DRM — key technical features

Jonathan Stott
BBC Research & Development

Digital Radio Mondiale is a system which promises to re-invigorate the use of the broadcasting bands below 30 MHz. It offers a dramatic improvement in audio quality, not only improving the audio bandwidth and signal-to-noise ratio, but also countering the effects of selective fading and audible interference from other stations. It is also designed to support various features that will make receiver operation more user-friendly.

This article describes the basic "mechanics" of DRM and its features, which include station identification, alternative frequency lists and support for time-varying frequency schedules – of particular importance in HF broadcasting.

At present, the various broadcasting bands below 30 MHz are used in much the same way as they always have been since the birth of radio broadcasting over 80 years ago. Audio is conveyed on the radio-frequency carrier in the form of *amplitude modulation* (AM) [1]. In consequence these broadcasting bands are often called the "AM bands" – distinguishing them from the very-high-frequency (VHF) broadcasting band where *frequency modulation* (FM) is used.

The AM bands are very useful to broadcasters because the modes of propagation used (ground and sky wave) give coverage to large or remote areas – the latter being invaluable for international broadcasting.

As you might suppose for a system with such early beginnings, a very simple receiver can be used. This, together with the fact that most listeners have grown up taking its existence for granted, helps to explain the very large number of AM receivers estimated to be in use: over 2 billion.

1. To be more precise, the technique used is "double-sideband AM with carrier".

However, radio listeners are becoming more sophisticated, more discerning and have more sources of entertainment and information becoming available to them all the time. Present AM radio cannot match the sound quality of FM radio, let alone CDs and other digitally stored or received formats, so listeners have every incentive to desert AM radio.

So, broadcasting in the "AM bands" must evolve to meet the new circumstances. There are very good reasons to keep these bands for broadcasting purposes because of the types of coverage that they, uniquely, can provide. What is needed is a way to improve the audio quality, and indeed the whole listening experience (ease of use, supplementary information, presenting the choices available for easy selection …) – in other words, to be more "listener-friendly".

The way to achieve this is to apply digital processing techniques, just as has been done with other bands and other media. This has been done by the efforts of many organizations within the Digital Radio Mondiale (DRM) consortium [1]. The system that they have developed has already been recognized by the International Telecommunications Union (ITU) in a draft Recommendation [2] for a world standard for digital broadcasting below 30 MHz.

This article describes some of the features and principles of operation of the DRM system. But first it is useful to recap what AM does now, and what the bands have to offer.

The "AM Bands"

Current uses of the AM bands

The AM bands serve a wide range of purposes, taking advantage of the properties of the two modes of propagation at these frequencies: *ground wave* and *sky wave*.

The lowest-frequency broadcasting band – LF, often known as the *long-wave* band (LW) – can provide extended coverage by ground wave from a single transmitter, and is thus very effective for national coverage and beyond. The long wavelength implies large transmitting-antenna structures, so this band is normally only used where a large coverage area is desired and can thereby justify the investment in the antenna. LW is popular in Europe especially, but is not available for broadcasting in the Americas.

The medium-frequency (MF) or *medium-wave* (MW) band is available throughout the world and has a wide range of uses. Ground-wave propagation is slightly less effective at this frequency than at LF, but daytime coverage is still good. One transmitter or a small network of a few transmitters can provide national coverage (depending on the size of country) although the band is also used for local coverage with a single low-power transmitter. At night (when absorption in the D layer of the ionosphere is reduced) sky-wave propagation occurs in addition to ground wave. This can be considered a hindrance or a help. The enhanced propagation can bring *interference*; e.g. a distant co-channel station unheard during the day becomes audible in the background of the wanted station. There can also be *self-interference* from the sky wave of the wanted station – this arrives later than the ground wave and thus causes interference as a form of multipath, resulting in fading or distorted sound. However, the enhanced propagation can also be exploited positively for *international broadcasting*, to places that the ground wave cannot reach.

The HF (or *short-wave*, SW) bands are mostly known for facilitating international broadcasting, using sky wave. Broadcasts can be targeted at distant countries with the advantage that there is no gatekeeper to obstruct, or other operator to charge for, the delivery of the service. Signals can even be sent halfway round the world, although most major broadcasters arrange to have transmitters somewhat nearer the target in order to improve reliability and give more choice of operating times and frequencies.

Abbreviations

AAC	(MPEG-2/4) Advanced Audio Coding
AFS	Alternative frequency switching
AM	Amplitude modulation
BER	Bit-error ratio
CD	Compact disc
COFDM	Coded orthogonal frequency division multiplex
DAB	Digital Audio Broadcasting
DC	Direct current
DRM	Digital Radio Mondiale
DVB	Digital Video Broadcasting
DVB-T	DVB - Terrestrial
FAC	Fast access channel
FFT	Fast Fourier transform
FM	Frequency modulation
HF	High-frequency
ITU	International Telecommunication Union
ITU-R	ITU - Radiocommunication Sector
LF	Low-frequency
LW	Long-wave
MF	Medium-frequency
MLC	Multi-level coding
MPEG	Moving Picture Experts Group
MSC	Main service channel
MW	Medium-wave
NVIS	Near-vertical-incidence sky-wave
OFDM	Orthogonal frequency-division multiplex
PA	Power amplifier
PM	Phase modulation
QAM	Quadrature amplitude modulation
RDS	Radio Data System
RF	Radio-frequency
SBR	Spectral band replication
SDC	Service description channel
SFN	Single-frequency network
SSB	Single side-band
SW	Short-wave
UEP	Unequal error protection

The lower-frequency end of the SW band is also used to provide national coverage of tropical or large countries (see *Appendices A* and *B*). Both applications rely on sky-wave propagation and thus depend on the state of the ionosphere, which changes daily, yearly and according to the 11-year sunspot cycle – with some random variations thrown in. For this reason, use of the HF band involves the broadcaster in changes of frequency during the day, and thus in the course of a broadcast unless broadcasting hours are short. In addition, uncertainties in forecasting, together with the fact that a large target country may require different modes of propagation in order to reach its various parts, mean that the use of two or more frequencies in parallel is common.

What are the pros and cons of the present AM bands?

Advantages

1) Coverage can be national or international, with few transmitting sites needed;
2) International broadcasting is achieved without gatekeepers;
3) Receivers are simple, cheap, readily available and work anywhere.

Disadvantages

1) Audio quality (bandwidth, noise, effects of interference and fading) is limited by modern standards [2];
2) Frequent changes of frequency (and complicated schedules) in HF broadcasting are confusing for the listener, who may find the receiver difficult to tune;
3) Poor image when compared with more modern digital technologies.

Could things be improved?

The first two advantages above are fundamental to the propagation characteristics of the bands themselves. They would therefore be retained whatever the technology – this is why these bands are so important to broadcasters. So, the challenge is to find a

2. In principle, amplitude modulation can support an adequate audio bandwidth. However, pressure on the spectrum means that the channel spacing has been restricted in order to accommodate more stations. This in turn means that the bandwidth of receivers has been reduced accordingly – in order to minimize adjacent-channel interference. In parts of the world where the spectrum is less crowded, wider transmitter and receiver bandwidths could in principle be used, thereby increasing the received audio bandwidth. The distortion caused by the selective fading that results from multipath reception can be reduced by using a synchronous detector instead of the common envelope detector. However, the audio signal-to-noise and signal-to-interference ratios remain inextricably linked to the corresponding RF ratios.

way to overcome the disadvantages without losing the essentials of the third advantage. Digital technology is the answer.

DRM – the digital alternative

What is DRM?

DRM is a consortium which brings together a wide spread of relevant experience, having members (see *Panels 1* and *2*) drawn from broadcasters and transmitter operators, manufacturers of transmitters and receivers, and research organizations. They recognized that introducing a new broadcasting system involves more than just developing "clever technology" – it must be the right technology, which solves the real problems. They therefore adopted a structure loosely based on that used by the DVB [3] project, having a Steering Board, Commercial Committee and Technical Committee. The Technical Committee is tasked with developing technical solutions to the requirements identified by the Commercial Committee.

Sub-groups of the Technical Committee propose and develop the various parts of the system. In addition, a System Evaluation sub-group tests what has been developed, both in the laboratory and the field. Field tests have the added and important function of measuring the characteristics of typical LF/MF/HF broadcasting paths. This provides a level of detail that was not generally available but is essential in order to tailor the design to match what is needed.

Discussion of some key requirements

The official User Requirements for DRM are set out in a formal document which can be found on the DRM website [1]. Let us briefly consider some of the key ones (as expressed in the author's words).

⇨ **The audio quality, as perceived by the listener, must be improved over that achieved by AM, so that DRM stations can compete with other stations using FM which may also be available to the listener**

The most obvious quality limitation of present AM is the audio bandwidth. In addition to this "intrinsic" limitation, AM quality can be impaired by limited signal-to-noise ratio, slow or fast fading, selective fading from multipath, co- and adjacent-channel interference and other interference from man-made and natural sources. Robustness against all of these is important, in addition to the "intrinsic" quality achievable under good conditions.

3. Digital Video Broadcasting. The DVB project has developed standards (which have been widely adopted) embracing all aspects of digital video broadcasting, including cable, satellite and terrestrial delivery. See http://www.dvb.org.

Panel 1
DRM Members

- Academy of Broadcasting Science of China (China)
- Antenna Hungaria (Hungary)
- Atmel ES 2 (France)
- British Broadcasting Corporation (UK)
- Broadcasting Centre Europe (Luxembourg)
- Coding Technologies Sweden AB (Sweden)
- Comatlas (France)
- Continental Electronics Corporation (USA)
- Deutsche Telekom AG (Germany)
- Deutsche Welle (Germany)
- DeutschlandRadio (Germany)
- Egyptian Radio and TV Union (Egypt)
- Europe 1 (C.E.R.T. – Germany)
- Fraunhofer Institute (Germany)
- Harris Broadcast Corporation (USA)
- International Broadcasting Bureau (USA)
- JVC Victor Company of Japan, Ltd. (Japan)
- Kymenlaakso Polytechnic (Finland)
- LSI Logic Europe (UK)
- Main Centre for Control of Broadcasting Networks, Voice of Russia (Russia)
- Merlin Communications International Ltd. (UK)
- Micronas Intermetall GmbH (Germany)
- NHK Japan Broadcasting Corporation (Japan)
- Norkring AS (Norway)
- Nozema (Netherlands)
- RadioScape Ltd. (UK)
- Radio Canada International (Canada)
- Radio France (France)
- Radio France Internationale (France)
- Radio Nederland Wereldomroep (Netherlands)
- RAI Radio Televisione Italiana (Italy)
- Radio Sweden Intl. (Sweden)
- Retevisión (Spain)
- Riz Transmitters (Croatia)
- Robert Bosch GmbH (Germany)
- Roke Manor Research Ltd (UK)
- Sangean America, Inc. (USA)
- Sony International Europe (Germany)
- SWR Südwestrundfunk (Germany)
- Technisat (Germany)
- Technology for Communications International (USA)
- TéléDiffusion de France (France)
- Telefunken Sendertechnik GmbH (Germany)
- Teracom SE (Sweden)
- Thomcast (France)
- Voice of Nigeria (Nigeria)

This list reflects the membership at the time of writing (January 2001) – please refer to the DRM website for up-to-date information.

⇨ **The DRM signal must fit within the present channelling arrangements in the AM bands, but with flexibility to permit exploitation of possible rearrangements which may take place during the life of the standard**

The changeover from AM to DRM will clearly be a slow process, with so many existing transmitters and receivers in place. At the same time, the bands are extremely crowded. So, for changeover to be possible, we must ensure that any one AM station can be converted to DRM, without upsetting or being disturbed by its existing (geographical and spectrum) AM neighbours. This requirement implies not only that the RF bandwidth fits the channelling, but also that the protection ratios (AM-DRM, DRM-AM & DRM-DRM) and the necessary carrier-to-noise ratio are compatibly related.

⇨ **It should be possible to convert existing modern transmitters to DRM operation**

There is a wide range of AM transmitter types in service, some very old indeed. A majority uses some variety of high-level modulation *(see Fig. 1)*, whereby the amplitude of the RF signal is modulated by varying what is, in effect, the power

Panel 2
DRM Associate Members

- Arab States Broadcasting Union (Tunisia)
- Asia Pacific Broadcasting Union (Malaysia)
- CCETT (France)
- Christian Vision (UK)
- Communications Authority Hungary (Hungary)
- DLM Direktorenkonferenz der Landesmedienanstalten (Germany)
- European Broadcasting Union (Switzerland)
- ESPOL (Ecuador)
- Friedrich Ebert Stiftung (Germany)
- HFCC (Czech Republic)
- Institut für Rundfunktechnik (Germany)
- ICRC International Committee of the Red Cross (Switzerland)
- International Telecommunications Union (Switzerland)
- Landesrundfunkanstalt Sachsen-Anhalt /Projektbüro Digitaler Rundfunk (Germany)
- National Association of Shortwave Broadcasters (USA)
- Radio New Zealand International (New Zealand)
- University of Applied Sciences, FH Merseburg (Germany)
- University of Hannover (Germany)
- University of Ulm (Germany)

This list reflects the membership at the time of writing (January 2001) – please refer to the DRM website for up-to-date information.

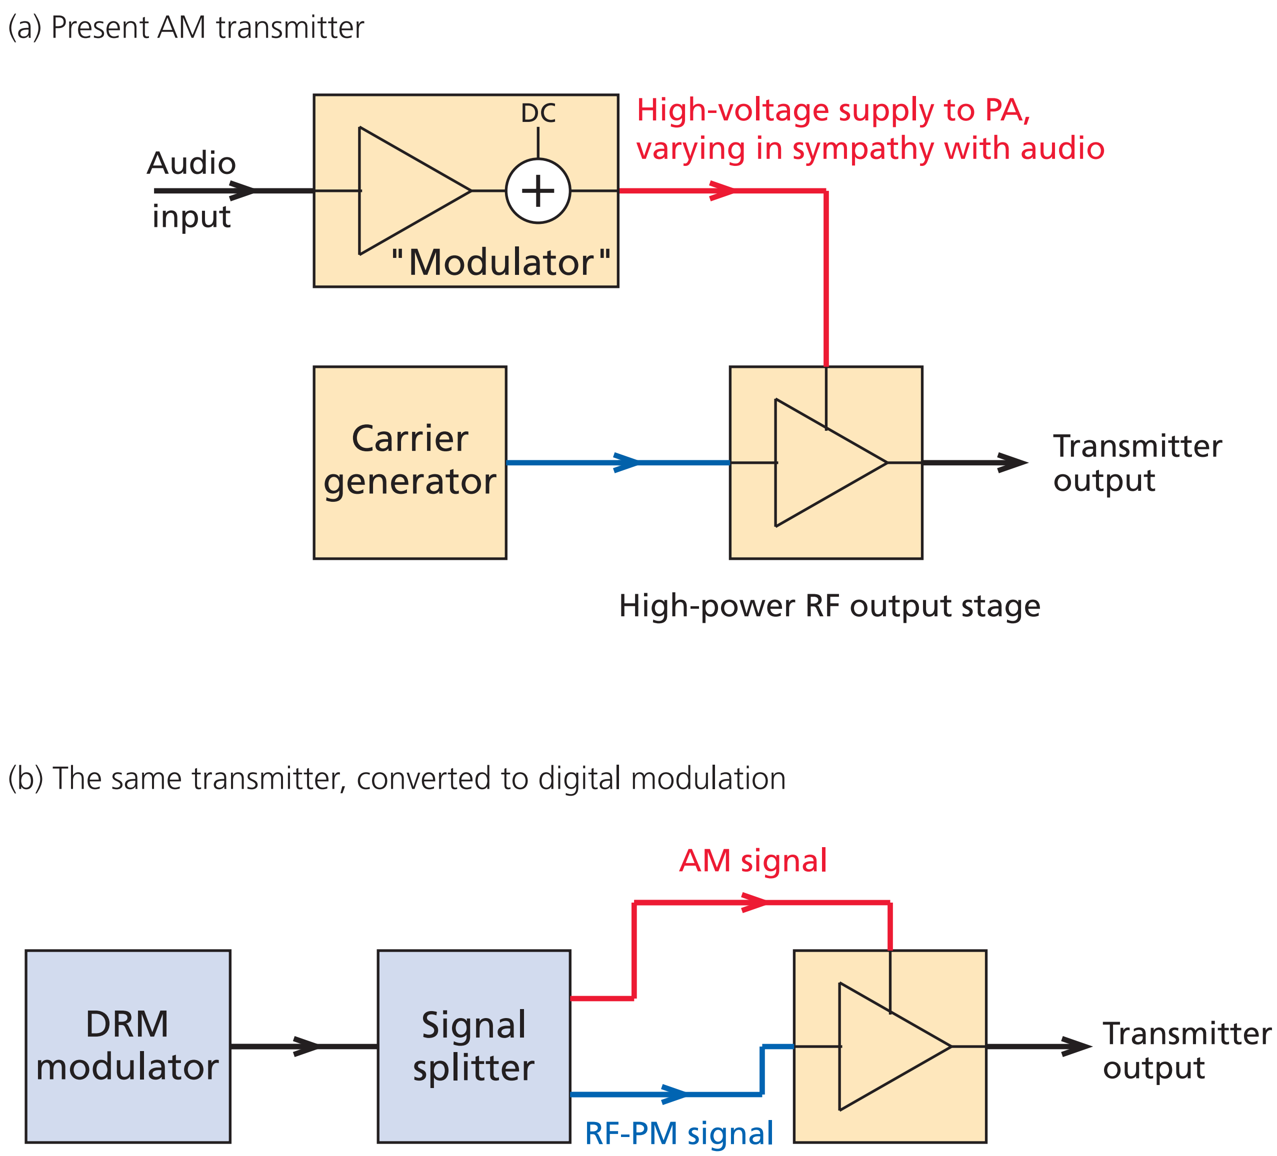

Figure 1
The use of existing AM transmitters for DRM.

supply to a non-linear RF amplifier (commonly a valve operating in Class C). A digitally-modulated signal [4] can be considered to be equivalent to a phase-modulated signal (of constant amplitude), in turn also modulated in amplitude. Such a signal can be "smuggled through" an AM transmitter structure by separating the signal into its AM and PM components. The amplitude-modulating signal simply takes the place of the previous audio modulation input, while the carrier-frequency input (normally derived from a synthesizer or other stable frequency source) is replaced by the phase-modulated RF signal. The bandwidth of the amplitude-modulated signal [5] in this arrangement is greater than that of the audio signal it was designed for. In addition, the modulator must work down to zero frequency.

4. Or indeed any signal in general, *other* than double-sideband AM, which has a symmetrical spectrum and thus a purely real baseband equivalent. A similar process is used when such transmitters are used for single-sideband (SSB) transmission.

In some cases [6] the modulator would have to be replaced to achieve this, while in others it may sufficient to modify a filter [7].

⇨ **The DRM signal should support operation of a single-frequency network (SFN)**

SFNs are sometimes used with AM. For LF/MF networks (e.g. for national coverage) they use spectrum efficiently but at the price of distorted reception in *mush areas* at the places (between transmitters) where similar strength signals are received from two of the transmitters. For this reason, planners try to arrange for the mush areas to fall in areas of low population – a serious planning constraint. Designing the DRM system to support SFNs would bring the bonus that networks which currently use two or more frequencies (because the mush-area problems of an SFN would be too great with AM) could be converted to SFN operation, in addition to existing SFNs. There could also be a benefit at HF, where SFNs are sometimes used for international broadcasting with the risk, while using AM, of much of the coverage area becoming a mush area.

⇨ **Increased user friendliness**

This term encapsulates a wide range of things to improve the listener experience. The listener chooses to listen to a particular station or service, but currently to do this with AM reception needs detailed knowledge of the band and frequency – which may change with time. The band and frequency are things of no intrinsic interest – and a DRM receiver should largely insulate the listener from having to deal with them. Suitable data transmitted along with the signal can tell the *listener* (via display, jingle, speech synthesizer …) what the station is – and tell the *receiver* how to find it again, following any scheduled changes in frequency, as well as advising what alternative frequencies may be available. It is also possible to give some information about future programmes and their timing. In effect this is like RDS [8] – updated, expanded and tailored for digital delivery in the AM bands.

5. The bandwidth of the phase-modulated signal is also greater than that of the finally radiated signal, but this seems to pose less of a problem in that most transmitters seem to be able to support it. Note that limitation in the bandwidth of either path will lead to the transmitter radiating unwanted signals outside its allocated channel bandwidth. For the same reason, delays through the two paths must also be carefully matched.
6. Especially those using a modulation transformer.
7. Modern types of modulator use some kind of switching process (in effect a sampled process) and thus require a low-pass filter. The cut-off frequency of this filter, and perhaps the switching frequency itself, has to be increased.
8. Radio Data System – the system used in Europe and beyond to send station identification, alternative frequencies and other data alongside an FM signal.
   See: **http://www.rds.org.uk/rds98/rds98.htm**

⇨ **Flexibility for the broadcaster to trade intrinsic quality or data capacity against ruggedness**

A local service provided by ground wave, which happens to have little co- and adjacent-channel interference, may need little protection against errors (and could thus support a high data rate). By way of contrast, a long-distance sky-wave service, for example, might have to cope with a large multipath delay spread, interference and a poor signal-to-noise ratio simultaneously. The broadcaster should be able to exploit the differences and be able to maximize the payload within the capabilities of the channel [9].

The system developed by DRM takes into account all of these key requirements.

The main processes

A digital sound broadcasting system comprises conceptually distinct transmission stages.

1) The audio signal must first be converted to digital form. Since the raw bit-rate that results is impracticably high, a form of bit-rate reduction tailored to the signal properties is then applied. This is usually referred to as *source coding*.
2) The source-coded data is then *multiplexed* together with any other data that forms part of the payload.
3) The multiplexed data of the payload is subjected to *channel coding* [9] to increase its ruggedness.
4) The channel-coded data is modulated onto the RF signal for transmission.

Note that source coding *reduces* the data rate, while channel coding *increases* it.

At the receiving end, the receiver first acquires synchronization with the signal, then reverses the transmission stages by means of the following processes:

1) demodulation;
2) channel decoding (correcting the transmission errors);
3) demultiplexing the transmitted data into component streams;
4) source decoding (to obtain an audio signal from the audio stream).

9. It is unfortunate that the word "channel" gets over-worked in RF engineering and this article – in following standard usage – compounds the felony. It can mean (a) the chunk of spectrum (in this case 9 or 10 kHz wide) around the nominal transmission frequency (as in "channelling"), or (b) the combination of all the impairments that the signal suffers along the path from the transmitter to the receiver (as in "fading channel", "channel coding").

In practice these processes do not separate quite so neatly in every particular case, but the concept is nevertheless useful. Indeed, a similar division was used to divide the work of DRM development between teams with relevant expertise.

Source Coding

Requirements

The capacity available for audio within a single 9 or 10 kHz channel is distinctly limited – at best in the mid-20s of kbit/s, and perhaps as little as 10 kbit/s for some extremely unfavourable examples of HF paths. This clearly represents a serious source-coding challenge for DRM, which expects to deliver good audio quality for both speech and music.

Key technologies

DRM has picked up on work already done in the development of source coding elsewhere and has fine-tuned it to the particular application. For coding most broadcast programme material, a "waveform" coder is needed in order to cope with the arbitrary mix of speech, music and incidental background sounds. DRM uses for this purpose Advanced Audio Coding (AAC), supplemented by Spectral Band Replication (SBR).

In principle, an alternative for speech-only programming is to use a coder designed expressly for speech, in which case the bit-rate can be reduced more than is possible with a waveform coder, while retaining the same speech quality. Although this could offer broadcasters further flexibility, there is however some doubt whether this approach would be used much in practice. Even notionally "speech-only" broadcast material contains jingles, background sounds in interviews and so on – all of which cause serious problems to speech coders.

How it works

Waveform coders, like AAC, work by analysing the content of each part of the audio spectrum and describing each no more accurately than is needed in order to satisfy the ear of the listener. Sounds that are masked by nearby louder sounds are discarded altogether. AAC follows in the tradition of MPEG-1 Layer 2 and MP3 in this regard, and forms part of the MPEG-4 standard.

However, even with the advances made, it is difficult to deliver an "FM-like" 15 kHz bandwidth using AAC alone at the very low bit-rates envisaged, without introducing audible artefacts. The results are better if AAC is used to deliver a more modest bandwidth – but this would mean that we fail in one of our key objectives: a noticeable increase in audio bandwidth compared with AM.

The answer lies in the combination of AAC with the SBR technique.

The SBR technique synthesizes the sounds which fall within the highest frequency octave-and-a-bit. Sounds in this range are usually either:

a) **noise-like** (sibilance, percussion instruments such as shakers, brushed cymbals etc.), or

b) **periodic** and related to what appears lower in the spectrum (overtones of instruments or voiced sounds).

At the sender, the highest-frequency band of the audio signal is examined to determine the spectral distribution and whether it falls into category (a) or (b) above. A small amount of *side information* is then prepared for transmission to help the decoder. The highest-frequency band is then removed before the remaining main band of the audio signal is passed to the AAC coder, which codes it in the conventional way.

At the receiver, the AAC decoder first decodes the main band of the audio signal. The SBR decoder then adds the synthetic upper band, helped by the instructions sent in the side information. Overtones are derived from the output of the AAC decoder, while noise-like sounds are synthesized using a noise generator with suitable spectral shaping.

The possibility of stereo operation is foreseen, although this would only be sensible where it was possible to use a double-width channel of 18 or 20 kHz RF bandwidth.

Multiplexing

With the limited bit-rate available, it is important to strike the right balance between flexibility and efficiency while protecting each bit of information to an appropriate degree. A distinction is therefore made between main payload data and the various types of data that the receiver needs to help it find and decode the desired programme. Furthermore, an option is provided for unequal error protection (UEP) of the payload data itself, so that the greatest protection is given to the data whose corruption would cause the worst impact on the audio signal.

The main payload is called the Main Service Channel (MSC). Two subsidiary channels are also provided, namely the Fast Access Channel (FAC) and the Service Description Channel (SDC). These two are key to ensuring a simplicity of operation of the receiver and are therefore designed to be reliably received in adverse conditions, with different forward error-correction from the MSC.

The FAC is intended to be decoded quickly by the receiver on first acquiring the signal (at switch-on, or during scanning). It carries a minimum of constantly-repeated data which might be essential at this stage, informing the receiver what bandwidth option is in use, what modulation is used for the SDC and MSC, which length of interleaver is used for the MSC, etc.

The SDC contains more data, also sent repeatedly but in a longer cycle so as to maintain efficiency. It contains an identification of the service(s) available in the MSC, together with further information to instruct the receiver how to decode each service. It is here that lists of alternative frequencies and frequency schedules would be transmitted if appropriate.

Finally, the bulk of the signal conveys the MSC. With the limited bit-rate available within one 9 or 10 kHz channel, this would normally be used to carry one audio programme, together with a modicum of data. Nevertheless, there is a degree of flexibility, so the MSC may contain between one and four streams of data. *Streams* and *services* are distinguished as follows. An audio service consists of one stream carrying audio, and optionally one stream carrying data. A data service consists of one stream carrying data.

Channel coding and modulation

Key technologies

Once sky-wave propagation or an SFN is involved, there is a clear need to be able to cope with multipath propagation. The delay spread can be as much as many milliseconds. This is long in comparison with (1/bandwidth), which makes DRM another application, following DAB and DVB-T, for which the multi-carrier modulation known as COFDM is appropriate [10]. Care is nevertheless needed as in some situations Doppler spread is also high, making DRM in some ways a more difficult application than DAB or DVB-T.

The *C* of COFDM stands for the essential forward error-correction *C*oding. The redundancy this provides enables the receiver to cope with both noise and the effects of fading and other channel impairments. The coding is based, like DAB and DVB-T, on the use of a convolutional code, but in a slightly more elaborate arrangement called Multi-Level Coding (MLC).

10. A single-carrier modulation system would require that the receiver had an adaptive equalizer with a large number of taps. The larger this number of taps, the more favourable the comparison becomes to COFDM, in terms of complexity and feasibility.

Since nearly-flat fading can occur in this application, as well as frequency-selective fading, both *time* and *frequency interleaving* are used.

COFDM

COFDM should by now be well known from its use in DAB and DVB-T. For a fuller explanation of the basic principles, please see an earlier article by the author [3]. The data to be transmitted are coded, using a convolutional forward error-correcting code, and distributed over a number of carriers for modulation and transmission. Each carrier thus has only a relatively small bit-rate to convey. This makes it possible to ensure that multipath spread is small compared with symbol length, so that, with the addition of a small *guard interval*, inter-symbol interference in the presence of multipath can be eliminated.

Furthermore, the carriers are arranged to be spaced in frequency by the reciprocal of the length in time of the window in which the receiver takes a "snapshot" of the waveform corresponding to each symbol – this ensures absence of crosstalk between carriers (in the absence of Doppler spread). (This is known as the orthogonality criterion).

The process of modulating and demodulating the many carriers is achieved far more easily than it is described, using a Fast Fourier Transform (FFT).

COFDM is resistant to common transmission impairments. For example, if one carrier is severely attenuated by selective fading, the receiver can note this and flag the coded data bits that are demodulated from it as being potentially of low reliability. The error-correcting decoder can then take this channel-state information into account in the decoding process. Similarly, provided there is sufficient interleaving, a similar approach could be taken when all carriers in a symbol are disturbed by a brief flat fade or impulsive disturbance.

Reference cells (like the "scattered pilots" of DVB-T) are inserted so that the channel response can be measured, and some further reference cells are provided to aid synchronization.

Flexibility

The advantages of COFDM can only be obtained if its parameters (carrier spacing, guard-interval length) are appropriately chosen to match the channel impairments. If the guard interval is too short, then performance will be lost because of the inter-symbol interference; conversely, for it to be too long is wasteful as the data capacity is less than it could be. If the carrier spacing is too small, then the ability to withstand Doppler spread will be insufficient.

Considering the wide range of broadcasting uses to which the LF, MF & HF bands are put, and the different propagation modes, DRM provides a range of "OFDM parameter sets" (carrier spacing / guard-interval combinations) so that broadcasters have the flexibility to adapt to their circumstances while maximizing capacity. The code rate and constellation type (16-QAM or 64-QAM for the MSC) can also be chosen as part of the capacity/ruggedness trade-off.

Two of the OFDM parameter sets are expected to cover most applications. One has a guard interval of $2\frac{2}{3}$ ms, together with a carrier spacing of $41\frac{2}{3}$ Hz. This is described as intended for "Gaussian channels, with minor fading", and is thus particularly suitable for local or national coverage at LF/MF, although it may also be useful in some longer-distance applications. The guard interval is sufficient for SFN operation.

Another parameter set is described as intended for "Time and frequency selective channels, with longer delay spread". It has a guard interval of $5\frac{1}{3}$ ms, with a carrier spacing of $46\frac{7}{8}$ Hz and a higher density of pilots. The longer guard-interval copes with greater multipath spread, as can be caused during multi-mode, multi-hop sky-wave propagation, while the greater carrier spacing and pilot density give greater tolerance to Doppler spread.

It will now be clear why measurements of delay spread and Doppler spread have been made during DRM field trials.

For each parameter set, a total number of carriers is specified to make the whole signal occupy just less than the nominal bandwidth available. There are bandwidth options corresponding to a normal channel, a double channel and one-half channel. As both 9 and 10 kHz channelling exists in the world (see *Appendix A*), this makes six nominal bandwidths in all: 4.5, 5, 9, 10, 18 & 20 kHz.

Where the data streams go

The FAC and SDC are fitted into the signal in ways that support their functionality (see *Fig. 2*).

As the FAC tells the receiver which bandwidth option is in use, it follows that the FAC must be confined within 4.5 kHz, whatever the bandwidth option in use, so that the receiver can find it. The same subset of the spectrum also contains certain synchronization references for the same reason.

There is also a division in time. A transmission superframe is 1.2 s long and is divided into three transmission frames of 400 ms, each containing a whole number of OFDM symbols.

The pattern of reference cells continues throughout; however, data cells are treated slightly differently. Symbols at the beginning of each transmission superframe are used to carry the SDC alone, while all the other symbols are used to carry the MSC and FAC. This time-division arrangement is intended to facilitate alternative-frequency switching (AFS) in simple receivers.

The SDC contains only data that is sent repeatedly. Once the receiver has all this information, it does not need to decode the SDC data in every transmission superframe. This means that there is the opportunity (every 1.2 s) for a receiver which has a single front end to "take a look" at what is happening on another frequency, returning to the main frequency in time to continue receiving the MSC without losing any data.

Various strategies are open to the receiver designer as to what is meant by "taking a look". At a minimum, the receiver could determine whether a signal is present, and estimate its strength. However, it would be preferable to:

⇨ confirm that the signal found on the alternative frequency carries the same service as the first (if this is the case, both signals will carry the same SDC and hence have the same emitted waveform);

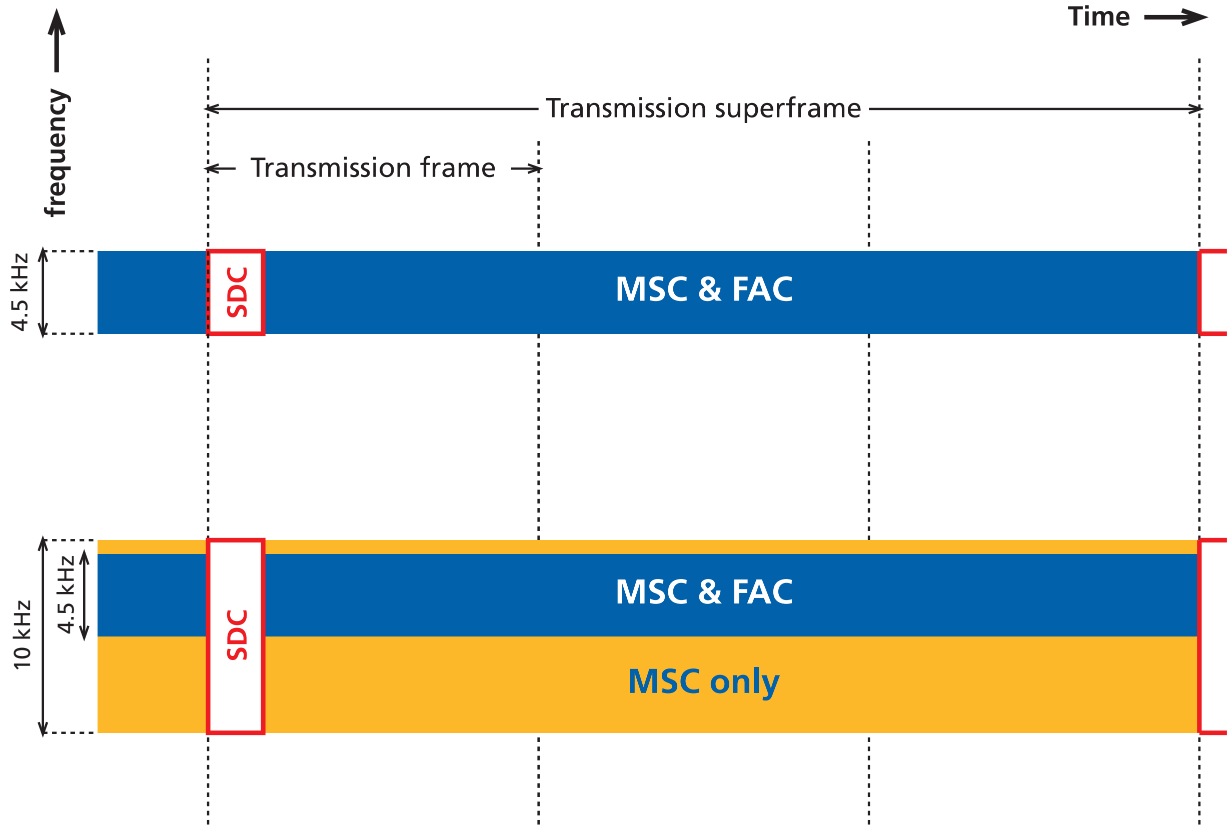

Figure 2
DRM frame structure, with two bandwidth options, showing that the FAC information is always to be found in the same part of the spectrum, and that the SDC is periodically inserted.

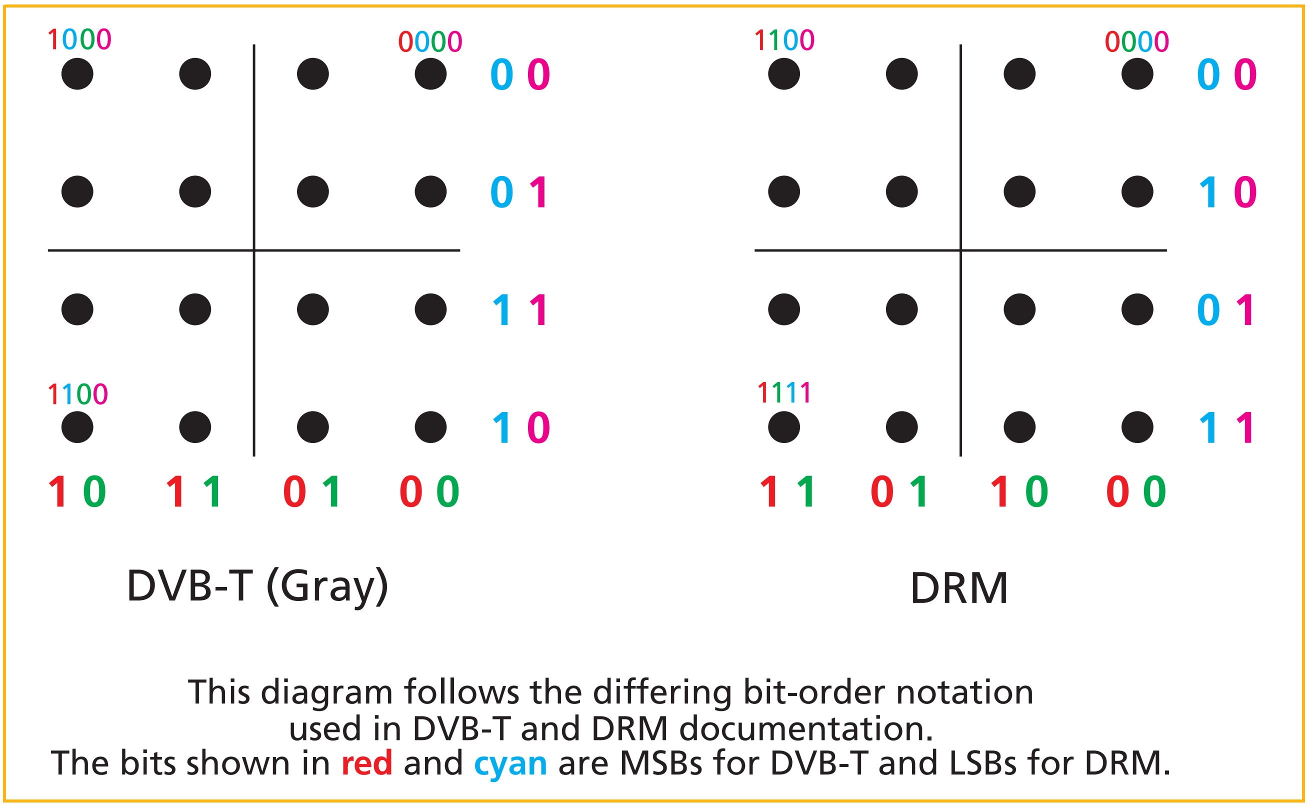

Figure 3
Comparison of mapping applied to a 16-QAM constellation in DVB-T and DRM.

- estimate the quality of reception;
- approximately determine the relative synchronization.

In this way the receiver can decide whether the signal available on the alternative frequency is worth switching to, and do so at the next opportunity.

Multi-level coding

DRM, like DVB-T, modulates the MSC data onto the many OFDM carriers using a QAM constellation (either 16-QAM or 64-QAM). Several coded bits are carried by each constellation. Both DRM & DVB-T also have "hierarchical" modes, which this article will not discuss further – the following applies to "normal" modes only.

DVB-T uses Gray coding of the constellation points [11] (see *Fig. 3*). All the bits used to determine which point is sent for each data cell are taken from the same coded data

11. The Gray coding reduces the likelihood of hard-decision errors occurring simultaneously in more than one bit.

Initially, the receiver has no knowledge which of the four states has been transmitted:

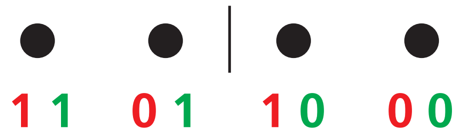

The receiver first decodes the LSBs from all the data cells in a coded block. Once this is done, it knows (within the level of reliability given by the strong coding applied to the LSBs) which LSB was sent for each constellation, simplifying the choice for the MSB:

if the LSB was decoded as **0**

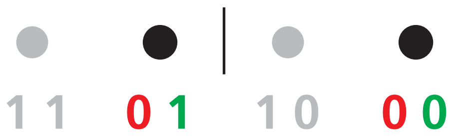

if the LSB was decoded as **1**

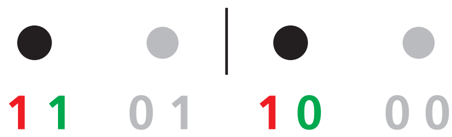

In deciding whether a 0 or 1 was transmitted for the MSB, the receiver now only has to distinguish between 2 more-widely spaced possibilities, which should be within the capability of the weaker coding applied to the MSBs.

The process can be repeated in further iterations: knowing the MSBs, the LSBs can be decoded again, and so on.

Similarly, for 64-QAM, the three levels LSB, CSB and MSB are successively decoded.

Figure 4
Steps in multi-level coding (one axis only shown).

stream and thus have the same degree of protection from the coding. However, the nature of the mapping is such that the bits are not equally robustly conveyed – a coded LSB [12] is more likely to be wrongly interpreted than a coded MSB. This imbalance is not as bad as it sounds, because what is fed to the Viterbi decoder is not a hard decision (0 or 1) but instead a soft decision, in which the likely reliability is described. On average, MSBs will have soft-decision values indicating greater certainty than LSBs. The Viterbi decoder takes all this into account during decoding. Nevertheless, another approach is possible.

12. Here I mention only LSBs and MSBs, as arise in 16-QAM, which maps two bits-per-axis. The argument is similar for the three bits-per-axis of 64-QAM. Note also that I use LSB, MSB to refer to bits (some authors use LSb, etc, reserving "B" for byte).

The alternative is multi-level coding. In this case, the differing frailty of constellation LSBs and MSBs is taken into account by taking them from two streams that have been coded using different code rates. The MSBs are coded at a higher rate – i.e. less strongly coded – in recognition that the receiver is less likely to get them wrong.

In principle this would give an improvement in efficiency – more capacity in bit/s for the same bit-error ratio (BER) at the same signal-to-noise ratio – on its own. However, provided that the mapping is appropriately chosen, a further refinement can be invoked in the receiver. For simplicity, 16-QAM will be taken in the following explanation, which is an example of the approach that can be followed, once this kind of mapping and separate coding is adopted.

The receiver first decodes the coded LSBs from the constellations of a block of received data cells. This gives the corresponding (smaller) stream of data originally sent (these bits were strongly coded, remember). The stream can then be re-encoded [13], so that on re-visiting the same received constellation points, the receiver knows (to a certain reliability) which LSBs were mapped onto them. This simplifies the decision as to which MSB must have been sent (see *Fig. 4*). In effect the distance between a 0 and a 1 for the MSB has been increased, and decoding thus made more reliable.

It is possible to perform multiple iterations of this decoding process, in each case using the recently decoded results to improve the reliability of the next step of decoding. In doing this a modest performance benefit can be obtained.

Interleaving

For simplicity, the explanation just given of Multi-Level Coding neglected any mention of *interleaving*. Most systems using forward error-correction coding need interleaving – although the error-correction decoder can cope with a significant proportion of bits being corrupted, it would be overwhelmed if the corrupted bits came in a group together. Certain probable propagation effects have the potential to cause this kind of difficulty: selective fades can consistently degrade groups of nearby carriers, while a flat fade can degrade all the carriers simultaneously in one OFDM symbol. For this reason DRM applies both time and frequency interleaving. The main interleaving is done on a cell-wise basis, using a convolutional structure in order to minimize the total through delay. Two lengths of time interleaver are provided, so that a broadcaster expecting stable propagation conditions (e.g. an MF ground-wave service) can offer the listeners a shorter delay before hearing the programme after switching on, or selecting a different service. Taking account also of the fact that audio data is grouped into frames, the fundamental through delay (sender and receiver) is of the order of 0.8 or 2.4 s for the short or long interleaving options respectively.

13. Various de-interleaving and re-interleaving steps are omitted for simplicity, although they are essential in practice.

Further short interleavers are provided on the individual levels of the MLC.

Conclusions

Digital Radio Mondiale is a system that promises to re-invigorate the use of the broadcasting bands below 30 MHz. It offers a dramatic improvement in audio quality, not only improving the audio bandwidth and signal-to-noise ratio, but also countering the effects of selective fading and audible interference from other stations. It is also designed to support various features that will make receiver operation more user-friendly. Such features as identifying the station, listing alternative frequencies and supporting the provision of time-varying frequency schedules will transform the nature of listening, especially to broadcasts using short-wave bands.

All this is made possible by the adoption of state-of-the-art digital techniques, from audio source coding to channel coding and modulation.

DRM is the result of co-operation between a large number of key players – broadcasters, transmitter operators, manufacturers (of transmitters, receivers and semiconductors) and research organizations. It is flexible, so as to suit the varied requirements of different broadcasters – and the environment in which they operate. It has been recognized by the International Telecommunications Union in a draft ITU-R Recommendation [2].

Acknowledgements

The author wishes to thank all those in the DRM member organizations who have taken part in the development and evaluation of DRM. Without their collaborative efforts there would be no DRM system to form the subject of this article.

Jonathan Stott studied Engineering and Electrical Sciences at Churchill College, Cambridge University, graduating with Distinction in 1972. He has performed research for the BBC ever since, almost exclusively on the application of digital techniques to broadcasting. He is currently a Project Manager in the Spectrum Planning Group of BBC R & D.

Jonathan was deeply involved with the development of the DVB-T standard for digital terrestrial television, and moved on from this to the team at BBC R&D which is working with the DRM Consortium to develop a standard for broadcasting in the bands below 30 MHz. While working on DRM, he has also become aware of the potential threats to this part of the spectrum from various new types of interference, such as PLT and xDSL, so "protection of the spectrum" has become a parallel task.

Bibliography

References

[1] DRM website: http:/ /www.drm.org

[2] Draft new Recommendation ITU-R BS.[Doc.6/63]: **System for digital sound broadcasting in the broadcasting bands below 30 MHz**
ITU, Geneva, 2000.
http://www.itu.int/itudoc/itu-r/sg6/docs/sg6/2000-03/contrib/063e.html

[3] J.H. Stott: **The How and Why of COFDM**
EBU Technical Review No. 278, Winter 1998
http://www.ebu.ch/trev_278-stott.pdf

Further reading

Hall, Barclay and Hewitt (Eds): **Propagation of radiowaves**
IEE Publications, Stevenage, UK. 1996.
http://www.iee.org.uk/publish/books/elecwave.html#Propagation

Kenneth Davies: **Ionospheric Radio**
IEE Publications, Stevenage, UK. 1991.
http://www.iee.org.uk/publish/books/elecwave.html#Ionospheric_radio

Appendix A: Frequency allocations in the AM bands

Broadcasting frequency allocations are set out in the International Radio Regulations, being revised from time to time at World Radio Conferences. There are some variations according to the part of the world, with the following examples reflecting the situation for Europe unless otherwise stated.

- **LF:** channels at 9 kHz spacing, centre frequencies 153 to 279 kHz
- **MF:** channels at 9 kHz spacing, centre frequencies 531 to 1602kHz

The Americas do not use LF, and have 10 kHz channelling at MF.

Several HF bands are available for broadcasting, around 4, 6, 7, 9, 12, 13, 15, 17, 19, 21 and 26 MHz. 5 kHz channelling is used, but with 10 kHz nominal RF bandwidth. In addition, some “Tropical Bands” at the low-frequency end of the HF band, around 2, 3 and 5 MHz, are supposed to be reserved for use by countries in the Tropics for their national broadcasting, whereas the other bands can be used for national or international broadcasting.

Appendix B: Ground- and sky-wave propagation

A ground wave “hugs” the curvature of the Earth (see *Fig. B1*). The conductivity and permittivity of the ground significantly affect the propagation. It should really be called a *surface wave*, especially as propagation is at its best over the salty water of the sea! Propagation is better over wet ground than say the dry sand of a desert. Vertical polarization is used since horizontal polarization would be severely attenuated.

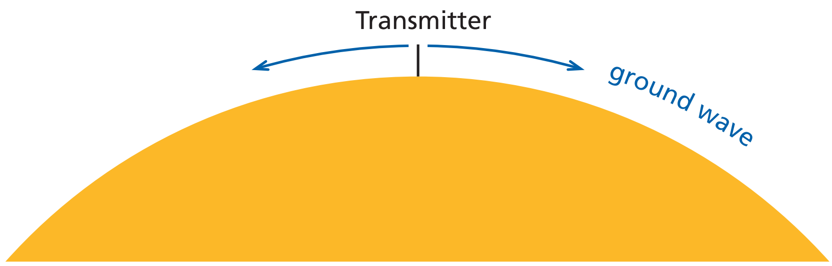

Figure B1
Ground-wave propagation.

Sky-wave propagation is possible because of the ionosphere. This contains layers of charged particles (ionised by the actions of the Sun’s rays) which interact with radio waves. Depending on the degree of ionisation, the gas density and the frequency of the radio wave, the wave may be *refracted* (bent) or *absorbed*. The relatively dense, lower, D layer mostly absorbs (especially daytime MF) but the higher E and F layers can bend the wave so much that it returns to Earth. As viewed from the ground, it is as if the wave had travelled in straight lines but was *reflected* from a fictitious, slightly higher layer. If the wave is reflected by the ground on its return, then further “reflection” by the ionosphere can occur so as to reach far-distant locations by *multiple hops* (see *Fig. B2*).

If the frequency is too high, the wave is not bent sufficiently to return – there is a critical frequency which is the highest for which reflection occurs. This critical frequency is least for vertical incidence.

For international broadcasting the transmitting antenna is designed to favour the appropriate elevation of the wave so that the distant location can be reached by one or more hops. However, another type of broadcasting deliberately aims the wave *upward* so that it is reflected back down to the region surrounding the transmitter (see *Fig. B3*). This means that it is suitable to provide national coverage, and is useful in conditions where ground-wave coverage would be difficult. This *Near-vertical incidence sky-wave* (NVIS) broadcasting requires the frequency to be below the critical frequency;

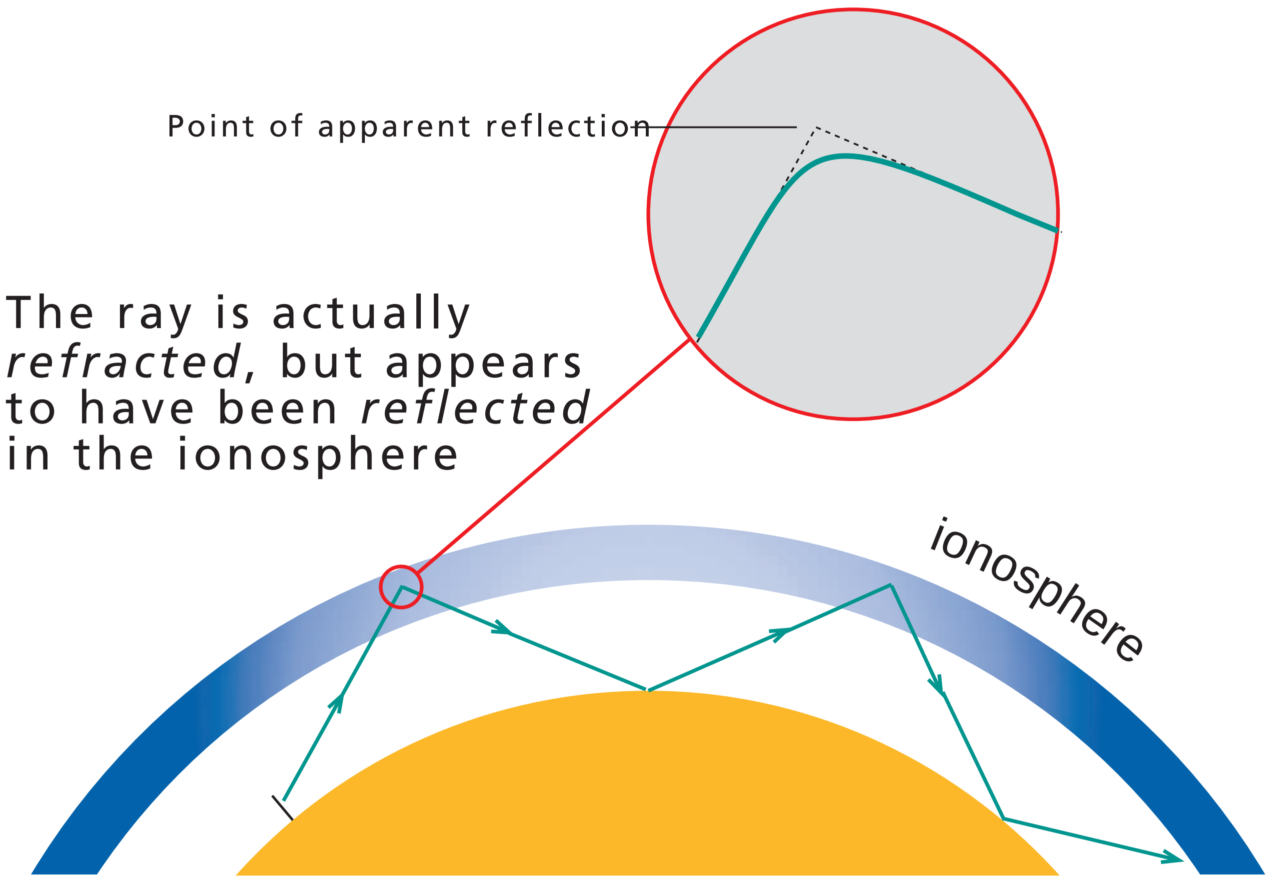

Figure B2
Sky-wave propagation.

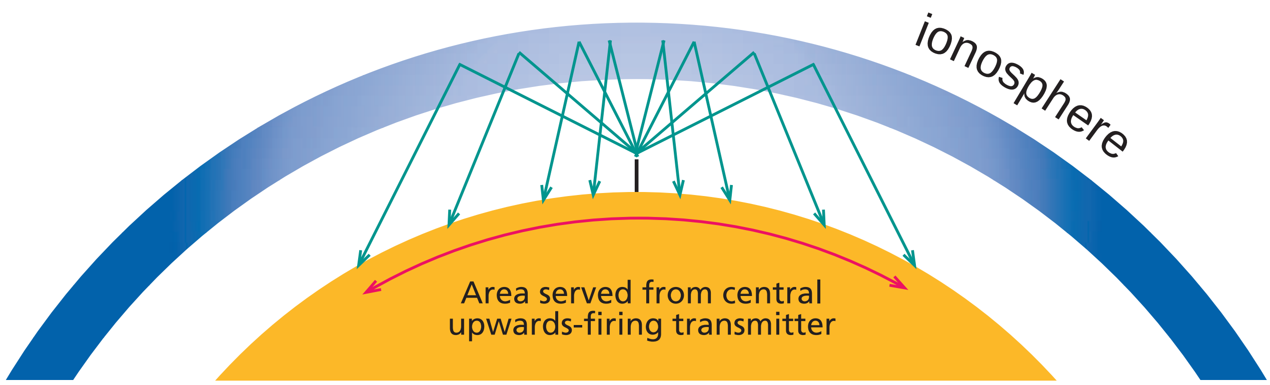

Figure B3
Near-vertical-incidence sky-wave propagation (the multiple hops after ground reflection are not shown for clarity).

the relatively lower frequency Tropical HF Bands are thus set aside for this purpose. A particular feature is that when losses are low enough to support multi-hop propagation, there will be severe multipath, as the first, second, etc. hops can all be received within the service area (together with the ground wave if close to the transmitter).

This description of ground- and sky-wave propagation is very much simplified, reflecting the author's level of understanding of what is a very complicated subject – see the Bibliography for some further reading on this topic.