



The how and why of COFDM

J.H. Stott

BBC Research and Development

Coded Orthogonal Frequency Division Multiplexing (COFDM) is a form of modulation which is particularly well-suited to the needs of the terrestrial broadcasting channel. COFDM can cope with high levels of multipath propagation, with a wide spread of delays between the received signals. This leads to the concept of single-frequency networks in which many transmitters send the same signal on the same frequency, generating “artificial multipath”. COFDM also copes well with co-channel narrowband interference, as may be caused by the carriers of existing analogue services.

COFDM has therefore been chosen for two recent new standards for broadcasting – DAB and DVB-T, both of which have been optimized for their respective applications and have options to suit particular needs.

The special performance of COFDM in respect of multipath and interference is only achieved by a careful choice of parameters and with attention to detail in the way in which the forward error-correction coding is applied.

1. Introduction

Digital techniques have been used for many years by broadcasters in the production, distribution and storage of their programme material. They have also been used in “supporting roles” in broadcasting itself, with the introduction of Teletext and digital sound (NICAM) for television, and the Radio Data System (RDS) to accompany FM sound broadcasts. These have all used relatively conventional forms of digital modulation.

Sound and television terrestrial broadcasting is now entering a new age in which the main audio and video signals will themselves be broadcast in digital form. Systems for DAB [1] and DVB-T [2] have been standardized by ETSI for use in Europe and elsewhere in the world. These systems have been designed in recognition of the circumstances in which they will be used:



- ⇒ DAB (unlike its AM and FM predecessors) was especially designed to cope with the rigours of reception in moving cars – especially the problem of multipath reception which, in this case, is time-varying;
- ⇒ For DVB-T, a higher capacity than DAB was essential, mobile reception was not a priority, but multi-path tolerance was still important because of the widespread use of set-top TV antennas.

A new form of modulation – COFDM – was chosen for both systems, albeit with differences in detail, and with appropriate changes of parameters, to suit the different requirements of DAB and DVB-T. Both systems include a degree of flexibility.

COFDM involves modulating the data onto a large number of carriers using the FDM technique. The key features which make it work, in a manner that is so well suited to terrestrial channels, include:

- ⇒ *orthogonality* (the “O” of COFDM);
- ⇒ the addition of a *guard interval*;
- ⇒ the use of *error coding* (the “C” of COFDM), *interleaving* and *channel-state information* (CSI).

This article sets out to explain these features of COFDM, and their significance.

2. Why use multiple carriers?

The use of multiple carriers follows from the presence of significant levels of multipath.

Suppose we modulate a carrier with digital information. During each symbol, we transmit the carrier with a particular phase and amplitude which is chosen from the constellation in use. Each symbol conveys a number of bits of information, equal to the logarithm (to the base 2) of the number of different states in the constellation.

Now imagine that this signal is received via two paths, with a relative delay between them. Taking transmitted symbol n as an example, the receiver will attempt to demodulate the data that was sent in this symbol by examining all the received information relating to symbol n – both the directly-received information and the delayed information.

When the relative delay is *more* than one symbol period – see *Fig. 1* (left) – the signal received via the second path acts purely as interference, since it only carries information belonging to a previous symbol or symbols. Such *inter-symbol interference* (ISI) implies that only very small levels of the delayed signal can be tolerated (the exact level depending on the constellation in use and the acceptable loss of noise margin).

When the relative delay is *less* than one symbol period, *Fig. 1* (right), part of the signal received via the second path acts purely as interference, since it only carries information belonging to the previous symbol. The rest of it carries the information from the wanted symbol – but may add constructively or destructively to the main-path information.

This tells us that, if we are to cope with any appreciable level of delayed signals, the symbol rate must be reduced sufficiently so that the total delay spread (between the first- and last-received paths) is only a modest fraction of the symbol period. The information that can be carried by a single carrier is thus limited in the presence of multipath. If one carrier cannot



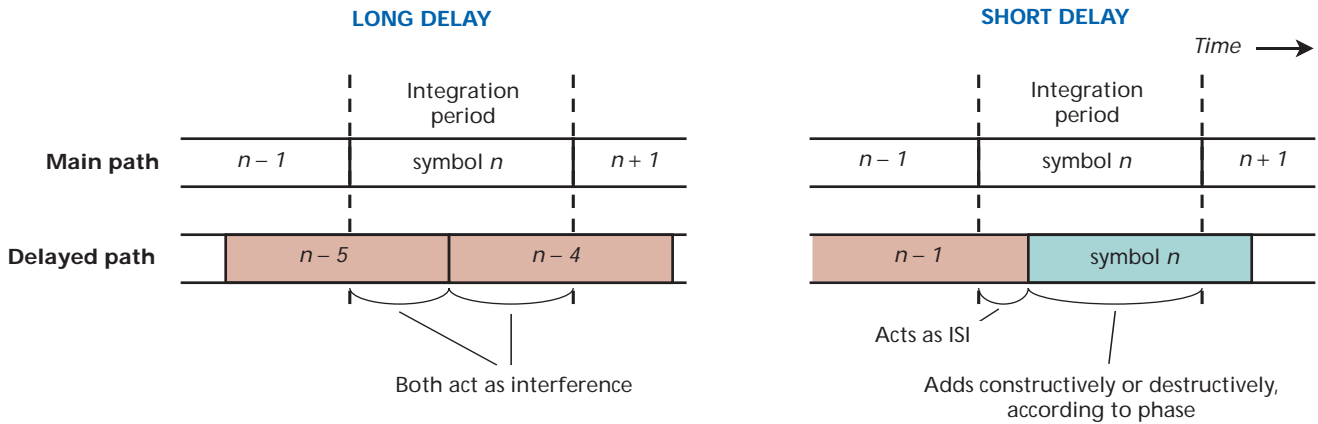


Figure 1
How a delayed path causes inter-symbol interference, with two examples of delay.

then carry the information rate we require, this leads naturally to the idea of dividing the high-rate data into many low-rate parallel streams, each conveyed by its own carrier - of which there are a large number. This is a form of FDM - the first step towards COFDM.

Even when the delay spread is less than one symbol period, a degree of ISI from the previous symbol remains. This could be eliminated if the period for which each symbol is transmitted were made longer than the period over which the receiver integrates the signal - a first indication that adding a *guard interval* may be a good thing. (We shall return to this idea shortly.)

3. Orthogonality and the use of the DFT/FFT

3.1. Orthogonality

The use of a very large number of carriers is a prospect which is practically daunting: *surely, we would need many modulators/demodulators and filters to accompany them?* It would also appear that an increase of bandwidth would be required to accommodate them. Both these worries can fortunately be dispelled if we do one simple thing: we specify that the carriers are evenly spaced by precisely $f_u = 1/T_u$, where T_u is the period (the “useful” or “active” symbol period) over which the receiver integrates the demodulated signal. When we do this, the carriers form what mathematicians call an *orthogonal set*:

The k th carrier (at baseband) can be written as:

$$\psi_k(t) = e^{jk\omega_u t}$$

where $\omega_u = 2\pi/T_u$, and the orthogonality condition that the carriers satisfy is:

$$\int_{\tau}^{\tau+T_u} \psi_k(t)\psi_l^*(t)dt = 0, \quad k \neq l$$

$$= T_u, \quad k = l$$



More intuitively, what this represents is the common procedure of demodulating a carrier by means of multiplying it by a carrier¹ of the *same* frequency (“beating it down to zero frequency”) and then integrating the result. Any other carriers will give rise to “beat tones” which are at integer multiples of ω_u . All of these unwanted “beat tones” therefore have an integer number of cycles during the integration period T_u , and thus integrate to zero.

Hence, without any “explicit” filter-ing², we can separately demodulate all the carriers without any mutual cross-talk, just by our particular choice for the carrier spacing. Furthermore, we have not wasted any spectrum either. The carriers are closely packed so that they occupy the same spectrum in total as would a single carrier – if modulated with all the data and subject to ideal sharp-cut filtering.

3.2. Preserving the orthogonality

In practice, our carriers are modulated by complex numbers which change from symbol to symbol. If the integration period spans two symbols (as for the delayed paths in Fig. 1), not only will there be same-carrier ISI, but in addition there will be *inter-carrier interference* (ICI) as well. This happens because the beat tones from other carriers may no longer integrate to zero if they change in phase and/or amplitude during the period. We avoid this by adding a guard interval, which ensures that all the information integrated comes from the same symbol and appears constant during it.

Fig. 2 shows this addition of a guard interval. The symbol period is extended so it exceeds the receiver integration period T_u . Since all the carriers are cyclic within T_u , so too is the whole modulated signal. Thus the segment added at the beginning of the symbol to form the guard interval is identical to the segment of the same length at the end of the symbol. As long as the delay of any path with respect to the main (shortest) path is less than the guard interval, all the signal components within the integration period come from the same symbol and the orthogonality criterion is satisfied. ICI and ISI will only occur when the relative delay exceeds the guard interval.




Figure 2
The addition of a guard interval.

The guard interval length is chosen to match the level of multipath expected. It should not form too large a fraction of T_u , otherwise too much data capacity (and spectral efficiency) will be sacrificed. DAB uses a guard interval of approximately³ $T_u/4$; DVB-T has more options, of which $T_u/4$ is the largest. To tolerate very long delays (as in the “artificial multip-

1. Actually a complex conjugate, corresponding to the standard I-Q quadrature demodulation process.
 2. In fact the “integrate-and-dump” process can itself be shown to be equivalent to a filter with a $\text{sinc}(\omega/\omega_u)$ characteristic, with nulls on all the carriers except the wanted one.
 3. Actually, it is precisely $63T_u/256 \approx 0.246T_u$.



ath" of an SFN), T_u must therefore be made large, implying a large number of carriers – from hundreds to thousands.

The paths of Fig. 2 may still add constructively or destructively. In fact it is possible to show that the signal demodulated from a particular carrier is very similar to the transmitted signal, but is simply multiplied by the effective frequency response of the (multipath) channel at that same carrier frequency⁴.

Many other things can cause a loss of orthogonality and hence also cause ICI. They include errors in the local-oscillator or sampling frequencies of the receiver, and phase-noise in the local oscillator [3][4]. However, in practice, the effects of these can, with care, be held within acceptable limits.

3.3. Use of FFT

We've avoided thousands of filters, thanks to orthogonality – what about implementing all the demodulating carriers, multipliers and integrators?

In practice, we work with the received signal in sampled form (sampled above the Nyquist limit, of course). The process of integration then becomes one of summation, and the whole demodulation process takes on a form which is identical to the Discrete Fourier Transform (DFT). Fortunately, efficient Fast Fourier Transform (FFT) implementations of this already exist (the integrated circuits are already available), so that we are able to build laboratory COFDM equipment reasonably easily. Common versions of the FFT operate on a group of 2^M time samples (corresponding to the samples taken in the integration period) and deliver the same number of frequency coefficients. These correspond to the data demodulated from the many carriers. In practice, because we sample above the Nyquist limit, not all of the coefficients obtained correspond to active carriers that we have used⁵.

The inverse FFT is similarly used in the transmitter to generate the OFDM signal from the input data.

4. Choice of basic modulation

In each symbol, each carrier is modulated (multiplied) by a complex number taken from a constellation set. The more states there are in the constellation, the more bits that can be conveyed by each carrier during one symbol – but the closer become the constellation points, assuming constant transmitted power. Thus there is a well-known trade-off of ruggedness versus capacity.

At the receiver, the corresponding demodulated value (the frequency coefficient from the receiver FFT) has been multiplied by an arbitrary complex number (the response of the chan-

4. For the mathematically inclined, the addition of the guard interval has in effect turned the normal process of "convolution of the signal with the impulse response of the channel" into a *circular convolution*, which corresponds to multiplication of the DFT frequency coefficients.
5. Note that this does not lead to any loss of capacity or inefficient use of bandwidth. It merely corresponds to "headroom" for the analogue filtering in the system.



nel at the carrier frequency). The constellation is thus rotated and changed in size. How can we then determine which constellation point was sent?

One simple way is to use *differential demodulation*, such as the DQPSK used in DAB. Information is carried by the change of phase from one symbol to the next. As long as the channel changes slowly enough, its response does not matter. Using such a differential (rather than a coherent) demodulation process causes some loss in thermal noise performance – but DAB is nevertheless a very rugged system.

When higher capacity is needed (as in DVB-T) there are advantages in *coherent demodulation*. In this, the response of the channel for each carrier is somehow determined, and the received constellation is appropriately *equalized* before determining which constellation point was transmitted, and hence what bits were transmitted. To do this in DVB-T, some pilot information is transmitted (so-called *scattered pilots*⁶) so that, in some symbols on

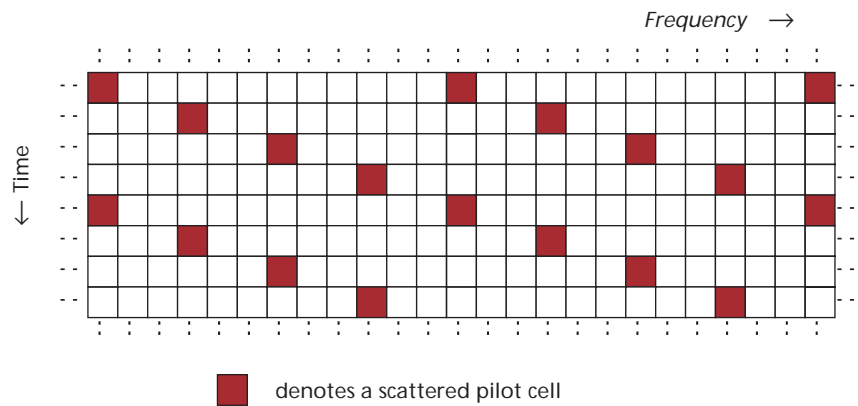


Figure 3
The pattern of inserting scattered pilot cells that is used in DVB-T.

some carriers, known information is transmitted (see Fig. 3) from which a sub-sampled⁷ version of the frequency response is measured. This is then interpolated, using a 1-D or 2-D filter, to fill in the unknown gaps, and is used to equalize all the constellations which carry data.

5. Use of error coding

5.1. Why do we need error coding?

In fact, we would expect to use forward error-correction coding in almost any practical digital communication system, in order to be able to deliver an acceptable bit-error ratio (BER) at a reasonably low signal-to-noise ratio (SNR). At a high SNR it might not be necessary – and this is also true for uncoded OFDM, but only when the channel is relatively flat. *Uncoded* OFDM does not perform very well in a selective channel. Its performance could be evaluated for any selective channel and for any modulation scheme, by:

- ⇒ noting the SNR for each carrier;
- ⇒ deducing the corresponding BER for each carrier's data;
- ⇒ obtaining the BER for the whole data signal, by averaging the BERs of all the carriers used.

Very simple examples will show the point. Clearly, if there is a 0 dB echo which is delayed such that every m^{th} carrier is completely extinguished, then the “symbol” error ratio (SER)

6. Some carriers always carry further *continual-pilot* information which is used for synchronization.

7. Sub-sampled in both frequency and time.



will be of the order of 1 in m – even at infinite SNR. (Here, “symbol” denotes the group of bits carried by one carrier within one OFDM symbol). An echo delay of say $T_u/4$ – the maximum for which a loss of orthogonality is avoided when the guard-interval fraction is 1/4 (as in DAB and some modes of DVB-T) – would thus cause the SER to be 1 in 4. Similarly, if there is one carrier, amongst N carriers in all, which is badly affected by interference, then the SER will be of the order of 1 in N , even with infinite SNR.

This tells us two things:

- ⇒ *uncoded* OFDM is not satisfactory for use in such extremely selective channels;
- ⇒ for any reasonable number of carriers, CW interference that is affecting one carrier is less of a problem than a 0 dB echo.

However, just adding hard-decision-based coding to this uncoded system is not enough, either – it would take a remarkably powerful hard-decision code to cope with an SER of 1 in 4! The solution is to use convolutional coding in conjunction with soft-decision decoding, *properly integrated* with the OFDM system.

5.2. Soft decisions and channel-state information

First let us revise, for simplicity, 2-level modulation of a single carrier: one bit is transmitted per symbol with, say, a “0” being sent by a modulating signal of -1 V and a “1” by $+1$ V. At a receiver, assuming that the gain is correct, we should expect to demodulate a signal always in the vicinity of either -1 V or $+1$ V, depending on whether a “0” or a “1” was transmitted. Any departure from the exact values ± 1 V would have been caused by the inevitable noise added during transmission.

A *hard-decision* receiver would operate according to the rule that negative signals should be decoded as “0” and positive ones as “1”, with 0 V being the *decision boundary*. If the instantaneous amplitude of the noise were never to exceed ± 1 V, then this simple receiver would make no mistakes. But noise may occasionally have a large amplitude, although with lower probability than for smaller values. Thus if say $+0.5$ V is received, it most probably means that a “1” was transmitted, but there is a smaller yet still finite probability that actually “0” was sent. Common sense suggests that when a large-amplitude signal is received we can be more confident in the hard decision, than if the amplitude is small.

This view of a degree of confidence is exploited in *soft-decision* Viterbi decoders. These maintain a history of many possible transmitted sequences, building up a view of their relative likelihoods and finally selecting the value “0” or “1” for each bit, according to which has the *maximum likelihood*. For convenience, a Viterbi decoder *adds* logarithmic likelihoods (rather than *multiplying* probabilities) to accumulate the likelihood of each possible sequence. It can be shown that, in the case of BPSK or QPSK, the appropriate log-likelihood measure, or *metric*, of the certainty of each decision is indeed simply proportional to the distance from the decision boundary. The slope of this linear relationship itself also depends directly on the signal-to-noise ratio. Thus the Viterbi decoder is fed with a *soft decision* comprising both the hard decision (the sign of the signal) together with a measure of the amplitude of the received signal.

With other rectangular-constellation modulation systems, such as 16-QAM or 64-QAM, each axis carries more than one bit, usually with Gray coding. At the receiver, a soft decision can be made separately for each received bit. The metric functions are now more complicated than



for QPSK, being different for each bit, but the principle – the decoder exploits its knowledge of the expected reliability of each bit – still remains.

Metrics for COFDM are slightly more complicated. We start from the understanding that the soft-decision information is a measure of the confidence to be placed in the accompanying hard decision.

When data are modulated onto a single carrier in a time-invariant system, then *a priori* all data symbols suffer from the same noise power on average; the soft-decision information simply needs to take note of the random symbol-by-symbol variations that this noise causes.

When data are modulated onto the multiple COFDM carriers, the metrics become slightly more complicated as the various carriers will have different signal-to-noise ratios. For example, a carrier which falls into a notch in the frequency response will comprise mostly noise; one in a peak will suffer much less. Thus, in addition to the symbol-by-symbol variations, there is another factor to take account of in the soft decisions: data conveyed by carriers having a high SNR are *a priori* more reliable than those conveyed by carriers having low SNR. This extra *a priori* information is usually known as *channel-state information* (CSI).

The CSI concept can be extended to embrace *interference* which affects carriers selectively.

The inclusion of channel-state information in the generation of soft decisions is the key to the unique performance of COFDM in the presence of frequency-selective fading and interference.

We now return to the simple example in which there is a 0 dB echo, of such a delay (and phase relationship) as to cause a complete null on 1 carrier in every 4. Fig. 4 illustrates the effect of this selective channel: 1 carrier in every 4 is nulled out, while another carrier in every 4 is actually boosted, and the remaining two are unaffected. Note that received power is shown, to which the SNRs of the carriers will be proportional if the receiver noise is itself flat, as is usual. The “mean power” marked on the diagram is the mean of all carriers. It is equal to the total received power (via both paths), shared equally between all carriers.




Figure 4
The effect of a channel with a single 0 dB echo of long delay, such that exactly 1 carrier in every 4 is nulled out.

Although only a few COFDM carriers are illustrated, the pattern repeats cyclically for all of them. The dotted curve represents the power frequency response of the channel formed by the two paths.

In COFDM, the Viterbi metrics for each bit should be weighted according to the SNR of the carrier by which it travelled. Clearly, the bits from the nulled carriers are effectively flagged as



having “no confidence”. This is essentially the same thing as an *erasure* – the Viterbi decoder in effect just records that it has no information about these bits.

There is another well-known case of regularly-occurring erasures, namely *punctured codes*. Typically, convolutional codes intrinsically have code rates expressed as simple fractions such as $1/2$ or $1/3$. When a code having a higher rate (less redundancy) is needed, then one of these lower-rate “mother” codes is *punctured*, that is to say certain of the coded bits are just not transmitted, according to a regular pattern known to the receiver. At the receiver “dummy bits” are re-inserted to replace the omitted ones, but are marked as erasures – bits having zero confidence – so that the Viterbi decoder treats them accordingly. Punctured codes obviously are less powerful than the mother code, but there is an acceptable steady trade-off between performance and code rate, as the degree of puncturing is increased.

Suppose we take a rate- $1/2$ code and puncture it by removing 1 bit in every 4. The rate- $1/2$ code produces 2 coded bits for every 1 uncoded bit, and thus 4 coded bits for every 2 uncoded bits. If we puncture 1 in 4 of these coded bits, then we clearly finish by transmitting 3 coded bits for every 2 uncoded bits. In other words we have generated a rate- $2/3$ code. Indeed, this is exactly how the rate- $2/3$ option of DVB-T is made.

Now return to our simple COFDM example in which 1 carrier in 4 is nulled out by the channel – but the corresponding bits are effectively flagged as erasures, thanks to the application of channel-state information. 2 out of 3 of the remaining carriers are received at the same SNR as that of the overall channel, while 1 carrier is actually boosted, having an improved SNR. Suppose that rate- $1/2$ coding is used for the COFDM signal. It follows that the SNR performance of COFDM with this *selective* channel should be very slightly better (because 1 carrier in 4 is boosted) than that for a single-carrier (SC) system using the corresponding punctured rate- $2/3$ code in a *flat* channel. In other words, the effect of this very selective channel on COFDM can be directly estimated from knowledge of the behaviour of puncturing the same code when used in an SC system through a flat channel.

This explains how the penalty in the required CNR, for a COFDM system subject to 0 dB echoes, may be quite small – provided a relatively powerful convolutional code is used together with the application of channel-state information.

5.3. Interleaving

So far we have considered a very special example in order to make it easy to explain – by invoking the close analogy with the use of code puncturing. But what of the other delay values?

If the relative delay of the echo is rather shorter than we have just considered, then the notches in the channel’s frequency response will be broader, affecting many adjacent carriers. This means that the coded data we transmit should not simply be assigned to the OFDM carriers in a sequential order, since at the receiver this would cause the Viterbi soft-decision decoder to be fed with clusters of unreliable bits. This is known to cause a serious loss of performance, which we avoid by *interleaving* the coded data before assigning them to OFDM carriers at the modulator. A corresponding de-interleaver is used at the receiver before decoding. In this way, the cluster of errors occurring when adjacent carriers fail simultaneously (as when there is a broad notch in the frequency response of the channel) is broken up, enabling the Viterbi decoder to perform better.



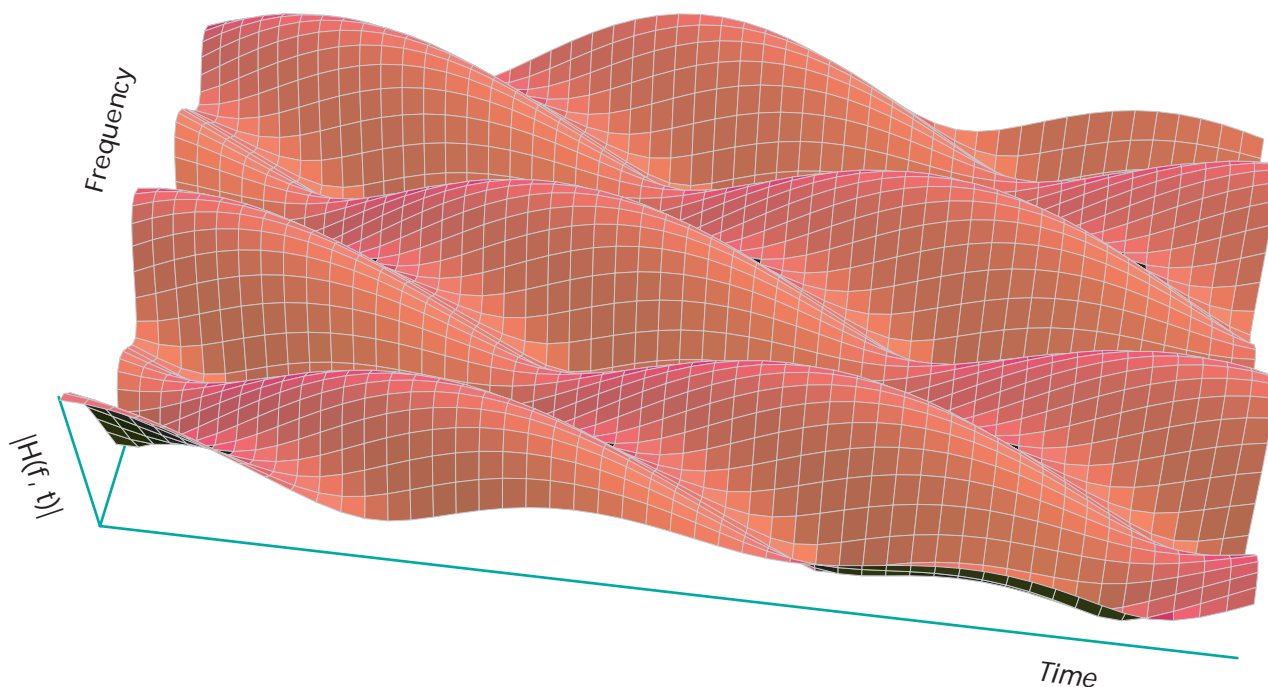


Figure 5
A time-varying channel example. There are two delayed paths, each with different Doppler shift, in addition to the “main” path. The z-axis represents the magnitude of the channel response.

As just described, the process could be called *frequency interleaving*. This is all that is needed if the channel only varies slowly with time, and that is why it is used in DVB-T. In mobile operation (a key application for DAB), we may expect the various paths to be subjected to different and significant Doppler shifts, making the frequency response vary with time (see Fig. 5). Furthermore, a vehicle may drive into shaded areas (such as underpasses) where all the signals are severely attenuated for a period (not shown in Fig. 5). For this reason, in the DAB system the coded data are also re-distributed over time, to provide *time interleaving*.

5.4. More coding

DAB conveys audio data which, despite being compressed in source coding, is relatively robust to the effects of transmission errors⁸. The BER remaining after correction by the Viterbi decoder is adequate. On the other hand, the compressed video data of DVB-T is more susceptible to errors so that the residual BER at the output of the Viterbi decoder is too high.

Thus DVB-T includes a second stage of error coding, called the “outer” coding, since in an overall block diagram it sandwiches the (“inner”) convolutional coding. Data to be transmitted are first coded with a Reed-Solomon code, interleaved with an additional “outer” interleaver, then passed to the “inner” convolutional coder. At the receiver, the Viterbi decoder is followed by an “outer” interleaver and the “outer” R-S decoder. The R-S decoder uses hard decisions, but is able to reduce the BER substantially, despite very modest extra redundancy having been added at the transmitter.

8. Some more-susceptible data have special treatment.



6. Single-frequency networks

Our simple example of a 0 dB echo often crops up when considering SFNs. If two synchronized COFDM transmitters operate on a common frequency, there will somewhere be locations where the two signals will be received at equal strength (and with a relative delay, depending on the geometry of the situation, which we assume to be within the system limits). An obvious question is: *does reception suffer or benefit from this situation?*

Clearly, compared with receiving either transmitter alone, the total received signal-to-noise power ratio (CNR) is doubled, i.e. increased by 3 dB. However, the presence of the two transmissions makes reception *selective* rather than *flat* (as we might hope to have with a single transmission, without “natural” echoes). This increases the CNR required to achieve the same BER, in a way which depends on the error-correcting code in use.

We have already seen a qualitative argument how this increase in CNR requirement may be closely related to the performance of punctured codes. Simulation shows that the increase in CNR requirement between flat and 0 dB-echo channels is just below 3 dB for a rate-1/2 code, while it is greater for higher-rate codes which have already been punctured. Practical experience supports the order of 3 dB for a rate-1/2 code while, for rate-2/3, the increase is of the order of 6 dB.

With rate-1/2 coding, when receiving two signals of equal strength (rather than either signal alone), the received CNR is increased by 3 dB: the CNR required for satisfactory reception (in the now highly-selective channel) is also increased by about the same amount. The performance is thus unchanged by adding the second path.

Abbreviations

16-QAM	16-state quadrature amplitude modulation	ETSI	European Telecommunication Standards Institute
64-QAM	64-state quadrature amplitude modulation	FDM	Frequency division multiplex
BER	Bit error rate	FFT	Fast Fourier transform
BPSK	Binary phase-shift keying	ICI	Inter-carrier interference
CNR	Carrier-to-noise ratio	ISI	Inter-symbol interference
COFDM	Coded orthogonal frequency division multiplex	NICAM	Near-instantaneous companding and multiplexing
CSI	Channel-state information	OFDM	Orthogonal frequency division multiplex
CW	Carrier wave	QPSK	Quadrature (quaternary) phase-shift keying
DAB	Digital Audio Broadcasting	R-S	Reed-Solomon
DFT	Discrete Fourier transform	RDS	Radio Data System
DQPSK	Differential quadrature (quaternary) phase-shift keying	SC	Single carrier
DVB	Digital Video Broadcasting	SER	Symbol error rate
DVB-T	DVB - Terrestrial	SFN	Single-frequency network
		SNR	Signal-to-noise ratio

Since, for most practical purposes, the case of the 0 dB echo appears to be more or less the worst one, this is very encouraging for the planning and developing of SFNs.

7. Summary of key DAB & DVB-T features

Both DAB and DVB-T have flexibility built in, to cope with a range of circumstances and uses.

DAB has four modes with 192, 384, 768 or 1536 carriers, and has corresponding guard intervals from 31 to 246 μ s. In each mode, the carriers occupy a total bandwidth of 1.536 MHz, they use DQPSK and they use both time- and frequency-interleaving.

DVB-T has two modes with either 1705 or 6817 carriers in a 7.61 MHz bandwidth, with a wide range of guard intervals from 7 to 224 μ s. Coherent demodulation is used, with QPSK / 16-QAM / 64-QAM constellations. In conjunction with options for inner-code rate, this provides extensive trade-off between ruggedness and capacity (from 5 to 31.7 Mbit/s). No time-interleaving is used. The convolutional inner code is supplemented by a Reed-Solomon outer code. (The figures quoted above relate to the use of nominally 8 MHz channels. The DVB-T specification can be adapted to 6 or 7 MHz channels by simply scaling the clock rate; the capacity and bandwidth then follow in the same proportion.)

8. Conclusions

COFDM, as used in DAB and DVB-T, is very well matched to the terrestrial channel, being able to cope with severe multipath and the presence of co-channel narrowband interference. It also makes single-frequency networks possible.

COFDM is also adaptable to various uses by making an appropriate choice of parameters. Both DAB and DVB-T have a range of options to facilitate this.

COFDM only works because all the key elements are correctly integrated. These elements include many orthogonal carriers, added guard intervals, interleaving, soft-decision Viterbi decoding and the use of channel-state information.



Jonathan Stott studied Engineering and Electrical Sciences at Churchill College, Cambridge University, graduating with Distinction in 1972. He then joined the BBC Research Department (now BBC R&D), where he is a Project Manager in the Spectrum Planning Group. Most of his career has been taken up with applying digital techniques to broadcasting.

In recent years Mr Stott has been deeply involved with the development and introduction of digital terrestrial television, starting with participation in the European RACE dTTb project. This led to his becoming a member of the Task Force on System Comparison which, under the leadership of Lis Grete Møller from Denmark, drew up the DVB-T specification for modulation and coding of digital terrestrial television. He led the theoretical and simulation work within the BBC R&D team that is at the forefront of digital television developments in Europe. He now leads a BBC R&D team which is looking at the application of digital techniques to sound broadcasting at frequencies below 30 MHz, in collaboration with the DRM consortium.

Acknowledgements

This article is an updated version of a paper which was given as part of the IEE Summer School on “New Broadcast Standards and Systems”, held in Surrey, UK, from 7 to 11 July 1997.

The author wishes to thank the many colleagues, within the BBC and collaborating organizations throughout Europe, who have helped him to develop his understanding of COFDM.

Bibliography

- [1] ETS 300 401 (1994): **Radio broadcast systems; Digital Audio Broadcasting (DAB) to mobile, portable and fixed receivers**
<http://www.etsi.fr>
- [2] ETS 300 744 (1997): **Digital broadcasting systems for television, sound and data services; framing structure, channel coding and modulation for digital terrestrial television**
<http://www.etsi.fr>
- [3] J.H. Stott: **The effects of frequency errors in OFDM**
BBC Research and Development Report No.RD 1995/15.
http://www.bbc.co.uk/rd/pubs/reports/1995_15.html
- [4] J.H. Stott: **The effects of phase noise in COFDM**
EBU Technical Review, No 276 (Summer 1998).
<http://www.bbc.co.uk/rd/pubs/papers/pdffiles/jsebu276.pdf>

The following further reading is recommended:

- [a] **Advanced digital techniques for UHF satellite sound broadcasting. Collected papers on concepts for sound broadcasting into the 21st century**
European Broadcasting Union, 1988
- [b] M.C.D. Maddocks: **An introduction to digital modulation and OFDM techniques**
BBC Research Department Report No. RD 1993/10.
- [c] J.H. Stott: **The DVB terrestrial (DVB-T) specification and its implementation in a practical modem**
Proceedings of 1996 International Broadcasting Convention, IEE Conference Publication No. 428, pp. 255 to 260.
- [d] A. Oliphant, R.P. Marsden, R.H.M. Poole and N.E. Tanton: **The design of a network for digital terrestrial TV trials**
Proceedings of 1996 International Broadcasting Convention, IEE Conference Publication No. 428, pp. 242 to 247.
- [e] L.G. Møller: **COFDM and the choice of parameters for DVB-T**
Proceedings of 20th International Television Symposium, Montreux.
http://www.bbc.co.uk/validate/paper_17.htm
- [f] J.H. Stott: **Explaining some of the magic of COFDM**
Proceedings of 20th International Television Symposium, Montreux, 1997.
http://www.bbc.co.uk/rd/pubs/papers/paper_15/paper_15.html
- [g] A. Oliphant: **VALIDATE – verifying the European specification for digital terrestrial TV and preparing for the launch of services**



Proceedings of 20th International Television Symposium, Montreux, 1997.

http://www.bbc.co.uk/rd/pubs/papers/paper_16/paper_16.html

- [h] A. Morello, G. Blanchietti et al.: **Performance assessment of a DVB-T television system**
Proceedings of 20th International Television Symposium, Montreux, 1997.
 - [i] J. Mitchell and P. Sadot: **The development of a digital terrestrial front end**
Proceedings of 1997 International Broadcasting Convention, IEE Conference Publication No. 447, pp. 519-524.
http://www.bbc.co.uk/rd/pubs/papers/paper_12/paper_12.html
 - [j] C.R. Nokes, I.R. Pullen and J.E. Salter: **Evaluation of a DVB-T compliant digital terrestrial television system**
Proceedings of 1997 International Broadcasting Convention, IEE Conference Publication No. 447, pp. 331-336.
http://www.bbc.co.uk/rd/pubs/papers/paper_08/paper_08.html
 - [k] A. Oliphant: **VALIDATE – a virtual laboratory to accelerate the launch of digital terrestrial television**
ECMAST Conference, May 1998, Berlin, Germany.
<http://www.bbc.co.uk/rd/pubs/papers/ecmast22/ecmast22.html>
 - [l] C.R. Nokes: **Results of tests with domestic receiver ICs for DVB-T**
Proceedings of 1998 International Broadcasting Convention, pp. 294-299.
-