Broadcast Radio Transmission Systems

A Brief Guide for the Non-Technical Reader

Office of the Adjudicator –
Broadcast Transmission Services

February 2014

Introduction

Late last year (2013) Department of Culture, Media and Sport (DCMS) Minister Ed Vaizey announced that his Department would not be setting a firm date for the migration of UK radio services to a wholly digital platform. The reasoning behind this decision has been well rehearsed and will not be repeated here, save to say that among the factors were that the penetration of digital listening is unlikely to meet the government's target threshold in their original timescale, coverage issues related to network roll-out, and the relative scarcity of digital radio sets in cars.

All these issues can of course be mitigated in time, but not, it is felt, by the original target date of 2015. What this means for broadcasters is that the three major radio broadcast platforms (analogue AM and FM and DAB digital radio) are likely to co-exist for a number of years, and that, unlike terrestrial television, there will be no wholesale switch-over to a digital terrestrial radio platform in the short, and arguably, medium term.

The Office of the Adjudicator, Broadcast Transmission Services (OTA-BTS) feels that now would be an appropriate time to take stock of the three broadcast platforms in terms of how their engineering functions, and the pros and cons of each platform (because no single system is perfect). At no point is it our Office's aim to offer comment on the strategic, policy or regulatory arguments behind these affairs. That is simply not our job. Rather, the intention is to offer an overview of the broadcast engineering systems that have served radio broadcasters and listeners so well over the years, aimed at a lay, non-technical readership, and written in a readable and hopefully instructive style. We hope you enjoy the document and find it useful.

The Longevity of Broadcast Transmission Systems

If you walk into any electronics store and buy a new smart-phone or tablet computer, you can be pretty sure that in a few months time there will be a newer, smarter, faster and more glamorous item available. We live in a technological world where the only constant is change at an ever-accelerating rate. Unless you decide to build a log cabin in the woods and live off berries, there is no escape from this relentless march of progress, and most people would agree that the dazzling array of devices and services on offer is on balance a good thing.

But the world of broadcast engineering isn't like that. Broadcast transmission standards persist for many years, decades even. For example, the analogue PAL (Phase Alternation Line) colour TV standard was first brought into use in the UK and Europe in the 1960s, and was finally only consigned to history in 2013. That represents nearly forty years of adherence to the same basic technology. Such a situation in the world of computing or telecommunications would be unthinkable, yet in broadcasting there are very good reasons for why technical standards last such a very long time. These are detailed below.

The most important one is that broadcasters have an obligation to maintain universality of service over the years. In the BBC's case, this is written into the Corporation's Charter, and in the case of commercial broadcasters it is in the interest of advertisers and shareholders to maximise and maintain audience levels.

If a new technical standard were to be introduced that was not receivable on existing receiver sets without careful regard to this principle, then thousands, probably millions, of listeners and viewers would be left high and dry with redundant radios and televisions. A good historical example of this may be found in the transition of television from the original 405 line black and white system to the new 625 line colour system in the 1960s. Because there were many thousands of old 405 sets still in use by viewers who could not afford, or saw no need to, upgrade to a new colour set, the old 405 line system was broadcast in tandem with the new 625 line standard, to the point where eventually it was cheaper to offer the handful of 405 line viewers that remained a free colour television than it was to keep the old 405 line transmitters going.

A more recent example is the Digital Switch Over (DSO) project from analogue television to its new digital terrestrial platform. Arguably the biggest broadcast engineering project in history, the switch-over was accompanied by an extensive programme of marketing and education, free help lines and in some cases free installation services. All designed to set peoples' minds at rest and inform them about the benefits of the new digital television world. The marketing campaign even showed the oldest television set in the UK, dating from the 1930s, happily displaying the new digital signals. Via some very clever electronics of course, but the message was, "There is nothing to fear."

"I'm sorry, that version is no longer supported" is not something you will ever hear in the world of broadcast engineering!

Of course, nothing stands still forever, and over the years many improvements and enhancements have been introduced to radio and television broadcasts. But, adhering to the principle that legacy consumers must *never* be disadvantaged, without exception these enhancements were shoehorned into the existing standard. Consumers who wished to take advantage of innovations such as Teletext, NICAM (Near Instantaneous Companded Audio Multiplex) stereo sound and the Radio Data System (RDS) could, if they wished, buy a new set and enjoy the new features; those with older receivers were simply unaware that these new services were even there.

So broadcast transmission systems have a very long shelf life indeed. Let us now turn our attention to radio and take a look our old friends AM and FM, and new kid on the block, DAB (Digital Audio Broadcasting)

Carrier Waves and Modulation

Without exception, every broadcast signal consists of two basic components: a radio frequency carrier wave, on to which is piggy backed the essential programme content, in radio of course, audio. The process by which the carrier wave (or waves) is altered in sympathy with the wanted audio is known by engineers as ***modulation***. For the purposes of this document, the terms modulation and programme audio are interchangeable.

The differences between the three existing domestic radio broadcasting standards are essentially how the carrier wave is modulated, and how the receiver decodes the broadcast signal, throwing away the parts it doesn't need and leaving only the audio to come out of the loudspeaker. Once the carrier wave has done its job it is no longer needed, and while this may be regarded as a thankless job, it is a fundamentally important one.

So let’s take a look at a carrier wave. Figure 1 shows a typical carrier wave.

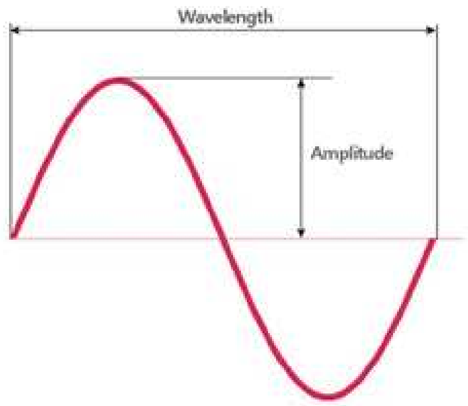

Figure 1: Basic form of a radio frequency carrier wave

Readers who were paying attention in the school maths room may recognise this as the mathematical function y = sinx, but if you don’t it’s no matter, and that is the only piece of mathematics in this document!

Better perhaps to regard it as an endlessly repeating succession of peaks and troughs – like any wave really. Engineers call the peaks “positive going excursions” and the troughs “negative going excursions”.

And now it is time to introduce the Holy Trinity of radio frequency sine waves: frequency, wavelength and amplitude.

Frequency

As can be seen from Figure 1, the carrier wave is an endlessly repeating succession of peaks and troughs. Imagine Figure 1 repeating itself on and on to the right of your computer screen. The pattern of rising from zero to a positive going peak, back down to zero and then the same thing again but in the negative direction is known as a ***cycle***. The number of cycles repeating per second is known as the wave’s ***frequency*.**

In the old days the unit of frequency was “cycles per second” (c/s) but this has long been superseded by the System Internationale scientific unit of the Hertz (Hz) named after the scientist Heinrich Hertz, an early radio pioneer. A couple of other things before we are done with frequency. You will see the Hz unit prefixed by letters such as k (short for kilo, meaning one thousand) and M (meaning mega, or one million). So, 100 kHz is one hundred thousand cycles per second, 100 MHz is one hundred million cycles per second. (Incidentally, if you ever bump into an engineer who talks about “kilocycles” and “Megacycles”, you can be sure you are talking to an old timer!)

Wavelength

Directly related to a wave's frequency is its ***wavelength*****.** The link between the two is a scientific constant called c, which is the speed of light, but let's not get too hung up on that (but yes, it is the same c as in Einstein's famous equation $e = mc^2$).

The wavelength is simply the distance in metres one cycle will travel in the course of its full positive and negative going excursions. (Think of the distance between the peaks of waves rolling onto a beach. That's wavelength).

So a radio frequency wave can be described in terms of its frequency or its wavelength, both properties being the equivalent of each other. In fact, in the glory days of medium wave broadcasting, stations often branded themselves by wavelength, mainly because radio dials were marked in metres not kHz. So, "the big 290" referred to a wavelength of 290 m, which is equivalent to a frequency of 1035 kHz. Same thing.

And here's something to commit to memory. The *higher* the frequency, the *shorter* the wavelength, and *vice versa*.

Amplitude

The final member of the Trinity is ***amplitude*****.** This simply refers to the "bigness" of the wave, or its size. It can be measured from the highest value of the positive peak to the lowest value of the negative-going trough, which is called, with great originality, "the peak to peak" value. Or, in some cases, we use a value called the "root mean square" (rms) value, which is a measure of the equivalent work an alternating current can do compared to a direct current. Huh? Don't worry about it, except that if you read an rms power in the spec of a hi-fi amplifier or set of loudspeakers, that's what they're talking about.

For now, just take on board that amplitude is the "A" in AM. And guess what, that's where we're going now.

Amplitude Modulation

Amplitude Modulation (AM) is the oldest of our broadcast transmission systems, with a pedigree stretching back to the earliest days of sound broadcasting in the 1920s. Figure 2 shows how it works. The first line on the graphic shows an unmodulated carrier wave, repeating itself selflessly as the poor old pack-horses of radio frequency carriers do. The problem is that no-one would hear it unless some programme audio is introduced. The middle line represents just such a modulation signal, which varies the amplitude of the carrier in sympathy with itself, to produce a transmitted signal as shown in the bottom line of the illustration. And that's the signal that gets broadcast.

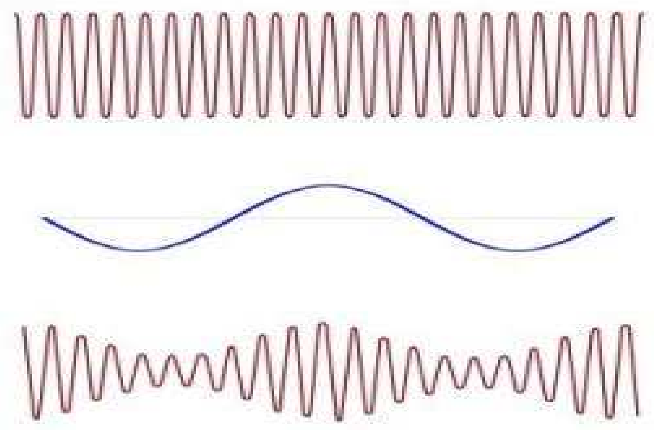

Figure 2: An amplitude modulated signal

Now, that's all well and good, but we need to have a way of listeners decoding this transmitted signal so they can hear the programme audio. We do this by a process called rectification (or detection – the two terms are interchangeable). The signal is passed through a one-way system that allows either the peaks or troughs to pass through (it doesn't matter which, the waveform is symmetrical), leaving us with Figure 3.

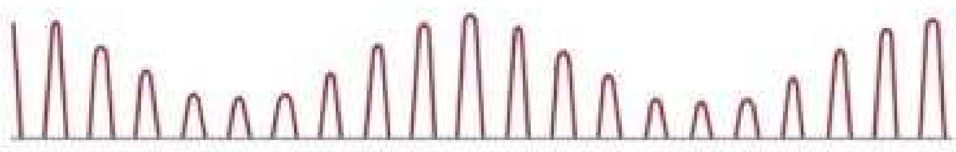

Figure 3: Detected (rectified) AM Signal

In the early days of AM broadcasting, this process was carried out by a crystal and a cat's whisker (hence crystal sets), a primitive precursor of the modern transistor, and until the invention of that device, thermionic valves. But valve is the key word: a device that allows only one-way traffic.

All that remains is to filter out the remains of the poor old carrier, and what remains coming out of the loudspeaker will be a faithful representation of what was said in the studio. Simple? Well, actually, no. No broadcast platform is perfect, so here's what's good and not so good about AM.

What's good about AM

Simplicity: it's a simple, well-understood technology with decades of service.

Wide area coverage: owing to the nature of Long Wave and Medium Wave frequencies employed by AM, the primary propagation mechanism is via Ground Wave, where the signal hugs the earth and therefore, to an extent, follows the curvature of the earth. This means that a small number of high-powered transmitters can provide virtually national coverage. It is no accident that the BBC established their AM site in Droitwich in the 1930s since it was the geographical heart of the country.

What's not so good about AM

Audio quality: for mathematical reasons, which we shall not go into (remember, no more maths?), AM has a restricted audio bandwidth, 3.4 kHz in the UK and Europe if you're interested, about the quality of a phone line. Fine for speech radio, but not so cool for music - like playing a CD on an expensive hi-fi down the phone.

It's mono only: theoretically AM can be used to broadcast stereo signals, and there were stations in the 70s and 80s who did just that. But they were mostly confined to territories like Canada, Mexico and Japan, but not in the US or Europe and certainly not in the UK. Receivers were scarce and relatively expensive, and in any case music stations were migrating en masse to FM. AM stereo is an interesting footnote in history, but for now and the foreseeable future, as Phil Spector once said, it's "back to mono."

Noise and interference: because AM is modulated in what engineers call the amplitude domain, any unwanted signals will also be detected. This is known as "impulse interference" and sources include electrical appliances, car ignition systems, and static atmospheric crashes. These will appear with the wanted audio as clicks, buzzes and crashes, and there is little that can be done to remove them.

Long range RF interference: although, as we have seen, the primary propagation mode for AM Long and Medium Wave broadcasts is surface wave, there is also a sky wave component, where the signal will reach upwards through the earth's atmosphere. Eventually it will meet the various layers of the ionosphere, a succession of belts of charged particles which have the property of reflecting radio waves back to earth. For AM broadcast signals this is not normally a problem during daylight hours, since the ionosphere will absorb or merely weakly reflect them, but after dark the reflections become much stronger, to the extent that continental transmissions can be heard on top of UK ones (and vice versa). The characteristics of the ionosphere's behaviour are well understood and predictable, and broadcasters can help each other out by mutually turning down their transmitter power at night. But it's not a complete fix, and long distance interference will always be an issue to contend with.

Power consumption: one of the advantages of AM is that very large geographic areas can be covered by a small number of high powered transmitters. The flip-side is, of course, that the transmitters will consume large amounts of electricity. The BBC's main Long Wave transmitter at Droitwich is capable of producing 500 kW (500,000 Watts) of power. To put this into context, that's equivalent to 5,000 100 W light bulbs burning continuously day and night. That's one heck of an electricity bill!

Real estate: almost without exception, AM transmitter sites use the actual mast (i.e. the bit you can see) as the radiating element of the antenna system. But the mast is "driven" against a counterpoise earth system, which you can't see because it is buried underground. The "earth mat" is usually a lattice of copper wires buried in the earth and can take up a considerable area of land. What's more, when AM is eventually switched off (as it surely will be), the land lease terms normally specify that the site be restored to its original condition as a turnip field or whatever, and these decommissioning costs will have to be borne by someone.

Frequency Modulation (FM)

FM was developed to address the shortcomings of AM, in terms of audio quality, noise and interference, and stereo capability. FM uses VHF frequencies, and the propagation mode is "line of sight", that is, the receiver has to be able to "see" (in radio frequency terms) the nearest transmitter.

The basic modulation concept for FM is shown at Figure 4 below.

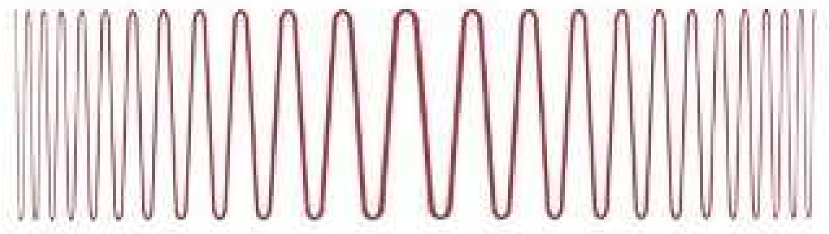

Figure 4: Basic FM modulation

It can readily be seen that although the *amplitude* of the carrier remains constant, its *frequency* varies in sympathy with the applied audio. This is manifested by the carrier waves being "squashed" or "stretched" depending on the nature of the programme content. The technical terms for this effect are compression and rarefaction. The receiver is designed to detect variations in frequency only, and is not concerned at all (within reason of course) to what is happening to the amplitude. This means that amplitude noise and interference artefacts that so bedevil AM will not be decoded at all, producing a noise and hiss free output. Remember that Steely Dan song "FM: No Static At All"? Well, that's what they were singing about.

The VHF spectrum used for FM broadcasts (Band II in the UK, 87.5-108 MHz) is more spacious in spectrum terms than Long and Medium wave, which means, as it were, that FM signals have more room to breathe, and can therefore deliver a much wider audio bandwidth than AM. In fact, the bandwidth of FM broadcasts extends from 30 Hz – 15 kHz, which is, to all intents and purposes, high fidelity.

After some years of development and appraisal, the Zenith-GE FM standard was chosen as the preferred system, and remains in universal use today. The signal that is broadcast is a composite one, consisting of a monophonic signal (Left + Right) which will be readily decoded by a mono-only radio. In addition, a *difference* signal (Left – Right) is also broadcast. A stereo radio will add the sum and difference signals to recover the Left channel and subtract the difference signal from the sum to recover the Right channel.

The calculations are:

$$(L+R) + (L-R) = 2L$$
$$(L+R) - (L-R) = 2R$$

The structure of the composite FM signal is shown in Figure 5.

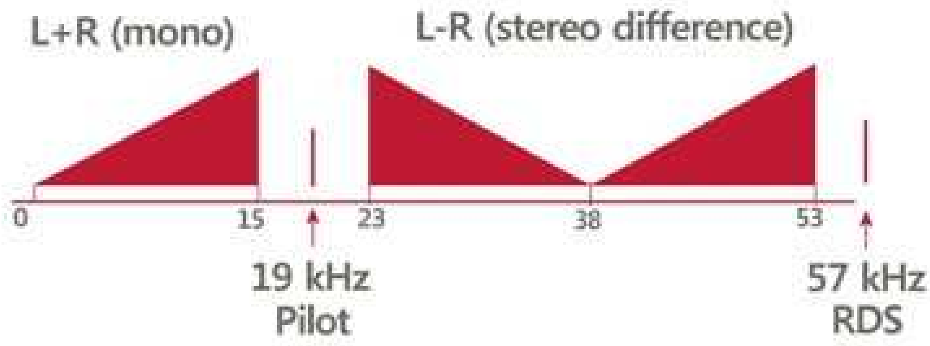

Figure 5: FM composite signal (horizontal scale: kHz)

Moving from left to right on this diagram, we can see first of all the L+R (mono) signal, which will de detected by mono receivers. Next we come across a pilot signal at 19 kHz. The purpose of the pilot signal, as its name implies, is to act as a kind of watch-keeper or reference signal to ensure that the finally decoded stereo signal will have the correct relationship with its constituent parts. Twice 19 kHz is 38 kHz (engineers call this the second harmonic) and on to this sub-carrier the L-R (stereo difference) signal is modulated. Finally the third harmonic of the pilot signal at 57 kHz is where the mysteries of the Radio Data System (RDS) lurk.

The final trick is to use ***pre-emphasis*** and ***de-emphasis***. Because noise is more noticeable in the higher frequencies of the programme audio, these frequencies are given a boost (pre-emphasis) so as to give them a fighting chance in battling the hostile environment of the transmission path. In the receiver, the reverse process occurs (de-emphasis), which restores the tonal balance of the original audio signal, but with a better signal to noise ratio than had this process not been applied.

This composite FM signal is then modulated on to the RF (Radio Frequency) carrier, and off we go.

In the early days of FM broadcasts, the expectation was that listeners would use a fixed rooftop antenna connected to a hi-fi tuner. Indeed, some audio buffs still do this, but the vast majority of FM radio is consumed via portable sets, car radios, mobile phones and so on, often in highly unfavourable environments, especially urban ones. This factor alone stretches the capability of the technology, and it is to the credit of the original system designers that FM still gives a very fair account of itself decades after the first broadcasts were made.

What's Good About FM

No static: a well designed FM transmitter and receiver can offer a listening experience virtually totally free from hiss and crackles.

Audio quality: the FM standard is able to reproduce audio frequencies up to 15 kHz, which is towards the upper threshold of human hearing, and can be considered as high fidelity. However, it should be observed that the original purist hi-fi ethos has long gone, since virtually all FM broadcasters use advanced audio processing devices such as the Orban Company's Optimod to "beef up" their signal on the FM dial and define a "tonal signature" for their station. Purists may deplore this, but it is a fact of life in this day and age. But in any case, a carefully set up audio chain is still more than capable of producing a pleasant and tuneful listening experience to all but the most "golden eared" of listeners.

Stereo: designed from the outset to offer stereo listening, while at the same time being decodable on mono-only radios.

Data capability: in the late 1980s, a system called Radio Data System (RDS) was introduced using a very low capacity (1187.5 bits per second) channel. But despite this (nowadays) laughably slow data rate, RDS is crammed with useful features such as station identification on radio displays, automatic re-tuning (if required) to local traffic announcements, automatic retuning to the strongest transmitter in a network, automatic retuning to other stations of similar format or genre, and many more possible features, some of which have been implemented by broadcasters and some not.

Graceful failure mode: if you drive away from an FM transmitter, the received signal will become less strong the further away you go. In fact, the signal diminishes by the inverse of the square of the distance travelled. So, travel two units of distance and the signal has reduced by one quarter, three and it is one ninth, four is one sixteenth and so on. Now, regulators and planners define FM service areas in terms of the radio frequency field strengths that equate to what would be an acceptable listener experience – hence the contour coverage maps with which most readers will be familiar.

So the further away from the transmitter, the weaker the signal. But what actually happens if you approach the limits of coverage? The first thing is that the stereo signal on your radio reverts to mono. Then you will hear background noise start to rise as the radio struggles to decode the increasingly weak signal. But, depending on the quality of the receiver, it may still be possible to make out what is happening with the programme even in the face of considerable noise and hash – important if, for example, you happen to be listening to a penalty shoot-out.

Engineers call this graceful degradation. DAB, for all its other advantages, does not fail in such a manner, as we shall see in due course.

What's not so good about FM

Multi-path Distortion

As we have seen, FM broadcasts rely on a line-of-sight path to the transmitter. This is not a problem if the receiver is in a fixed location and receiving an adequate signal. But FM on the move, most obviously in a car, can present some problems when the receiver antenna receives not only the direct signal from the transmitter but also one or more reflected signals bounced from buildings and other terrain features. This is illustrated in Figure 6.

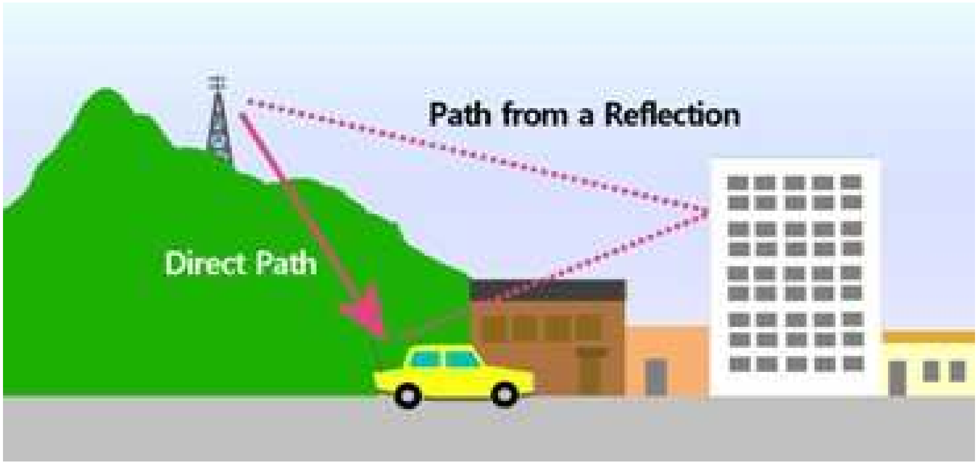

Figure 6: Multi-path distortion

The addition of the direct and reflected signals can produce some very unpleasant effects for the listener, such as audible distortion, phasing effects and in extreme cases complete cancellation of the signal. In today's increasingly high-rise cities, and consequent "urban canyons", multi-path is set to be an increasing problem for FM. (In the USA, of course, high-rise buildings are the norm, but because distances between major population centres are so great, broadcasters can simply blast out more power. As a crowded small island nation, we can't do that.)

A few words about digital systems

Before we delve into the intricacies of DAB, now would be a good time to say a few words about the basics of digital systems. Any digital system, no matter how complex, relies on one basic principle: the data that is created, stored and transmitted is made up of components that can only assume one of two states. We can express this concept in a number of ways: on/off, like a light switch; present or not present; positive or negative, and so on.

In computing, this concept has been formalised into the mathematical form of binary arithmetic, based on a number system to the base 2. We are all familiar with the decimal number system, which is numbers expressed to the base 10. This is convenient for humans since we have ten fingers and thumbs, but in reality it is possible to use *any* number as the base. We choose binary for digital systems, where a value can only be expressed as a 1 (one) or 0 (zero) because electronic switches (in the old days, electronic valves, nowadays transistors) are very good at flipping states from 0 to 1, and in performing mathematical transformations such as addition, subtraction, multiplication and division in binary.

This all happens unseen in the depths of your laptop, tablet, smart-phone (and DAB radio), where millions, even billions of tiny electronic switches race away to process data at almost unimaginable speeds. For example, the laptop computer on which this is being written, by no means a state of the art machine, processes data at the rate of 1.5 GHz. That's a clock ticking at the rate of one and a half thousand million times per second.

The basic unit of digital currency is the **binary digit**, invariably abbreviated to **bit.** A group of 8 bits is called a **byte**, and four bits (i.e. half a byte) is known as a **nibble**. Who says geeks don't have a sense of humour? These basic units are prefixed by multipliers such as Mega (one million), Giga (one billion), Tera (one thousand billion), and in the future who knows what the limit to these enormous numbers will be?

If you think about it (and those of a nervous disposition probably shouldn't), these dizzying numbers are nothing short of miraculous. But the key thing to remember is that no matter how large the data may be, it is, at its most basic level, made up of billions of electronic switches doing just one thing, uncomplainingly and forever.

The challenge of digital broadcast systems is to marry this digital technology to the familiar world of carrier waves and relay the digitised data through the airwaves so that it may be decoded in the receiver to sound and pictures. In radio, the platform adopted in the UK (and many other territories) is Digital Audio Broadcasting, DAB.

Digital Audio Broadcasting (DAB)

DAB was designed as a new radio platform for the digital age. Among its original design goals were highly robust mobile and portable performance, freedom from multi-path effects, high quality digital audio, spectrum efficiency via use of single frequency networks, more services for listeners to choose from, the ability to tailor capacity to programme format, and advanced data capabilities.

The system was developed by a European consortium of broadcasters and manufacturers under the banner of the Eureka 147 Project. The final system design was published in the mid 1990s as a European Telecommunications Standards Institute document, ETSI 300 401. This is a highly detailed and complex publication, and not for the faint-hearted. Even experienced broadcast engineers have been known to reach for the ice-pack after a few chapters.

As a consequence, this section can only hope to provide an overview of how DAB works – but we hope you will find it enlightening and useful.

How many carriers?

If readers take only one thing away from this section, it is that unlike AM and FM, which use only a single main carrier, DAB uses a multiplicity of them. The carriers are arranged like the teeth of a comb, as shown in Figure 7.

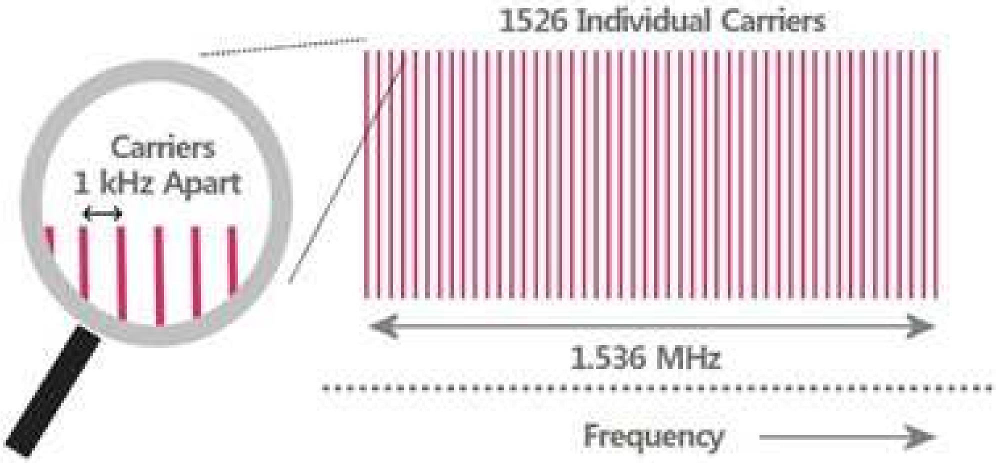

Figure 7: Arrangement of DAB multiple carriers

This is so unlike the simple arrangements of analogue systems, and may at first sight seem counter intuitive, but a system architecture like this is the only way the Eureka project could achieve its goals. Essentially the multiple services carried by a DAB multiplex or ensemble – the terms are interchangeable – are spread about this comb of carriers in such a way that if some of the carriers are lost or degraded by interference or multi-path effects, there is a very high probability that the signal can be recovered from the data on the carriers that are not affected. The basic technology that allows this to happen is called (get that ice-pack ready) **Coded Orthogonal Frequency Division Multiplexing**.

A brief COFDM primer

So, let's take this a step at a time:

Coded: this refers to the highly complex and sophisticated algorithms, that enable the very robust error protection and recovery mechanisms that protect the broadcast signal in a hostile environment. In simplistic terms, the job of the transmitter is to *encode* and the job of the receiver to *decode* them.

Orthogonal: orthogonal is a mathematical term that means "at right angles to." It can be proved via some highly complex maths that if all the carriers in the DAB "comb" exhibit electrical orthogonality, each carrier can be modulated separately without affecting any of its neighbours. This means that each separate carrier may be modulated at a relatively slow rate, but the combination of all the carriers represents a large amount of data, and hence a useful combination of radio programmes.

Frequency Division Multiplexing: don't worry, this isn't as scary as it sounds. It just means that the digital data representing the programme audio is scattered and spread across the "comb" of carriers in such a way to maximise the chances of successful recovery at the receiver after the transmitted signal has passed through a hostile radio frequency environment. If some of the carriers are degraded or even lost completely, the chances are that enough data is left on the "good" carriers for the receiver to reconstitute the signal.

DAB modulation scheme

The modulation scheme used in the DAB standard is called **Differential Quadrature Phase Shift Keying**, or DQPSK for short. (You may like to pause here for another visit to the freezer for more ice cubes).

An analogy may help here. Imagine a clock face but with only one hand, and that hand can only "tick" at the quadrants of the face. So, if the starting position was at 12 o'clock, the next tick would be at 3 o'clock, the one after that 6 o'clock, the one after that 9 o'clock and then finally back to 12 o'clock. This is illustrated in Figure 8.

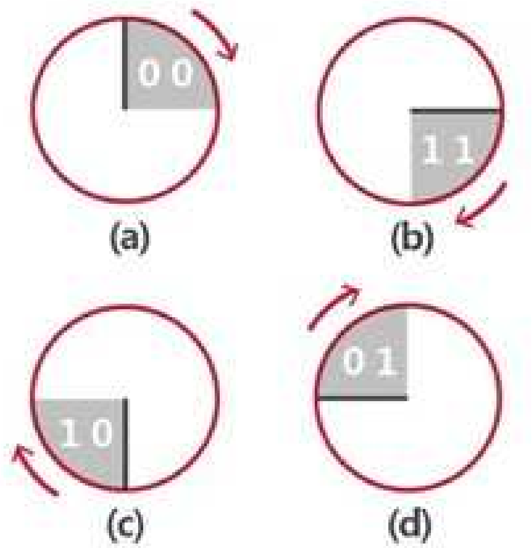

Figure 8: DQPSK data symbols

Each tick of the single-handed clock represents two bits of data, that is, the binary values 00, 01, 10 and 11. These values can be separately rendered if the carrier sine wave phase is shifted each time by exactly 90 degrees. Phase just means the horizontal position the wave adopts along the horizontal, or time, axis. A complete carrier cycle is equivalent to 360 degrees (like a circle, and indeed sine waves can be described in terms of a circle, like a clock face, but that requires a lot of maths, so let's not go there). Moving the phase of a carrier is called ***phase shift***, and it follows that a 90 degree phase shift equates to a shift of a quarter of a cycle. Two signals which are in a phase relationship of 90 degrees to each other are said to be in ***quadrature*** and that's the Q in DQPSK.

These phase shifts need to happen as quickly as possible to ensure the smooth working of the system. Again, an analogy may help. Most people now are familiar with smart-phone touch screens. Imagine that you had four photographs of the same sine wave on your phone, but each one 90 degrees out of phase with each other. If you were to flick the screen horizontally to switch backwards and forwards between the four images, you would be phase shifting the carriers in quadrature. The same as happens in DAB. Except if you tried to do it at the same rate as DAB does, you'd have a very sore thumb indeed. Fortunately, electronics are available to prevent any Repetitive Strain Injury episodes.

One last thing. The D in DQPSK stands for ***differential***. Our original clock face analogy assumed that the whole process started at exactly 12 o'clock and then worked its way in an orderly fashion round the dial. But in reality, it doesn't matter where the four quadrants start and finish, *provided that* the relationship is strictly maintained at 90 degrees. So, if for example, the initial start point happened at 23 minutes past the hour, the next phase shift has to happen at 38 minutes past, the next one at 53 minutes past and so on.

So there we have it – DQPSK in a nutshell.

Allocation of capacity

The total payload of a DAB multiplex is a function of the modulation rate on each carrier multiplied by the total number of carriers transmitted. Since each modulation "symbol" represents two bits, the sum is:

2 x number of symbols per second x 1526 carriers

If you do the maths (or get someone else to do them for you), that equates to a gross payload around 2 Megabits/second (2 Mbps). In these times of almost unlimited bandwidth, that doesn't sound a lot, but it is enough to enable several audio services per multiplex. For example, the national commercial radio multiplex, Digital One, currently carries 14 separate radio stations, a mix of music and speech services.

The DAB standard allows broadcasters to tailor the amount of multiplex "real estate" to the programme format. Thus, music services will occupy more capacity than speech services. The standard offers a range of bit-rates, which are directly related to audio quality. Typically, a pop music service (say BBC Radio 2) will be encoded at 128kbps whereas TalkSport occupies 64 kbps. There are no hard and fast rules about this, though, and some stations, such as Planet Rock, encode at 80 kbps, whereas BBC R3 uses up to 192 kbps. A variety of modes is also available: stereo and a variant called joint stereo, and of course good old fashioned mono. The choice of bit-rate is a balancing act between the total capacity available, the perceived audio quality that is deemed acceptable to listeners, and, of course, the cost of carriage.

Although some audio purists may carp at music being transmitted at the lower bit-rates (see following section), it's probably fair to say that DAB offers a flexible and realistic mechanism for maximising the number of services carried within the constraints of the total payload available.

Audio quality

Audio quality is probably the most contentious single issue to follow DAB around since the first full-time services were introduced in the mid 1990s. Early marketing material described DAB as offering "CD quality sound." While there is no evidence to suggest that this was meant to mislead, it simply isn't true. CD is based on a 16 bit linear architecture, whereas DAB uses digital compression. If it did not, it would be quite impossible to broadcast the range of services currently offered.

This is because, even nowadays, there will be an upper limit to the amount of capacity available, either on-line or on terrestrial broadcast platforms such as DAB and DTT

(Digital Terrestrial Television). Engineers had to come up with a way to counter this, and the answer was digital compression. A full exposition of this subject is well beyond the scope of this document, but suffice to say that compression takes advantage of characteristics of the human auditory and visual responses, such that certain items do not need to be coded. For example, in audio, the presence of very loud sound at the same time or nearby a much quieter sound, the human ear will not detect the quieter sound since it is to an extent “masked” by the louder one. This in turn means that it is not necessary to code both sounds equally (as a linear recording would), but that precious bits can be saved and therefore there are budget savings to be made in the transmitted bit-stream.

The audio coding system in DAB is commonly known as MPEG-2 or more accurately, MPEG1 Layer 2. (MPEG, by the way stands for Motion Picture Experts Group. Now don’t they sound like a fun bunch of guys?) At the time DAB was being developed, MP2 was state of the art, but it is showing its age now and has been overtaken by more efficient algorithms such as MP3 and MP4. We can’t just replace MP2 in DAB with a more modern coding system though, for the very good reason that it would render obsolete millions of older DAB sets that can only decode MP2 audio. The DAB+ standard does use a more modern coding system called AAC (MP4 by any other name) and many modern DAB sets can handle both coding standards. In time, DAB+ may be introduced in the UK, but not, Ofcom say, any time soon.

The fact is though, that we live in a world of digital compression. Every iTunes download, every YouTube clip, every programme watched on BBC iPlayer and similar platforms – all rely on compression. And while purist audiophiles can rightly claim that compressed audio can never be as good as CD (or, with an even more fanatical gleam in their eyes, vinyl!), the consensus among most sensible people is that the hugely increased range of content far outweighs these rather narrow concerns.

DAB data services

The DAB standard offers many opportunities to embed non-audio data in the signal. The designers deliberately provided these facilities as neutral “data hooks” rather than specify actual applications, which would be developed by ingenious people later. Various experiments in multi-media have been attempted over the years, some more successful than others. Perhaps the most common is Programme Associated Data (PAD), which can impart text data relating to what is being broadcast (song title and artist, presenter information, phone-in numbers and so on). Every DAB receiver on the market will display this information on a small screen, and most listeners find this useful.

Single Frequency Networks (SFNs)

From the outset, DAB was designed to use networks with all transmitters radiating on the same frequency. We can’t do this on FM, since neighbouring transmitters would interfere with each other to produce “mush zones”, but in DAB a clever combination of precision timing and guard intervals ensures that a receiver will be perfectly happy listening to more than one transmitter. In fact, multiple signals arriving at the receiver actively reinforce each other, leading to a benefit called “network gain”.

There are spectrum efficiency benefits accruing from this. A single DAB frequency block can enable national and large regional services without the need to source alternative frequencies. It also makes the task of network extension roll-outs much easier (subject to normal clearance and regulatory procedures) since low power infill transmitters can effectively be dropped into place.

DAB Failure Modes

Unlike FM radio, which degrades gracefully DAB degrades…well, not disgracefully exactly, but suddenly. This is because, being a digital system, in simple terms, it either works or it doesn't. Within its defined coverage area, DAB uses very robust error correction and protection mechanisms to ensure the integrity of the signal, especially in a mobile or portable environment. But as the receiver moves towards the edge of the coverage area, there is simply not enough radio frequency signal available for the receiver to make sense of the broadcast. When this happens, the signal simply stops working. The very early DAB car radios made no attempt to tidy this up, and produced horrific squeals, bangs and crashes as the mangled scale factors in the MPEG-2 coding were inverted or otherwise distorted. This was very unpleasant, as this writer's ears can testify, and modern sets simply mute the audio at the point where no useful signal is being received.

The "cliff edge" failure mode of DAB is not really a problem for national networks, since as the influence of one transmitter wanes; the radio will pick up the next one in the chain and seamlessly switch over. But for local networks the coverage boundary is much more tightly defined than with FM and the gradual deterioration we are used to in FM simply doesn't happen. This isn't really a disadvantage; it's just something listeners have to get used to.

Other digital radio systems

Before we wrap up, it is worth noting that DAB is not the only digital radio system in the world.

Digital Radio Mondiale

Digital Radio Mondiale (DRM) was developed to replace or at least augment analogue international broadcasts on Short Wave (i.e. sub-30 MHz). It uses OFDM (Orthogonal Frequency Division Multiplexing) technology to enable digital quality audio to be beamed round the world, and it does this very well, with international broadcasters such as the BBC World Service and Deutsche Welle routinely offering full time services.

As a system for domestic broadcasting though, it is more of an uneasy fit. Short Wave frequencies are prone to long range propagation effects called "skip" where signals bounce off the ionosphere, back to earth and back up again. In fact, it is this phenomenon that enables any sort of international radio communication in the first place. But long skip signals arriving in a domestic environment can cause catastrophic interference. Also, at the higher end of the Short Wave bands, towards 30MHz, propagation is affected by the 11 year sunspot cycle, with long distance signals waxing and waning with the sun's activity. While this is great fun for radio hams, it's less conducive to a stable domestic broadcast environment. That said, the BBC and others have conducted experimental broadcasts (notably in London and Devon) with reasonably promising results.

Also, DRM can be made to function at VHF frequencies (i.e. above 30 MHz), and again there has been some interesting experimental work carried out. But the main barriers to DRM finding widespread adoption as a domestic broadcast standard are likely to be economic, political and regulatory rather than technical. Put simply, DRM does what it was originally designed to do brilliantly, but is probably too late arriving at the party to be a serious contender for domestic broadcasting.

HD Radio

HD (High Definition) Radio is the digital radio standard in use in the USA. Unlike DAB, which uses a dedicated ensemble to enable multiple services, HD Radio fits the digital data "under the skirts" of the host FM analogue signal. Proponents say this enables existing FM infrastructure to be easily upgraded to digital; while critics claim that this "cuckoo in the nest" causes unacceptable interference to the analogue host and *vice versa*.

Either way, US broadcasters rejected the European DAB system (even though DAB was ready to roll long before HD Radio) because they did not like the "one size fits all" nature of DAB, where competing stations shared exactly the same coverage footprint; and one expects that the "not invented here" syndrome may have played a part too.

There are no plans to introduce HD Radio in the UK.

Satellite radio

It is of course possible to listen to radio on satellite platforms such as Sky in a fixed, sofa-based environment. Platforms such as Sky use "geostationary" satellites whose circular orbits are positioned at such a distance from the earth that they will always maintain a fixed relationship with the receiver dish as the earth revolves on its axis. This rather fanciful idea was first proposed by Arthur C Clark many decades ago: it is now a stone-cold reality.

But take satellite radio outside the living room into a car or hand-held radio, and there is no way on earth (no pun intended) that a geostationary system can deliver radio on the move. There are two fundamental reasons for this: firstly, the sheer distance of these satellites from earth (roughly 22,000 miles) requires an extremely high gain dish to recover the tiny signal beamed from space; and secondly, high gain also means highly directional, and while the dish mounted on your chimney breast can always easily "see" the satellite, a dish mounted on the roof of a car simply couldn't.

Should you want to enable mobile satellite radio, the only option is to adopt a Highly Elliptical Orbit (HEO) configuration. As its name suggest, HEO uses orbits stretched into a very tight ellipse. At the part of the ellipse nearest the earth, a mobile receiver will receive enough signal to successfully decode the signal, but only for a short time until the satellite whizzes off back out to space. So the only answer is to use a constellation of multiple satellites with overlapping orbits so that the receiver always has sight of at least one.

This of course is a fearsomely expensive venture, but it has been successfully developed in the US as the XM platform, which broadcasts multiple radio channels across the North American continent. But what works in the US won't necessarily work in Europe, for the following non-technological reasons:

Language: in the US, there are essentially two dominant languages, English and Spanish. In Europe, there are many. Providing multiple language services is expensive and bandwidth-hungry.

Copyright: in the US, the American Society of Composers Authors and Publishers (ASCAP) collect the vast majority of royalties on behalf of its members. This means that broadcasters have in essence one copyright body to deal with. In Europe, each country has its own systems and organisations, and negotiating arrangements with these multiple bodies represents a daunting task.

It is probably safe to say that satellite-delivered radio in Europe with the same capability as terrestrial systems is an unlikely prospect, especially given the rapid growth of broadband capacity to mobile devices.

Conclusions

The story of radio broadcast engineering has been one of gradual evolution, from early AM broadcasts received on crystal sets, through the introduction of FM stereo broadcasts and now embracing the digital age with its own custom-designed platform, DAB.

Radio as a medium has proved to be remarkably resilient over its history, successfully re-inventing itself to fend off challenges from talking pictures, television and more recently, the Internet. It seems people love radio: the latest audience research from industry body RAJAR (Radio Joint Audience Research) reveals that 91% of UK adults tune in to their chosen stations every week.

Listeners obviously care *what* they listen to, but are probably less concerned about the *how.* But engineering is the functional base for radio, and the choice of transmission systems is a crucially important one. Technical standards in broadcasting have very long lives: this is not some kind of Luddite resistance to change, but rather a careful and measured strategy to ensure continuity across platforms and ensure that listeners are never disadvantaged.

The three platforms described in this document – AM, FM and DAB – are in parallel use now, and it looks like for some time to come. Now is a good time to assess their characteristics and capabilities. We hope that this document has gone at least part of the way to fulfilling that.

Acknowledgements:
Graphics: Philip Bond
Layout & proofing: Sarah Peverell