now a power density: the amount of power that arrives within a small range of delays $d\tau$ centred around a delay τ , divided by $d\tau$. This isn't a problem: we can calculate a delay spread in exactly the same way: normalise the power-delay profile so that it has a total received power of one, and then calculate the standard deviation of the resultant probability density function. The only difference is that we now have to integrate over all possible delays, rather than just summing over the individual rays:

$$\Delta = \sqrt{\frac{\int_{-\infty}^{\infty} P(\tau)(\tau - \bar{\tau})^2 d\tau}{\int_{-\infty}^{\infty} P(\tau) d\tau}} \quad (0.6)$$

where:

$$\bar{\tau} = \frac{\int_{-\infty}^{\infty} P(\tau)\tau d\tau}{\int_{-\infty}^{\infty} P(\tau) d\tau} \quad (0.7)$$

See the problems for an example of this.

1.3.3 Problems with the rms Delay Spread

Sadly, the rms Delay Spread does have a problem when it is used in real life, and that is noise. If you try and measure a mobile channel, you'll get a result that has some noise on it, for example:

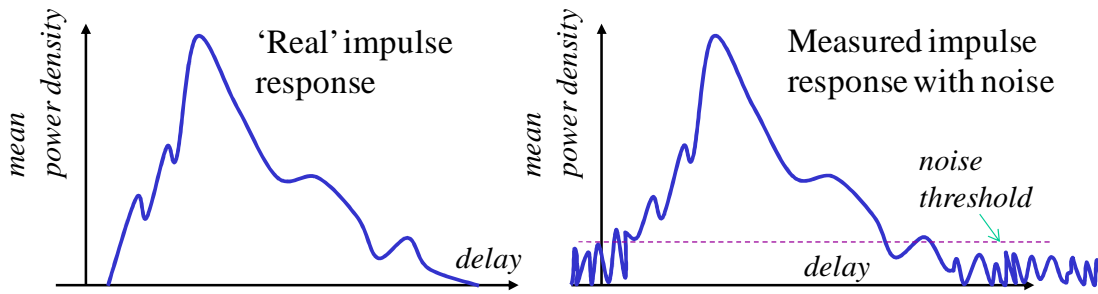


Figure 12 - Measured Impulse Response with Noise

With the noise, it is not obvious where the impulse response ends. If we include all the noise in the calculation of the rms Delay Spread, we'd end up with an answer that's much too big (in theory, infinite). On the other hand, if we use some 'noise threshold' (as shown in the figure above) and ignore everything below that limit, we would be missing some parts of the real impulse response, and that would (usually) give an answer that was too low.

Without some knowledge of the shape of the impulse response (for example, that it looks like an exponential decay), it can be very difficult to estimate a good value for delay spread in cases where there are a lot of very low energy rays arriving at different times (and hence a lot of energy arriving below this noise threshold).

1.4 The Coherence Bandwidth

Any parameter or result in signals and systems theory expressed in the time domain has a corresponding parameter or result in the frequency domain. In the case of the delay spread, the corresponding parameter in the frequency domain is the *coherence bandwidth*: the range of frequencies over which the gain of the channel remains 'about the same'.

1.4.1 A Simple Example of Coherence Bandwidth

Perhaps a simple example might help to illustrate the concept. About the simplest example I can think of is a channel impulse response with just two rays in it (one arriving after a delay time τ_1 , the other after a delay of τ_2), each with the same amplitude A . Input a single frequency cosine wave with unit amplitude $\cos(\omega t)$ into this channel, and the receiver will receive the sum of these two delayed rays:

$$r(\omega, t) = A \cos(\omega(t - \tau_1)) + A \cos(\omega(t - \tau_2)) \quad (0.8)$$

and using the formula for the sum of two cosines, this can be expressed as:

$$r(\omega, t) = 2A \left\{ \cos \left(\omega t - \omega \left(\frac{\tau_1 + \tau_2}{2} \right) \right) \cos \left(\omega \left(\frac{\tau_2 - \tau_1}{2} \right) \right) \right\} \quad (0.9)$$

which represents a cosine wave with an amplitude of:

$$\left| 2A \cos \left(\omega \left(\frac{\tau_2 - \tau_1}{2} \right) \right) \right| \quad (0.10)$$

and a phase relative to the transmitted signal $\cos(\omega t)$ of:

$$\phi = -\omega \left(\frac{\tau_2 + \tau_1}{2} \right) \quad (0.11)$$

(at least when $\cos(\omega(\tau_2 - \tau_1)/2)$ is positive; it's π greater than this when $\cos(\omega(\tau_2 - \tau_1)/2)$ is negative).

If we plotted the impulse response, and the amplitude of the output from this channel as a function of frequency, we'd get something like this:

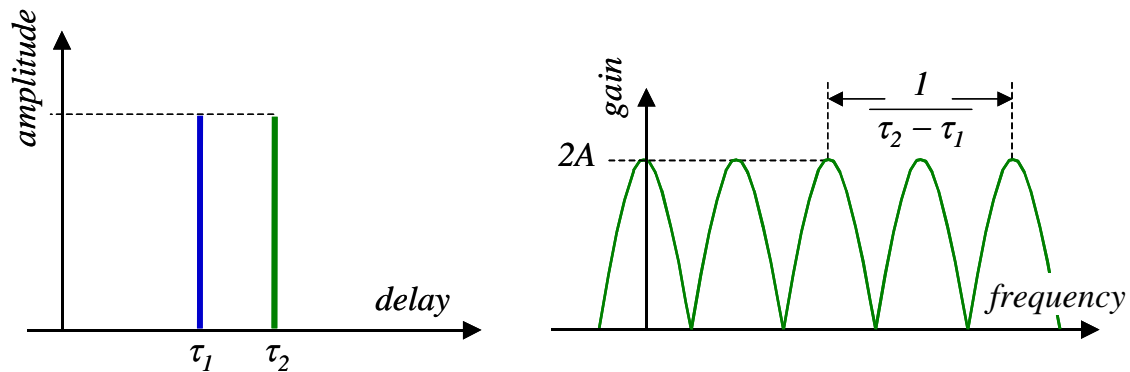


Figure 1-13 Impulse Response and Gain of Simple Channel

Put a second cosine wave at a slightly different frequency $\omega + \nu$ through the same channel, and the amplitude of the received signal becomes:

$$\left| 2A \cos \left((\omega + \nu) \left(\frac{\tau_2 - \tau_1}{2} \right) \right) \right| \quad (0.12)$$

The question we need to answer is: how similar is the response of this channel at these two frequencies? Clearly if $\nu = 0$, then the channel has the same gain for both frequencies, but what if ν is a small frequency difference? The channel gain changes smoothly, so the gain would be similar at two slightly different frequencies, although very different when the frequencies are more widely separated. Even this very simple case turns out to be rather awkward to solve, mostly due to the fact that we're trying to compare two oscillations that can differ in both amplitude and in phase.

The usual mathematical way of quantifying how similar two variables are is in terms of their correlation co-efficient, defined as the mean value of the product of the complex conjugate of one function (minus its mean) with the other function (minus its mean), divided by their standard deviations⁶:

$$\chi = \frac{E \left\{ (x - \bar{x})^* (y - \bar{y}) \right\}}{\sigma_x \sigma_y} \quad (0.13)$$

⁶ We have to take the complex conjugate of one of the functions in the case where the functions are complex. Doing so ensures that the correlation coefficient is one when the two functions are identical, since one number multiplied by its complex conjugate is always real.

This gives a value of one if the x and y are perfectly correlated (in the sense that y is identical to x or at least a constant multiple of x), and zero if the value of y is completely independent of the value of x .

When the two quantities can differ in both amplitude and phase, the task of deciding how similar they are has an added complexity. The usual technique is to use complex numbers to represent each component in the impulse response of the channel: the amplitude and phase of each oscillation being represented by the amplitude and phase of a complex number. For example, the signal in the first ray:

$$A \cos(\omega(t - \tau_1)) = A \cos(\omega t - \omega \tau_1) \quad (0.14)$$

would be represented by a complex number with amplitude A and phase angle $-\omega \tau_1$:

$$A \exp(-j\omega \tau_1) \quad (0.15)$$

and the real signal in this ray can be re-created by multiplying this complex number by a complex oscillation at the original frequency, and then taking the real part of the result:

$$\begin{aligned} \Re\{A \exp(-j\omega \tau_1) \exp(j\omega t)\} &= \Re\{A \exp(j\omega(t - \tau_1))\} \\ &= A \cos(\omega(t - \tau_1)) \end{aligned} \quad (0.16)$$

Using this notation, we can represent the sum of the two cosine waves in equation (0.9) in complex form as:

$$\begin{aligned} r(\omega, t) &= A \cos(\omega(t - \tau_1)) + A \cos(\omega(t - \tau_2)) \\ &= \Re\{A \exp(j\omega(t - \tau_1))\} + \Re\{A \exp(j\omega(t - \tau_2))\} \\ &= \Re\{A(\exp(-j\omega \tau_1) + \exp(-j\omega \tau_2)) \exp(j\omega t)\} \end{aligned} \quad (0.17)$$

and then use the complex representation of this signal:

$$r_{\mathbf{E}}(\omega, t) = A(\exp(-j\omega \tau_1) + \exp(-j\omega \tau_2)) \quad (0.18)$$

Similarly, we can use as complex representation of the channel's output at the higher frequency $\omega + \nu$, and use the complex representation of the sum of the two rays at this frequency:

$$r_{\mathbf{E}}(\omega + \nu, t) = A(\exp(-j(\omega + \nu)\tau_1) + \exp(-j(\omega + \nu)\tau_2)) \quad (0.19)$$

We then need to find the correlation coefficient between these two complex representations. One big advantage of this technique is that the mean value of both of these signals is zero: a complex exponential oscillation has a zero mean, and the mean of the sum of two variables is the sum of the means of each variable.

This leaves us with the problem of finding the correlation between $r_{\mathbf{E}}(\omega, t)$ and $r_{\mathbf{E}}(\omega + \nu, t)$ as a function of the frequency difference ν , averaged over all possible values of ω . The expectation value of the product of the complex conjugate of $r_{\mathbf{E}}(\omega, t)$ and $r_{\mathbf{E}}(\omega + \nu, t)$, from equations (0.18) and (0.19) gives:

$$E\left\{A^2(\exp(j\omega\tau_1) + \exp(j\omega\tau_2))(\exp(-j(\omega + \nu)\tau_1) + \exp(-j(\omega + \nu)\tau_2))\right\} \quad (0.20)$$

and multiplying out the brackets gives the four terms:

$$E\left\{A^2\left(\begin{array}{l} \exp(-j\nu\tau_1) + \exp(j\omega(\tau_2 - \tau_1))\exp(-j\nu\tau_1) \\ + \exp(j\omega(\tau_1 - \tau_2))\exp(-j\nu\tau_2) + \exp(-j\nu\tau_2) \end{array}\right)\right\} \quad (0.21)$$

The middle two terms are functions of ω , and average out to zero, leaving just the first and last terms when the expectation value is taken:

$$A^2(\exp(-j\nu\tau_1) + \exp(-j\nu\tau_2)) \quad (0.22)$$

Rather than calculate the standard deviations for the denominator of the correlation coefficient, we can take a short cut: we know the correlation coefficient will be one when $\nu = 0$, and when $\nu = 0$ this expectation value is $2A^2$, so we can write the final correlation coefficient of the frequency response of the two channels as:

$$\chi(\nu) = \frac{(\exp(-j\nu\tau_1) + \exp(-j\nu\tau_2))}{2} \quad (0.23)$$

This is a complex correlation coefficient: if we wanted to know over what range the absolute magnitude of the correlation coefficient had a value greater than χ , we'd have to use:

$$\begin{aligned} |\chi(\nu)|^2 &= \chi^*(\nu)\chi(\nu) \\ &= \frac{(\exp(-j\nu\tau_1) + \exp(-j\nu\tau_2))^* (\exp(-j\nu\tau_1) + \exp(-j\nu\tau_2))}{4} \\ &= \frac{2 + 2\cos(\nu(\tau_2 - \tau_1))}{4} \\ &= \frac{1 + \cos(\nu(\tau_2 - \tau_1))}{2} \end{aligned} \quad (0.24)$$

so to determine the range of frequencies over which the channel gain has a correlation with a magnitude of greater than χ , we'd have to use:

$$\nu = \frac{\cos^{-1}(2|\chi|^2 - 1)}{\tau_2 - \tau_1} \quad (0.25)$$

In this case (two equal powered rays arriving after a delay of τ_1 and τ_2 respectively), the delay spread is:

$$\Delta = \frac{\tau_2 - \tau_1}{2} \quad (0.26)$$

so we could write:

$$\nu = \frac{\cos^{-1}\left(2|\chi|^2 - 1\right)}{2\Delta} \quad (0.27)$$

which is interesting, since it shows that the coherence bandwidth is inversely proportional to the delay spread, the constant of proportionality depending on the correlation coefficient of the channel gains over the range of frequencies ν (in other words, how close to the same you want the channel response to be at the two frequencies considered).

Putting some numbers into this: consider an impulse response that looks like this:

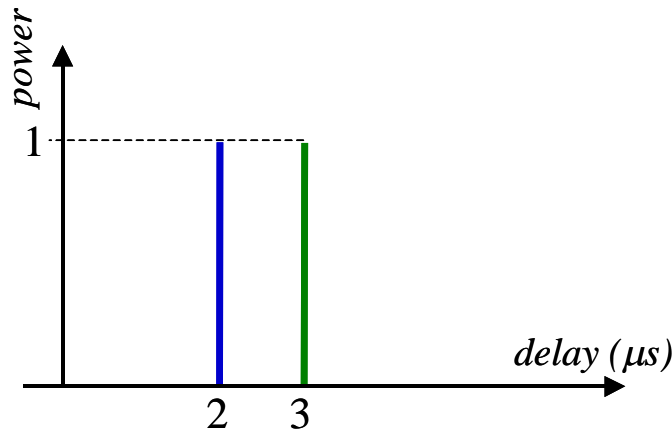


Figure 1-14 Simple Two-Ray Power Delay Profile

Then, if we wanted to know what range of frequencies could be used so that there was a correlation with a magnitude of at least 0.9 between any two frequencies, we could first calculate the delay spread ($0.5 \mu\text{s}$ in this case) and then:

$$\nu = \frac{\cos^{-1}\left(2 \times 0.9^2 - 1\right)}{2 \times 0.5 \mu\text{s}} = 0.902 \text{ Mrad/s} = 144 \text{ kHz} \quad (0.28)$$

but for a correlation of at least 0.5 between frequencies, the coherence bandwidth increases to:

$$\nu = \frac{\cos^{-1}\left(2 \times 0.5^2 - 1\right)}{2 \times 0.5 \mu\text{s}} = 2.09 \text{ Mrad/s} = 333 \text{ kHz} \quad (0.29)$$

This is a very simple case (two equal powered rays), but the general conclusion remains true for all other channels: the coherence bandwidth is inversely proportional to the delay spread, the constant of proportionality depending on the power delay profile of the channel, and the maximum correlation coefficient required between the channel gain at different frequencies within the coherence bandwidth.

1.4.2 Calculating the Coherence Bandwidth in the General Case

For the general case, a simple expression for the coherence bandwidth can be determined from the impulse response of the channel and some results from Fourier theory. Any channel with a time-dependent impulse response of $h(\tau, t)$ has a time-dependent frequency response of $H(\omega, t)$ obtained by taking the Fourier transform of the impulse response with respect to the delay τ , so we can calculate the channel frequency response at any given frequency using the Fourier transform integral:

$$H(\omega, t) = \int_{-\infty}^{\infty} h(\tau, t) \exp(-j\omega\tau) d\tau \quad (0.30)$$

and the correlation coefficient between the frequency response of the channel at two different frequencies is then:

$$\chi(\nu) = \frac{\mathbb{E}\left\{\left(H(\omega, t) - \overline{H(\omega, t)}\right)^* \left(H(\omega + \nu, t) - \overline{H(\omega + \nu, t)}\right)\right\}}{\sigma_{H(\omega, t)}^2} \quad (0.31)$$

Now for all channels we're interested in, the mean value of the frequency response is zero (the phase of the output from each ray constantly rotates in phase as the frequency changes, averaging out to zero over all frequencies), so we can simplify this to:

$$\chi(\nu) = \frac{\mathbb{E}\left\{H^*(\omega, t) H(\omega + \nu, t)\right\}}{\sigma_{H(\omega, t)}^2} \quad (0.32)$$

and since we know that the correlation co-efficient must be one when ν is zero (the same trick as we used above to avoid having to work out the standard deviation), we can replace the denominator:

$$\chi(\nu) = \frac{\mathbb{E}\left\{H^*(\omega, t) H(\omega + \nu, t)\right\}}{\mathbb{E}\left\{H^*(\omega, t) H(\omega, t)\right\}} \quad (0.33)$$

The next step uses a version of the Wiener-Khinchin theorem from Fourier theory: the autocorrelation of the frequency response is 2π times the Fourier transform of the square of the modulus of the impulse response (in other words, of the power delay profile)⁷:

⁷ Consider taking the Fourier transform of the autocorrelation of the frequency response (that's the autocorrelation function, not the correlation co-efficient), defined as:

$$\phi(\nu) = \int_{-\infty}^{\infty} H^*(\omega) H(\omega + \nu) d\omega$$

Taking the inverse Fourier transform of this autocorrelation function and reversing the order of integration gives:

[continued on next page...]

$$\int_{-\infty}^{\infty} H^*(\omega, t) H(\omega + \nu, t) d\omega = 2\pi \int_{-\infty}^{\infty} |h(\tau, t)|^2 \exp(-j\nu\tau) d\tau \quad (0.34)$$

Now the expectation value of $H^*(\omega, t) H(\omega + \nu, t)$ is the mean value of this product over all possible frequencies, which could be expressed as:

$$E\{H^*(\omega, t) H(\omega + \nu, t)\} = \lim_{s \rightarrow \infty} \frac{1}{2s} \int_{-s}^s H^*(\omega, t) H(\omega + \nu, t) d\omega \quad (0.35)$$

so we could express our normalised correlation coefficient as:

$$\chi(\nu) = \frac{E\{H^*(\omega, t) H(\omega + \nu, t)\}}{E\{H^*(\omega, t) H(\omega, t)\}} = \lim_{s \rightarrow \infty} \frac{\frac{1}{2s} \int_{-s}^s H^*(\omega, t) H(\omega + \nu, t) d\omega}{\frac{1}{2s} \int_{-s}^s H^*(\omega, t) H(\omega, t) d\omega} \quad (0.36)$$

which when cancelling out the factors of $2s$ and using the result from equation (0.34) gives:

$$\chi(\nu) = \frac{\int_{-\infty}^{\infty} |h(\tau, t)|^2 \exp(-j\nu\tau) d\tau}{\int_{-\infty}^{\infty} |h(\tau, t)|^2 d\tau} \quad (0.37)$$

$$\begin{aligned} \frac{1}{2\pi} \int_{-\infty}^{\infty} \phi(\nu) \exp(j\nu\tau) d\nu &= \frac{1}{2\pi} \int_{-\infty}^{\infty} \left(\int_{-\infty}^{\infty} H^*(\omega) H(\omega + \nu) d\omega \right) \exp(j\nu\tau) d\nu \\ &= \frac{1}{2\pi} \int_{-\infty}^{\infty} H^*(\omega) \int_{-\infty}^{\infty} H(\omega + \nu) \exp(j\nu\tau) d\nu d\omega \\ &= \int_{-\infty}^{\infty} H^*(\omega) \left(\frac{1}{2\pi} \int_{-\infty}^{\infty} H(\omega + \nu) \exp(j(\omega + \nu)\tau) d\nu \right) \exp(-j\omega\tau) d\omega \\ &= \int_{-\infty}^{\infty} H^*(\omega) h(\tau) \exp(-j\omega\tau) d\omega \\ &= h(\tau) \int_{-\infty}^{\infty} H^*(\omega) \exp(-j\omega\tau) d\omega \\ &= 2\pi h(\tau) h^*(\tau) = 2\pi |h(\tau)|^2 \end{aligned}$$

and if the inverse Fourier transform of the autocorrelation function is 2π times the power delay profile, then the autocorrelation function must be the Fourier transform of 2π times power delay profile.

$$\phi(\nu) = 2\pi \int_{-\infty}^{\infty} |h(\tau)|^2 \exp(-j\nu\tau) d\tau$$

This provides a simple way⁸ to calculate the coherence bandwidth: all you need to do is take the Fourier transform of the power delay profile of the channel and then normalise it.

1.4.3 An Example of Coherence Bandwidth

Consider an example: the amplitude of the impulse response, power delay profile, frequency response and frequency correlation of a radio channel are shown in the figure below:

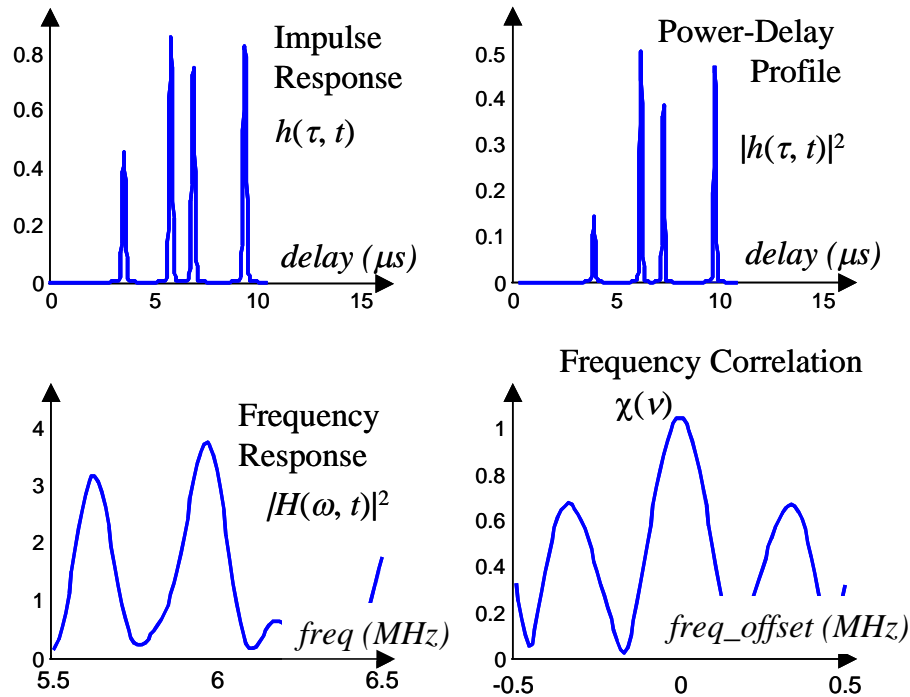


Figure 1-15 Correlation Bandwidth of Example Channel

(I've plotted the frequency response of this impulse response from 5.5 to 6.5 MHz, rather than around zero, since to a radio transmission, the only interesting part of the spectrum is around the carrier frequency, which I've assumed here is 6 MHz.)

Of course, this just provides a plot of the coherence between two frequencies. Again, just as with the delay spread, we'd ideally like a single number giving the range of frequencies over which the response of the channel is approximately the same. This channel has a delay spread of about 1.7 μs , and the channel changes significantly in gain in around 0.11 MHz (that's $1/5\Delta$). It's clear, I think, that the channel has a completely different (i.e. totally uncorrelated) gain if you move in frequency by 0.58 MHz ($1/\Delta$). If you want two frequencies that have approximately the same gain, you can't move more than about 0.058 ($1/50\Delta$) MHz; see the expanded figure below:

⁸ I should apologise to any mathematicians reading this who are probably spitting with fury at this point. I've taken a few short cuts with this derivation, for example, the autocorrelation as I've defined it can take infinite values, which means you can't take the Fourier transform. However, I hope this simplified derivation at least serves to indicate where the result comes from. For more discussion on this point, see the chapter on Fourier theory.

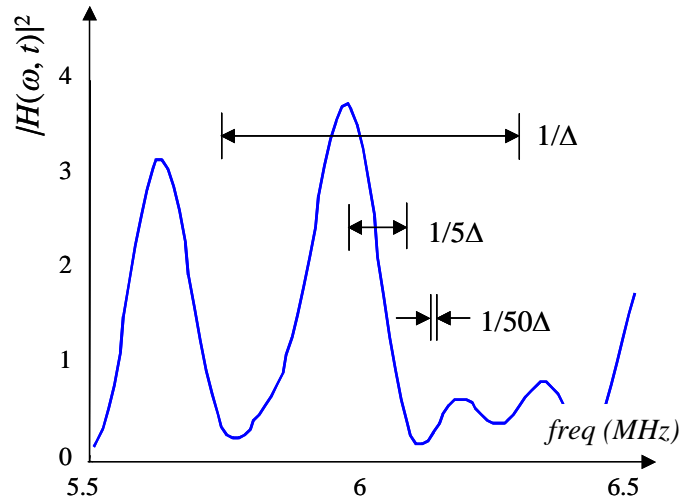


Figure 1-16 Channel Frequency Response and Coherence Bandwidths

Which figure is of the most use depends on what you want to use it for. If you want to know whether the channel looks flat, then you'll want a coherence bandwidth over which the gain of the channel is approximately the same, perhaps with a correlation coefficient of greater than 0.9. If you want to know how by how much you have to change the frequency to be confident that the gain of the channel has completely changed (perhaps the signal has faded, and you're thinking of choosing a new carrier frequency to use), you'll want a correlation coefficient of close to zero. If you want to know whether you need to use an equaliser or not, you'll want a correlation coefficient of about⁹ 0.5.

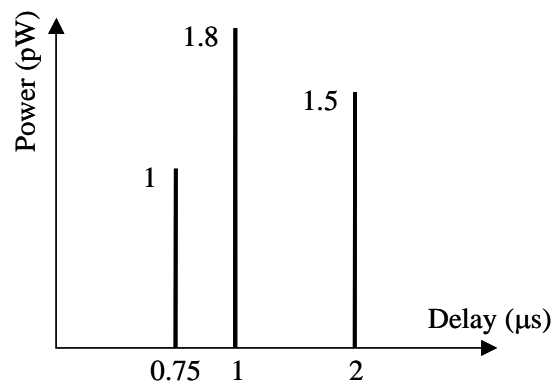
The exact relationship between the rms delay spread and the coherence bandwidth is dependent on the shape of the power delay profile, but in every case, the coherence bandwidth is inversely proportional to the delay spread. Increase the delay spread, and you reduce the range of frequencies over which the channel appears to have a similar gain. I'll finish this chapter with a summary of some useful approximations:

<i>Approximation</i>	<i>Correlation coefficient</i>	<i>Notes</i>
$B_{\omega 1} \approx \frac{1}{\Delta_{rms}}$	~ 0	The channel gain at frequency ω is almost entirely independent of the gain at $\omega + B_{\omega 1}$. Changing frequency by $B_{\omega 1}$ results in a channel with a completely different channel gain.
$B_{\omega 2} \approx \frac{1}{5\Delta_{rms}}$	~ 0.5	The channel gain at frequency ω is similar to the gain at $\omega + B_{\omega 2}$. If $B_{\omega 2}$ is greater than the symbol rate, an equaliser may not be required.
$B_{\omega 3} \approx \frac{1}{50\Delta_{rms}}$	> 0.9	The channel gain of frequency ω is almost exactly the same as the gain at $\omega + B_{\omega 3}$.

⁹ This is a rough rule-of-thumb only.

1.5 Problems

1) A channel has the impulse response shown in the figure below:



- i) What is the mean delay of this channel?
- ii) What is the delay spread?
- iii) Estimate the coherence bandwidth.

Estimate the maximum symbol rate that can be used through this channel without the need for an equaliser. (Assume an equaliser is not required if the delay spread is less than one-fifth of the symbol period.)

2) A radio channel has an impulse response containing two rays of equal power, one with a delay of $1 \mu\text{s}$, and the other with a delay of $3 \mu\text{s}$. Calculate the rms delay spread for this channel, and hence estimate the difference in frequencies for which the correlation between the gains is 0.5 by using $1/5\Delta$. Then calculate the exact difference in frequencies that gives an expected correlation of 0.5. How accurate is the approximate formula in this case?

What is the actual value of the correlation between two frequencies $1/5\Delta$ apart in frequency?

3) Repeat question 2 for a radio channel with an impulse response of $h(\tau) = \exp(-a\tau)$. How accurate is the approximation in this case?

4) What coherence bandwidth gives a maximum correlation of 0.9 between the channel gains for the impulse response in question 1?