



Julian Cable

BBC World Service

Broadcasters, both international and domestic, are busy turning their pilot DRM transmissions into real services. Meanwhile, the receiver manufacturers are trying hard to get their DRM radios into the shops as quickly as possible.

This article gives an insight into the DRM distribution chain used by BBC World Service, which is made up of a mixture of professional equipment and in-house developments. Playing an important role in the BBC's overall DRM strategy is the use of open-source software, wherever possible.

The DRM system

The DRM system has been designed as an end-to-end system spanning from content encoding to reception by the listener. As originally conceived and currently defined, it is optimized for use in the HF, MF and LF bands and is the digital replacement for AM in these bands.

The system uses state-of-the-art mechanisms to enable broadcasters to deliver an FM-like listening experience to fixed and mobile receivers over the challenging propagation conditions experienced in these frequency bands. DRM further exploits digital technology to provide previously impossible ease of use.

These mechanisms operate at many levels.

At the basic RF-channel level, DRM shares many characteristics with other modern systems such as DVB-T and DAB. It utilises Coded Orthogonal Frequency Division Multiplexing (COFDM) [1][2] to combat noise, narrowband fading and multipath. Broadcasts can be synchronized to support Single-Frequency Networks (SFNs) and Multi-Frequency Networks (MFNs). One DRM service normally occupies one AM channel. Alternative modes are available that use half a channel or two adjacent channels to support in-channel simulcast or enhanced audio services.

The RF channel carries three separately decodable signals:

- a very robust Fast Access Channel (FAC) which carries station identification and allows rapid band-scanning by receivers to determine the available services;
- a high capacity Main Data Channel (MDC) which carries the audio and/or data services;
- A Service Data Channel (SDC), intermediate in robustness and capacity, which carries meta-data for the main data channel.

Many of the usability features of DRM rely on data within the SDC. This data can include detailed language and genre information, alternative frequencies, date and time, traffic announcements and

many other items of metadata. This allows receivers to present the user with a very simple browsing interface to select the desired radio station. The receivers can also allow users to search for specific genres or can present only services in the users' preferred language. The SDC also enables seamless handover between frequencies in an MFN.




Figure 1
DRM spectrum occupancy

The main data channel can carry up to four simultaneous streams. This is not intended to provide a service multiplex such as DVB or DAB provides – the channel is not large enough for that. The ability to carry more than one stream is very useful, especially where one of the channels is a data service. This allows many applications, such as downloading of an Electronic Programme Guide (EPG) into the radio, providing visual enhancements to the audio, and pushing audio files into the receivers memory to enable “audio on demand” type services. It has also been used in conjunction with low-bitrate codecs to provide the same audio service in multiple languages within one channel.

The audio is encoded using the latest MPEG-4 implementations of CELP, HVXC and, most importantly, **AAC**. AAC supports music and speech down to around 16 kbit/s, which covers the main uses of DRM. Below this bitrate, CELP operates between 16 kbit/s and 6 kbit/s and then, HVXC takes over between 6 kbit/s and 2 kbit/s.

These lower bitrates are useful for the most robust DRM modes and the half-channel DRM modes, or if more than one audio stream is required in the channel.

Spectral Band Replication (SBR) can be used together with any of the codecs to enhance the high frequencies.

BBC World Service broadcasts

DRM is important to BBC World Service (BBC WS) because it may provide a distribution platform for broadcasting to listeners in digital “FM-like” quality over very large geographical regions. It also allows the BBC to continue to broadcast to areas where it would not be welcome to install its own transmitters. DRM broadcasting on HF shares many of the same characteristics as satellite radio, but with economics that do not require a subscription model.

The core audience for BBC World Service on DRM is eventually likely to be in Africa and Asia. However, mass-market consumer radios are needed to make this possible. The interest in making DRM radios is initially in Europe and currently our main aim is to provide compelling content to stimulate the production and sales of consumer radios. The experience of DAB in Europe has shown broadcasters that we have to invest in transmissions – even when there is no real audience – in order to provide receiver manufacturers with the confidence to proceed with design, manufacture and marketing activities. In partnership with the other members of the DRM consortium, we are trying to create a virtuous circle of confidence in DRM.

The transmitter sites and approximate footprints are shown in *Fig. 2*. Reception is often available over a wider area, including parts of Italy. We have even had reception reports from Athens,



Figure 2
BBC World Service DRM broadcasts

although the received quality there is not reliable.

BBC World Service currently transmits DRM for 18 hours per day on the medium-wave (MW or MF) band from Orfordness and 12 hours per day each from Rampisham and Kvitsoy. All the transmissions have the same core content:

- English for Europe, using AAC+SBR in mono at approximately 20 kbit/s;
- Dynamic Text Messages;
- 7-day Electronic Programme Guide.

We occasionally supplement this with special data services, such as tests of weather information, etc.

BBC World Service DRM broadcast chain

Overview

BBC World Service generates a distribution feed at the main BBC World Service offices at Bush House in London which consists of the pre-coded and multiplexed signal for transmission. This signal is then modulated at each transmitter site. The overall distribution therefore looks like the diagram shown in Fig. 3.




Figure 3
DRM signal distribution via satellite

The Bush House infrastructure is similar to that shown in Fig. 4. The audio and data sources are aggregated and conditioned before the DRM-specific multiplexing and coding equipment produces the multiplex distribution signal. This is encapsulated for broadcast via satellite.

The individual processes and sources are described in detail below.

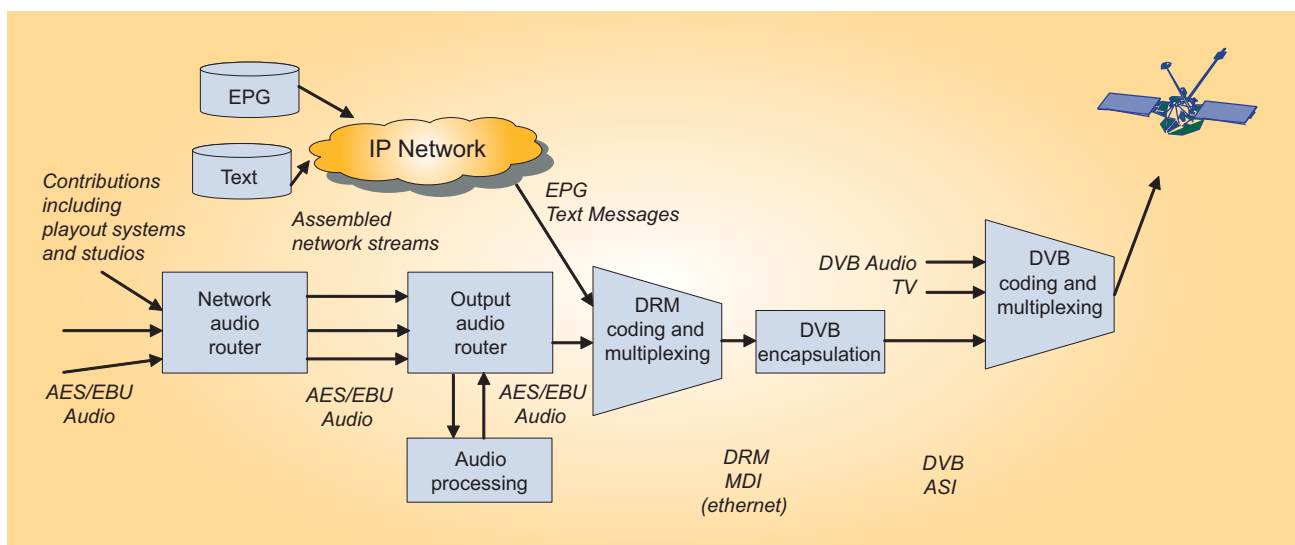


Figure 4
DRM content aggregation, coding and multiplexing

Content conditioning and aggregation

The content of our current core DRM transmission is all re-purposed from our UK DAB content offering. The re-purposed DAB content consists of the following:

- the European English stream of World Service audio;
- the dynamic text messages carried as the “dynamic label service” on DAB;
- The electronic programme guide.

All of these need some conditioning actions prior to aggregation by the content server.

We supplement the DAB content with DRM-specific metadata and with ad-hoc content for specific demonstrations. For example, we have been transmitting the UK Navtex service, re-purposed from the BBC Weather website.

Audio

BBC WS broadcasts, including the European English stream, comprise a fairly broad mix of content based on a core of world news and information. This includes well-controlled studio broadcasts, pre-recorded music, phone-ins, location reports and jingles. Our audience is generally thoughtful, adult and discriminating and so we avoid trying to create an obviously synthetic and fatiguing “commercial” sound, which could be a distraction from the content. Each broadcast stream is assembled with the intent of creating the best quality, linear, digital audio. The European English stream is broadcast on MF and HF AM, on DAB, DVB-C and DVB-T (Freeview) in the UK, on DVB-S to the UK and Europe, and on the internet. It is clear that when coding the linear stream for the different transmission media, the coding process affects different programme material in different ways.

With the exception of DRM, the transmission and/or distribution feeds for all of these services uses MPEG-1 layer II codecs at bitrates of between 64 kbit/s and 192 kbit/s. In order to get the best sound to the listener via each medium, we apply separate audio pre-conditioning prior to encoding. The main aim of the audio pre-conditioning is to reduce the occurrences of coding artefacts in the broadcast audio. Relatively heavy pre-conditioning, including a low-pass filter at 11.5 kHz, is applied to the DAB stream since it uses an inefficient codec (in comparison with AAC), has a low bitrate of 64 kbit/s, and will often be listened to in high-noise environments. We don't use any pre-conditioning for DVB-S since the end-to-end bitrate is high enough to produce good quality results.

For DRM, in common with most broadcasters, BBC WS encodes general audio broadcasts using the AAC codec with Spectral Band Replication (SBR) to improve the high-frequency listening experience. With only about 21 kbit/s for the audio channel, the pre-conditioning is important in order to achieve acceptable results. In our experience, the AAC coding produces the best results when fed with a band-limited signal. On the other hand, the SBR coding needs a full-bandwidth signal in order to analyse the spectral content. The situation is currently further complicated by the fact that many listeners at present are highly technical and are using software receivers which can display the audio spectrum. Their listening experience may be coloured by their visual perception of the presence or absence of high-frequency content.

In front of the coder sits a professional multi-band audio processor. The processor is configured with fairly slow release characteristics. It controls levels over around 10 dB to mask any large differences between sources without creating too dense a sound, especially at the top end of the frequency range. We have also minimized the amount of hard limiting and clipping so that audio transients survive the processing well. The AAC+SBR coder seems to sound best when it has some sharp peaks to “get its teeth into”.

Text messages

The text messages transmitted on DAB are originated with a UK audience in mind. They are generated in real-time by an in-house content production and scheduling system operated by the Radio & Music division of the BBC and delivered to the content server using a push mechanism over the BBC corporate intranet. The content server performs the following conditioning actions on the text messages:

- translating from ISO Latin 1 to UTF-8;
- translating UK local telephone numbers to international format numbers;
- translating local times to UTC (GMT).

Electronic Programme Guide

The DAB/DRM electronic programme guide is in the final stages of ETSI standardization. No further technical changes are anticipated in the final standard. BBC World Service transmitted an interim-format guide as a modified Broadcast Web Site from September 2005 until February 2006. Since then we have transmitted a standards-compliant stream. This stream can be decoded by listeners using the Dream open-source software radio [3] using a decoder contributed by BBC World Service which we hope will be supported by the first combined DAB/DRM radios. The DAB and DRM streams were upgraded in March 2006 to incorporate the latest TV-Anytime Genre information. This highlighted a small problem with the Dream receiver, which was fixed and released within 48 hours. The advantages and challenges of open source are discussed later in this article.

DRM metadata including Alternative-Frequency Signalling

In addition to the live audio and data streams entering the coding and multiplexing equipment, DRM carries a significant amount of metadata. We currently enter this data manually into the configuration files of the Coding and Multiplexing equipment. This is feasible since the data is quasi-static. The most common changes are to the alternative-frequency lists and when a new application is added.

The alternative-frequency information is carried in the DRM Service Data Channel (SDC). Broadcasts of BBC WS list the frequencies, start times and durations of all DRM and AM transmissions targeted at mainland Europe. The receiver manufacturers are currently working on automatic retuning using this information. This is extremely important to provide a seamless listening experience on short-wave where transmitters must change frequency throughout the day to take advantage of the changing propagation conditions.

Coding and Multiplexing

In DRM, the device responsible for Coding and Multiplexing is generally referred to as a “content server”. A typical content server has one or more audio inputs, a LAN input to receive text messages, multimedia data and metadata, and a LAN output to transmit the multiplexed stream.

The design of the multiplex is a trade-off between the audio quality of the main audio service and the time taken to deliver any additional services. *Table 1* shows the information currently transmitted by BBC World Service in the multiplex.

Our DRM broadcasts carry traditional radio programmes and it is important to our listeners that we maintain the audio quality. *Fig. 5* illustrates that 97 % of the main service channel is used for audio.

The coding and multiplexing process inevitably introduces a delay. For DRM when coding at the studio, this delay is a minimum of one DRM audio superframe or 400 ms. This is a consequence of transmitting the multiplexed stream in packets describing one such frame. The actual coding delay

Table 1
Content rates and update intervals

Content	Quantity	Channel	Rate	Source update interval	Carousel period
Audio	-	Main Service Channel Stream 0	As high as possible	Live	-
Text Messages	4 bytes per audio superframe	Main Service Channel Stream 0	80 bit/s	Live	-
Electronic Programme Guide	≈ 90 kbytes	Main Service Channel Stream 1 packet id 0	490 bit/s	Twice daily	25 minutes
Navtex	≈ 1 kbyte	Main Service Channel Stream 1 packet id 1	50 bit/s	Twice daily	3 minutes
Alternative Frequency Data	≈ 100 bytes	Service Data Channel	≈ 200 bit/s	Twice yearly	4 seconds

is mostly absorbed in this packet assembly delay, since the audio is coded in audio frames, of which there are typically 5 or 10 in an audio superframe. The true coding delay is therefore less than 100 ms. This means that the packet may be sent when the newest audio is around 100 ms old, but it will contain audio which is 400 to 500 ms old. A DRM audio superframe is also known as a DRM frame. In principle, a coding and multiplexing implementation, integrated with the modulation functions in a transmitter, could have a minimum lower delay.

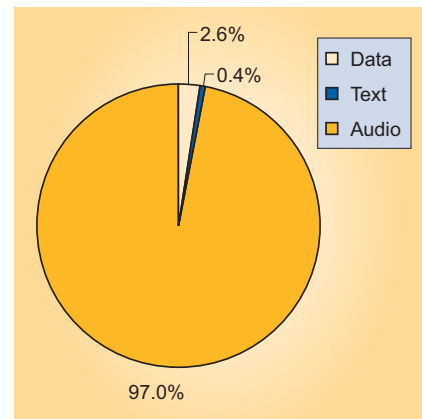


Figure 5
Sharing the channel capacity

Distributing the multiplex to the transmitters

Overview

DRM is designed to meet the needs of both small and large broadcasters. For small broadcasters, such as those providing metropolitan-area broadcasts in the 26 MHz band, the coding and multiplexing, modulation and amplification may be all in a single box. Larger broadcasters can obtain distribution bandwidth efficiencies by performing coding and multiplexing at source. For Single Frequency Networks and Synchronous Multi-Frequency Networks, this is essential to ensure all transmitters produce identical bit-streams.

BBC World Service codes and multiplexes at its primary site in London and delivers the multiplex streams via an IP multicast network to the transmitting stations. Two streams are generated, one for the HF transmissions and one for the MF transmissions. The total distribution bandwidth for this is 64 kbit/s. The HF stream is delivered to two transmitting stations, one at Rampisham in the UK and one at Kvitsøy near Stavanger in Norway. These form a synchronous multi-frequency network.

The DRM Multiplex Distribution Interface

The DRM Consortium has produced a very successful standard for the distribution of DRM multiplex signals to transmitters. This is called the DRM Multiplex Distribution Interface or MDI. The MDI carries the separate components to be broadcast in their own parts of the message, but integrated into a single message every DRM frame (400 ms) for error protection and synchronization purposes.

The MDI protocol is layered to allow selective use of different components and to allow the general-purpose parts to be re-used for other purposes. For the MDI there are three layers of note. The Top layer is the MDI itself. This consists of a number of named data fields carrying, for example, the Fast Access Channel (FAC) data for the frame, the Service Data Channel (SDC) data for a superframe and the Main Service Channel (MSC) data for the frame. In order to be able to carry this data over any arbitrary channel, an Application Framing (AF) layer is added below the application data. This provides a means to delimit packets, synchronize to the start of a packet and detect missing or out-of-order packets. The AF layer assumes an underlying highly reliable pipe which delivers bits in order, such as is provided by a local serial data link, or a TCP link. It can also be used over packet channels such as UDP, provided the maximum packet size is large enough to carry the AF packets.

If the underlying channel is not suitable for carrying AF packets, then an optional additional layer can be used. There are several circumstances when this additional Packet Fragmentation and Transport layer is appropriate:

- high residual error rate;
- delivery of packets out of order;
- packet link and maximum packet size too small;
- need to carry more than one channel and no multiplexing capability.

The PFT has specific optional mechanisms for each of these concerns. If multiplexing is needed, for example over a local serial data link, then the addressing mechanism can be used. This provides a 16-bit source address and a 16-bit destination address which can be used to distinguish different streams. If the maximum packet size of the underlying network is too small for a full AF frame, then the PFT provides a fragmentation and reassembly facility. This uses sequence numbers in each fragment which can be used to re-assemble fragments even if they are delivered in the wrong order. Finally, a Reed-Solomon forward error correction (FEC) facility is provided which incorporates optional interleaving. This is beneficial when the probability of burst errors or packet loss is high.

BBC WS current implementation

The signals to be broadcast must be distributed error-free and in time. A suitable residual error rate is less than one lost packet per day. "In time" means that the packets must arrive at the transmitter sufficiently in advance of the required transmit time to be handled by the DRM modulator. Any variation in the distribution delay must be added to the overall delay budget for the system.

BBC World Service uses a distribution system based on terrestrial leased lines and DVB-S for its primary distribution feed. The end-to-end delay has been measured at 400 ms, or equivalent to one DRM frame. The packet loss rate is much less than one per day.

It is clear that dedicated satellite provision is not viable for all broadcasters. World Service is lucky to have small amounts of capacity available on its main distribution system to permit this solution. The satellite delivery mechanism has two extremely important benefits:

- it provides an extremely low error rate;
- it is inherently point-to-multipoint, allowing additional transmitters and stations to receive and broadcast the signal at minimal cost per receiver.

Early on, we took the decision to transport the MDI over the satellite in a standards-conformant manner. We carry the MDI as an IP multicast stream using DVB-MPE (multi-protocol encapsulation). This allows the use of low-cost, off-the-shelf equipment at both the coding and multiplexing stage and at the receivers. Indeed, the MDI can be received by a standard PC with a domestic PCI or USB satellite receiver and standard software. This has proved extremely useful in deploying and monitoring the system. We also encourage other members of the Consortium to receive the satellite stream for use in equipment testing.

Where possible, we use the simplest version of the MDI, without using the PFT part. However in Mode A with a 10 kHz bandwidth, the main service channel is large enough that the MDI packet does not always fit in one UDP datagram. In this case, the PFT is used without forward error correction or addressing to provide a fragmentation and reassembly function. This shows just how flexible the DRM distribution protocols are.

Incorporating IP into the distribution system also facilitates monitoring and resiliency. For example, we can use standard packet monitoring and analysis tools to monitor and trouble-shoot the system and, in principle, the internet can be used as an alternative distribution system.

Alternative models

Various alternatives to the present system used in BBC World Service are possible. For example, the low bitrate of the multiplex allows the broadcasts to be carried over ISDN or even dial-up lines.

Transporting the audio and data separately to the transmit site is also possible and, indeed, we use it as a back-up mechanism. We are able to do this because we already broadcast the audio via satellite in good quality for distribution to cable head-ends. Distributing high-quality linear audio to the transmitting sites, specifically for DRM, would be expensive.

When using local encoding as a back-up, we lose the synchronization between the stations. This is acceptable for short periods and has the advantage that there are fewer common-mode failures than if we just provided alternative routes for the MDI.

Transmission

Suitable transmitters

One of the big questions surrounding DRM is whether broadcasters' legacy transmitters are suitable for DRM. Our experience is that only transmitters with a very linear response make good DRM transmitters. In practice this means that old transmitters can be used, but only at low powers and extremely poor efficiencies.

The most modern transmitters, designed with DRM in mind, provide both high output powers and high efficiencies. However, the prospective purchaser is encouraged to examine the transmitter performance, especially the compliance to the spectrum mask, at low output power, as well as at maximum output power. Digital modulators can lose resolution and introduce quantization noise at low output powers.

Single-Frequency Networks

DRM supports single-frequency networking in a similar manner to DAB. BBC World Service does not currently use any single-frequency networks, preferring to provide frequency diversity in a multi-frequency network.

Multi-Frequency Networks and staggered re-tuning

We have experimented with a multi-band synchronous multi-frequency network using the medium-wave and both short-wave transmissions. This involves running the medium- and short-wave transmissions in identical modes. Running the medium-wave transmission at 10 kHz bandwidth has presented few problems. After some experimentation we have chosen to keep the medium-wave transmission in Mode A and the short-wave transmissions in Mode B. Mode B gave disappointing results on medium wave and the channel bandwidth is higher in Mode A, permitting more bandwidth for audio and data.

The medium-wave transmission is on a fixed frequency of 1296 kHz but the short-wave transmitters must change frequency throughout the day to take advantage of the changing propagation conditions. We stagger the changes so that there is always one transmitter on a stable frequency. This is extremely important to provide a seamless listening experience.

For example, in the winter 2005 season, BBC World Service broadcast the following schedule on short wave (the light-yellow rows):

Table 2
BBC World Service DRM – 2005 winter schedule

Start	Stop	Frequency	Station
05:00	23:00	1296 kHz	Orfordness
07:00	09:00	5875 kHz	Kvitsøy
07:00	15:00	7320 kHz	Rampisham
09:00	14:30	9470 kHz	Kvitsøy
14:30	19:00	7465 kHz	Kvitsøy
15:00	19:00	5875 kHz	Rampisham

A listener receiving the broadcast on 9470 kHz at 14:30 would have to re-tune to continue listening. Short-wave transmitters cannot re-tune seamlessly. The BBC specification for transmitter re-tuning for AM broadcasts has always been 15 seconds. For DRM with the current equipment it takes approximately one minute to re-tune the transmitter and then resynchronize the DRM stream.

If a receiver could pre-emptively and seamlessly change frequency to 7320 kHz at 14:28 and then change again at 14:58 to 7465 kHz the listener would not perceive any change to the broadcasts.

To make this possible, we broadcast metadata carried in the Service Data Channel to let the receivers know of all available frequencies and their schedules and we run the transmissions as a synchronous multi-frequency network. This means that the data received on the new frequency will be identical to that which would have been received on the old frequency.

More sophisticated receivers could also simultaneously receive multiple transmissions and use diversity and/or combining techniques to maintain good listening under difficult reception conditions. BBC R&D has produced a prototype receiver which does this.

If DRM reception fails completely, the receiver also knows of AM broadcasts carrying the same audio stream which might, under some conditions, still be intelligible if DRM falls off the “digital cliff”. BBC World Service broadcasts an equivalent metadata stream on its AM broadcasts to Europe. This is intended to guide multi-standard receivers to DRM, either if the AM broadcast was selected in ignorance of DRM, or once the propagation channel recovers sufficiently to support DRM.

This AM metadata is broadcast using a new ETSI standard AM Signalling System (AMSS) developed specifically to make multi-standard receivers attractive and to ease the transition to DRM. The technique has already been described in a previous Technical Review article [4].

Use of open-source, in-house and professional equipment

BBC World Service uses a mixture of sources for its equipment. Whilst it would be wrong to claim that there was a grand plan behind the selection of each part of the system, some aspects of the software and hardware are potentially interesting.

We do try to follow some distinct policies:

- use a common computer platform wherever possible;

- isolate broadcast equipment from the general IT network infrastructure;
- use open standards for the interfaces between equipments;
- don't use operating systems which are likely to become unsupported;
- use systems with web-based management interfaces.

With a system such as DRM, the BBC is not just an early adopter but is helping to develop the system whilst being an early user of it. The BBC has world-class Research & Development resources in this area and we use these to make the system meet our needs and to help make it a success.

One possible approach to the development of a new system would be to participate in the standards process but to allow the equipment manufacturers to produce all the elements. Another approach, which is the one we take, is to produce implementations of key elements ourselves.

Why develop in-house?

Why does the BBC develop implementations? There are a number of reasons and a number of secondary benefits. The primary reason is to ensure that the standards we help to develop can be implemented and will be fit for the purpose. In the area of coding and modulation, for example, without a reference implementation, it is hard to estimate the processing power needed, and therefore the cost. Whilst these could be partial implementations (and sometimes are), completing the implementation to allow it to be used in a live broadcast gives more confidence and allows others to check the implementation.

Again, in the area of protocol design, it is all too easy to design protocols with faults that only become apparent in-use. Being able to test multiple reference implementations together for interoperability is invaluable.

Once we are able to use in-house products, then new features can be added without needing to wait for product release schedules.

Abbreviations

AAC	(MPEG) Advanced Audio Coding	HVXC	(MPEG) Harmonic Vector Excitation Coding
AF	Application Framing	IP	Internet Protocol
AM	Amplitude Modulation	ISDN	Integrated Services Digital Network
AMSS	AM Signalling System	ISO	International Organization for Standardization http://www.iso.org
ASI	Asynchronous Serial Interface	MDC	Main Data Channel
CELP	Code-Excited Linear Prediction	MDI	(DRM) Multiplex Distribution Interface
COFDM	Coded Orthogonal Frequency Division Multiplex	MF	Medium-Frequency
DAB	Digital Audio Broadcasting (Eureka-147)	MFN	Multi-Frequency Network
DRM	Digital Radio Mondiale http://www.drm.org/	MOT	Multimedia Object Transfer
DVB	Digital Video Broadcasting http://www.dvb.org/	MPE	(DVB) MultiProtocol Encapsulation
DVB-C	DVB - Cable	MPEG	Moving Picture Experts Group http://www.chiariglione.org/mpeg/
DVB-S	DVB - Satellite	MW	Medium-Wave
DVB-T	DVB - Terrestrial	MSC	Main Service Channel
EPG	Electronic Programme Guide	PFT	Packet Fragmentation and Transport
ETSI	European Telecommunication Standards Institute http://pda.etsi.org/pda/queryform.asp	SBR	Spectral Band Replication
FAC	Fast Access Channel	SDC	Service Data Channel
FEC	Forward Error Correction	SFN	Single-Frequency Network
GMT	Greenwich Mean Time	TCP	Transmission Control Protocol
HF	High-Frequency	UDP	User Datagram Protocol
		UTC	Universal Co-ordinated Time
		UTF-8	8-bit Unicode Transformation Format

Another very important function provided by a BBC R&D implementation is to accelerate the route to market for equipment manufacturers. The BBC does not turn its implementations into products for sale, but is keen to license its technology on a non-exclusive basis to suitable equipment manufacturers. For DRM this has been extremely successful, with one transmitter and five receiver manufacturers taking out licences for BBC R&D-generated implementations. Whilst this provides a useful revenue opportunity for the BBC, our goal is to help these companies bring products to market.

For some developments, charging licence fees is either impossible or inadvisable. Two examples of this are described in the section on Open Source below.

There are also cost benefits when using in-house developments. One example of this is that we can afford to deploy several instances of an in-house application without additional cost, for example for back-up or disaster recovery facilities. However, the long-term cost of maintaining in-house developed systems is often higher than that of commercial products. There comes a time in most systems where the equipment has become sufficiently commoditised that we can buy better and more cheaply than we can maintain the in-house systems. This is a sure sign of a successful technology.

The role of open-source software

As well as in-house developments and fully commercialised products, open-source software is playing an increasing role in our broadcast chain.

The most significant piece of open-source software is definitely **Linux**. As an embedded operating system kernel, Linux has many advantages over its competitors, of which the most significant are:

- **Reliability** – computers running Linux can run for years without needing to be re-booted;
- **Stable and future-proof** – an appropriate distribution can provide an execution platform which can stay fixed, with no mandatory upgrades or changes. Old versions continue to be available;
- **Embedded** – Linux appliances need no graphical user interface and can be managed and monitored remotely;
- **Compliant** – full implementations of networking stacks and protocols are inherently available, which is not always true of traditional embedded systems;
- **Portable** – applications designed to run on Linux are generally portable to the preferred platform.

The main disadvantage of Linux is that new hardware generally comes without drivers for Linux. This is especially true of audio hardware such as sound cards, and much care must be taken when building Linux machines for audio processing. As an example, the DRM broadcast chain uses Linux Content Servers, IP Encapsulators and Modulators. Of these, the Content Server is an in-house application whilst the Modulator and IP Encapsulator are fully commercial products using Linux as a component. These are all very stable. In contrast, the AMSS Generator runs another operating system and must be re-booted on a weekly basis. The next version of the AMSS Generator will run on Linux.

Interestingly, in deployment, it is best to treat Linux as a “closed source”. In-house developments which change Linux itself are likely to compromise the very benefits that Linux brings in terms of stability and reliability.

When it comes to application programs, there are few rules regarding the suitability of open-source or commercial solutions.

Occasionally a piece of open-source software has so many advantages that its use becomes obvious. In our DRM work, there are two open-source programs which we both use and contribute to: **Dream** and **Ethereal**.

Ethereal (www.ethereal.com)

Ethereal is a LAN Packet Analyser. It is an excellent tool “out of the box” but one of its great strengths is that decoders for new protocols can be written according to a reasonably well-documented programming interface and installed as plug-ins. The BBC has produced such plug-ins for the DRM Multiplex Distribution Protocol. Since Ethereal is licensed under the GNU Public License (GPL), it is not appropriate to charge for licences to use the plug-ins. The DRM-enhanced Ethereal can be used to monitor and debug DRM software such as Content Servers, Modulators and Receivers. Ethereal has a text-only version called *tethereal* which we run in an embedded mode to monitor for packet loss on the distribution network.




Figure 6
Caption ...

The intent is to submit the new protocols for inclusion in the standard distribution of Ethereal. In the meantime, they can be obtained by contacting the author. A sample of the levels of decode available is shown in *Fig. 6*.

Dream (drm.sourceforge.net)

Even more significant than Ethereal is the Dream open-source DRM receiver. This outstanding piece of software was originally developed by Alexander Kurpiers and Volker Fischer at the Darmstadt University of Technology in Germany but has had contributions from many authors. Contributions made by BBC staff to Dream include:

- implementation of the Receiver Control and Status Interface;
- implementation of a decoder for AMSS;
- implementation of a decoder for the Electronic Programme Guide;
- enhancements to the MDI for multicast;
- enhancements to the MOT protocol decoder.

The DRM Consortium represents an official collaboration between broadcasters and equipment suppliers to produce the DRM Standard and make DRM broadcasting a reality. In a similar, but totally unofficial way, the Dream developers' and users' community provides a very significant resource working towards the same goal. Whilst the Dream developers make their contributions mainly in an amateur way, many of them are professional engineers and developers. Without Dream, it is certain that the BBC World Service broadcasts of the Electronic Programme Guide would currently be received only by the receiver manufacturers!

The contributions made by the BBC have been matched and improved upon by other Dream developers. Feedback from the enthusiastic community of Dream users helps us to improve our broadcasts.

In summary, good open-source software is as good or better than commercial software. The right open-source software opens up the opportunity to collaborate with the developers and users of the software. This collaboration is as useful to the BBC as the revenue we obtain from more conventional licensing arrangements.

Sometimes professional equipment is the answer

Much of the preceding discussion has highlighted the benefits of in-house or open source. While cost is a factor, there are other more-compelling reasons to take these routes. However, there are some areas where the only sensible approach is to use fully commercialised professional equipment.

One example is the transmitters themselves. High power transmitters represent a safety-critical environment and using only professionally certified equipment is absolutely essential to provide a safe working environment. The broadcasts are internationally regulated and for DRM, for example, compliance to the Spectrum Mask is an absolute requirement. Even if the DRM Modulator is an in-house design, it is important that the equipment used for compliance monitoring is certified and calibrated.

Another quite different example is where sophisticated user interfaces are required to exercise the features of complex software applications. The DRM Content Server, for example, has hundreds of options. When used exclusively for one broadcast, as in the BBC, a simple in-house equipment has many advantages. However, in the case of a Broadcast Services Provider, the same content server is used to provide broadcasts on behalf of many broadcasters. In this environment, a rich well-designed and documented user interface is essential.

Conclusions

In conclusion, our experience of DRM so far is that, as a set of Digital Radio standards, it works superbly. We have a small and loyal audience for the broadcasts which helps to keep going our enthusiasm as broadcasters, while we wait for the consumer radios.

In describing our distribution chain, the author has tried to show that DRM needs both traditional audio processing and IT / internet skills to produce a rich and reliable service. IP networks in particular play a vital role in both content aggregation and distribution. IP multicast is the right distribution mechanism, but over reliable IP networks.

Open-Source software can play an important role in broadcast systems, both “under the hood” and as a real collaborative effort between professional and enthusiast communities.

Indeed, The communities working to make DRM a success – both professional and enthusiast – work together in a most synergistic and enjoyable way.



Julian Cable is a Senior Project Manager within the BBC World Service where he looks after DRM infrastructure developments.

He joined the BBC in February 2003. As well as project managing the HF Transmitter procurement activities, he is responsible for the technical design of the DRM broadcast chain. He is a principal contributor to the Distribution Interfaces working group within the DRM consortium and works closely with the DRM team at BBC R&D.

Prior to his current role, Mr Cable worked for over 20 years in the private sector, covering development, system integration and research activities in wireless and wire-line communications technologies such as Infosec, VoATM, VoIP, MPLS and Network Management. He holds a degree in Physics and is a Chartered Engineer and a Member of the Institute of Electrical Engineers.

Acknowledgments

The author would like to thank especially the following people for their support of DRM and in the creation of this article. If your name isn't here and it should be, its my fault!

- The BBC R&D Team: Andrew Murphy, Jonathan Stott, John Elliot, Lindsay Cornell, Oliver Haffenden;
- The BBC World Service Team: Babis Metaxas, Colin Wilding, John Sykes, Ketan Chaukar, Margaret Cole, Mark Deutsch, Murray Holgate, Neil Curry, Steve Martin;
- The DREAM developer and user community, especially Andrea Russo, Simone Stöppler, Volker Fischer;
- And, finally, the many people in Coding Technologies, Deutsche Welle, Fraunhofer Institute, Norking, VT Communications and all the other members of the DRM Consortium.

References

- [1] J.H. Stott: [The how and why of COFDM](#)
EBU Technical Review No 278, Summer 1998.
 - [2] J.H. Stott: [DRM – key technical features](#)
EBU Technical Review No 286, March 2001.
 - [3] Dream DRM Receiver: <http://drm.sourceforge.net/>
 - [4] A. Murphy and R. Poole: [The AM Signalling System, AMSS — does your radio know what it is listening to?](#)
EBU Technical Review No 306, January 2006.
-