BBC

BBC Research

Why should I use Dirac?

Dirac offers you a state of the art video compression system. The demonstrations on our stand show that

- Dirac is here.
- Dirac works.
- Dirac is ready for deployment.

But there are many video compression systems, you may want to know what is special about Dirac.

If you are the sort of person that hopes that your content has a long lifetime, consider the upgrade policy of your existing suppliers. Can you handle a format conversion on a regular basis? Can you handle the lack of support for 'obsolescent' technologies? If you cannot, then Dirac offers a long-term solution for you. Because the source code for the software is available for operation on a variety of platforms, you can have confidence that Dirac can stand the test of time.

If you want quality, in your own specialist area, Dirac provides a simple, but powerful range of options which can be selected to fine tune the performance for your content.

If you want a cost effective solution, then the Open Source licence means that you are a customer that has to considered respectfully by your suppliers. You will have a choice of supplier, so you can select one who offers a service which you need, not which they want you to buy.

Why is it called Dirac?

The Dirac codec is named after the British physicist Paul Dirac. Dirac (1902–1984) was one of the most influential scientists of the 20th century. In 1933 he shared the Nobel Prize for physics with Erwin Schrödinger for his contributions to quantum mechanics.

What are the origins of Dirac?

Although named after Dirac, there is no direct connection between Dirac the scientist and Dirac the video coding algorithm.

Dirac was originally developed at BBC Research at Kingswood Warren in the UK. The BBC have specialists in compression technology, having been working on digital television and its ramifications since the 1960s.

The ideas behind Dirac were originally developed as part of an idea to deliver high definition services as a layered addition to conventional broadcasts. This rapidly developed into a much broader system which became today's Dirac.

Within the BBC, Tim Borer is the project leader for the Dirac Project, and Thomas Davies is the guru who understands the algorithm in more depth than most.

http://dirac.sourceforge.net

Tim Borer
tim.borer@rd.bbc.co.uk

Thomas Davies
thomas.davies@rd.bbc.co.uk

Peter Wilson
Peter.wilson@hddc.co.uk

Stuart Sommerville
Stuart.sommerville@numediatechnology.com

DIRAC

the video compression family using open technology

Are you paying out very large amounts of money for Codec license fees? Are you in search of the ideal video compression format? Do you find that one of your biggest challenges is the myriad of proprietary media formats? Does the stringent and complex licensing environment put you off making decisions? Do your customers operate on a range of different platforms?

...if the answer is 'yes', then read on...

Dirac is a general-purpose video compression family suitable for everything from internet streaming to HDTV.

Dirac achieves state of the art performance – good quality at low bit rates, leading to lower costs.

Dirac is an open technology – removing licensing costs on software, hardware and content flow. Dirac's technical flexibility offers a versatile package, facilitating ease of operation over many applications and therefore saving money.

Now we are announcing the launch of Dirac PRO – a low delay version of Dirac which offers near-lossless modes for improved storage and which allows the reuse of existing plant.

CONTENTS

What is Dirac?

Dirac is a very versatile compression family. It includes a range of tools which gives flexibility in performance to match the environment.

Dirac has the capability of compressing high resolution files for production, compression for broadcast content, and compression for web 2.0 applications. Compression can be either lossless or visually lossless, or it can exploit lossy compression using long-GOP formats for distribution. The compression efficiency is similar to that of AVC/H264 but without the encoding complexity or licence burden.

Most of the normal formats are built into Dirac as default settings. For the more adventurous, most parameters can be varied simply and effectively.

- Direct support of multiple picture formats 4 K E-Cinema through to QCIF.
- Supports I-frame only up to long GOP structures.
- Direct support of multiple chroma formats, e.g. 4.4.4/4.2.2/4.2.0.
- Direct support of multiple bit depths, e.g. 8 bit to 16 bit.
- Direct support of interlace via metadata.
- Direct support of multiple frame rates from 23.97 fps to 60 fps.
- Definable pixel aspect ratios.
- Definable 'Clean Area' for inputs within larger containers.
- Definable signal ranges and offsets.
- Multiple colour spaces with metadata to describe:
  - Colour Primaries
  - Colour Matrices
  - Transfer Functions
- 32 bit frame numbers (more than two years at 60 fps) in both I frame only and Long-GOP.
- Definable wavelet depth.
- A choice from multiple wavelet filters (including filters optimised for down-conversion).

Applications

Dirac suits a wide range of applications from 2 K formats used for E-Cinema through to QSIF found on hand-held devices.

The new Dirac PRO offers high quality and low latency, with the high bandwidth operations found in production and post production, archiving and contribution links.

Dirac offers high compression for narrow bandwidth environments such as broad-casting and internet downloads, podcasting, peer-to-peer distribution and access services.

Whatever the application, it is possible to select parameters which offer the solution you need.

- Clip Distribution.
- Live Streaming Video.
- Pod Casting.
- Creative Archive.
- Peer to Peer transfers.
- HDTV with SD Simulcast capability.
- Higher density channel packing.
- Desktop Production.
- News links.
- Archive Storage.
- Digital Intermediate Film Out file storage.
- PVRs.
- Multilevel Mezzanine 3 GBit/s into 1.5 GBit/s, 1.5 GBit/s into 270 MBit/s, etc.

Products

The main specification is now written and available on SourceForge, the Open Source website.

Quality portable software to implement Dirac is being delivered both on the Dirac and Schrödinger areas of SourceForge.

In hardware, we are starting to see the first specialist products, such as low latency, high quality codecs from Numedia Technology Ltd.

Dirac PRO

Dirac PRO is new at IBC. Dirac PRO is specially developed for high end production. Dirac PRO is a natural development of Dirac and is designed to fill a need for light compression and low latency in Digital Cinema mastering, broadcast production, studio operations, post production and archive.

Dirac PRO is very flexible and can support many image formats, sampling formats and colour spaces. For higher resolutions such as 4 K, a layering process can be used such that a 2 K base layer can be the norm with an enhancement layer. Sampling formats can be 4.4.4, or 4.2.2; bit depth can be chosen up to 16 bit quantisation, and metadata can signal the colour coordinates in use.

The main difference between Dirac and Dirac PRO is in the treatment of the final process in compression – the arithmetic coding. Arithmetic coding is processing intensive and introduces delay. These are features that are undesirable in high end production work. The arithmetic coding produces most efficiency savings with highly compressed material. There is little benefit to be gained with the low compression used in top-end production. Dirac PRO therefore omits the arithmetic coding.

What we're demonstrating

We are demonstrating Dirac in its different profiles:

- Dirac for transmission.
- Dirac for production and post.
- Dirac for low-latency compression.

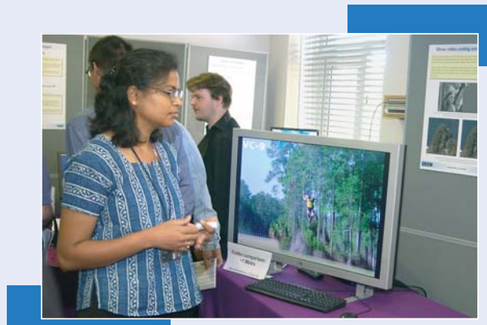

The main demonstration shows pictures from the BBC's recent trailers for its high definition broadcasts, compressed to around 18 MBit/s. The same Dirac compression could be used for any format from D-Cinema Post Production to hand-held.

Three specialised applications show how different parameters can be selected to provide optimum properties for different areas of production.

- The second specialised application is a low compression option designed to deliver nearly lossless coding. This lets us deliver 1080 progressive formats over infrastructure designed for 1080 interlaced.
- The third is an option to compress a 1.5 GBit/s studio interface to 270 MBit/s to allow the use of existing infrastructure for HDTV contribution or distribution.
- Two of these specialised applications are based on developments of MPEG technology, but which have been simply adapted to use the power of Dirac.

Dirac

The demonstration of Dirac itself shows how it works with high definition. We have chosen some of the most interesting sequences from the BBC's recent trailers for its high definition broadcasts. These sequences were originally artistically chosen to sell high definition to our public. This meant that the sequences were full of impact and detail. In technical terms, they present a challenge for any coding system.

We are deliberately not going to do a direct comparison with other coding systems. We do our own comparisons in house, and find that Dirac offers similar performance to the best.

Where there are compression artefacts visible, those in Dirac are much more benign. The block artefacts of DCT-based systems are much more easily picked up by the eye than the wavelet artefacts seen in Dirac. For this reason, some of the technical metrics such as PSNR measurements are not directly relevant.

Remember, the system that is seen on the stand can provide great performance with other formats. It supports post production and archive applications at up to 4 K resolution 16-bit RGB/YUV/XYZ. For distribution Dirac achieves comparable or better compression efficiency to the best proprietary or patented codecs from HDTV to QSIF resolutions.

Dirac for production and post

High-end HDTV production is rapidly migrating to high quality 1080 P50/60. But this format requires a higher data rate than the existing 1.5 GBit/s HDSDI infrastructure. The Dirac PRO profile supports the transport of these high quality images over conventional high definition infrastructure. This is an evolution of work on Mezzanine coding (aka SMPTE VC-2) which originally used DCT as the transform. Now we are embracing the concept within Dirac PRO and using wavelets as the prime compression tool. It is also suitable for quality coding of video for 270 MBit/s links.

Typical applications may be lossless or visually lossless compression for archives or mezzanine compression for reuse of existing equipment, such as 1080P 50 Hz carried in a 1080I 25 Hz channel.

For existing standard definition links, compressing 1.5 GBit/s HDSDI links into 270 MBit/s SDI or SDTI would facilitate the use of standard infrastructure for routing HD signals. Likewise, compressing HDSDI signals to be carried on Gigabit Ethernet (at circa 600 MBit/s) would also allow high definition working on a cheaper network infrastructure. Dirac PRO gives excellent quality at these levels of compression.

The demonstration shows high definition progressive signals compressed to fit a 1.5 GBit/s channel, as would be used for an SDI interface. The processing is carried out live using hardware from NuMedia Technology which is displayed on the stand.

Low latency compression

Low delay video coding is needed for live production, especially for wireless applications. Any system needs to be capable of integration into a equipment which exhibits a small footprint and uses little power.

The demonstration shows a prototype of such a system. This is working live, using a camera on the stand, so visitors can see the very low latency we achieve. This makes the system ideal for adoption in live broadcasts.

Technology

The original Dirac development has three main strands:

- A compression specification for the bytestream and the decoder.
- Software for compression and decompression.
- Detail which allows efficient hardware development.

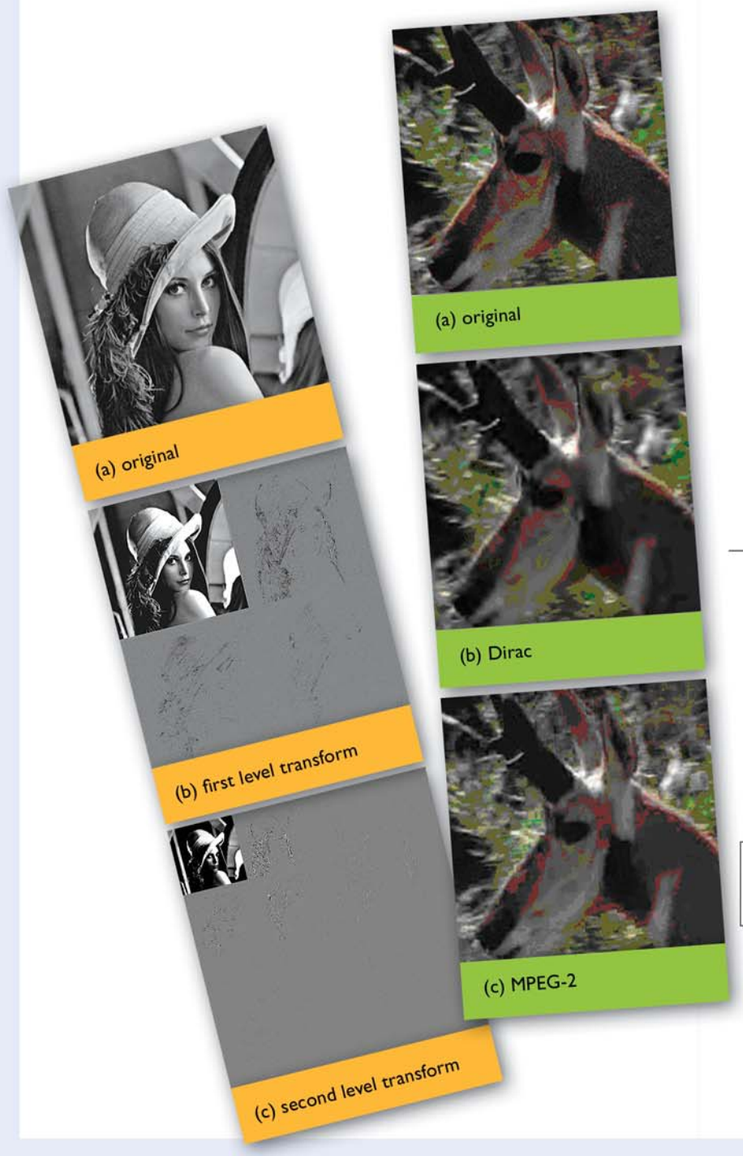

Unlike many of the other standard video compression systems, the software is not intended simply to provide reference coding and decoding, but also as a prototype implementation that can freely be modified, enhanced and deployed. The decoder implementation in particular is designed to provide fast decoding whilst remaining portable across software platforms.

Real time decoding of modern compression systems is difficult without extensively exploiting hardware support (in coprocessors and video cards) or assembly-language code, but these features can easily be added to Dirac's modular codebase.

Architecture

Dirac is similar to many of the established video coding systems. However we have adopted established technologies which combine effectiveness, efficiency, and simplicity. Together, these features give us a quality system which is not encumbered with patents.

First we use motion compensation to make use of the correlation between picture frames. Good motion compensation can dramatically reduce the amount of data required to code a picture.

Then we use wavelets (not the more conventional DCT) to transform the residual error signal.

Motion Compensation

In Dirac, frames have two essential properties. Firstly, they are either predicted from other frames (Inter) or not (Intra). Secondly they can be used to predict other frames (Reference) or not (Non-reference). All combinations of these properties are possible, and any Inter frame can be predicted from up to two reference frames. This means that Dirac can support conventional MPEG-style structures (Group of Pictures or GOP), but also any other prediction structure that may give better performance.

When we get down to pixel level, we can define the reference pixels for motion compensation, either at the global level (through pan, tilt, zoom, rotate, etc., commands) or by reference to pixels chosen by calculating specific motion vectors for the local block of pixels.

Transform coding

Wavelets have commonly been used for still image compression (a recent example is the core of JPEG2000). Now the power of modern chips allows us to use wavelets for motion pictures. The wavelet transform repeatedly filters the signals into low and high frequency parts. This repeated split concentrates the important data in one subband which can be efficiently encoded. We apply different degrees of quantisation to the transformed data. The human eye appears to be insensitive to coarse quantisation in some of the higher wavelet bands, and we exploit this ruthlessly to achieve high compression efficiency.

One of the weaknesses of MPEG-2 is the way that the picture goes all blocky when the coder is being worked hard. The use of the Discrete Cosine Transform (DCT) to transform the residual error limits the flexibility of the blocks used in the processing. By using wavelets, we can use varying sizes of blocks, and overlap them to mitigate the impact of block edges. This block structure also results in better motion predictions, again yielding improved compression.

Entropy coding

The transformed data still has redundancy. Entropy coding is used to reduce the band-width. The entropy coding technique used in Dirac is arithmetic coding. This is efficient and flexible. Arithmetic coding separates statistical modelling from the compression process itself, and better compression is afforded when the inter-dependence of data is exploited by switching between models based on previously-coded data.

Dirac applies entropy coding to the motion vectors and the output of the wavelet transform process.

Bytestream

The whole of the compressed data is packaged in a simple bytestream. This has synchronisation, permitting access to any frame quickly and efficiently - making editing simple. The structure is such that the whole bytestream can be packaged in many of the existing transport streams, such as MPEG, MXF, IP, Ogg, etc.

This feature means that we are able to use a wide range of sound coding options, as well as easy access to all the other data transport systems required for production or broadcast metadata.

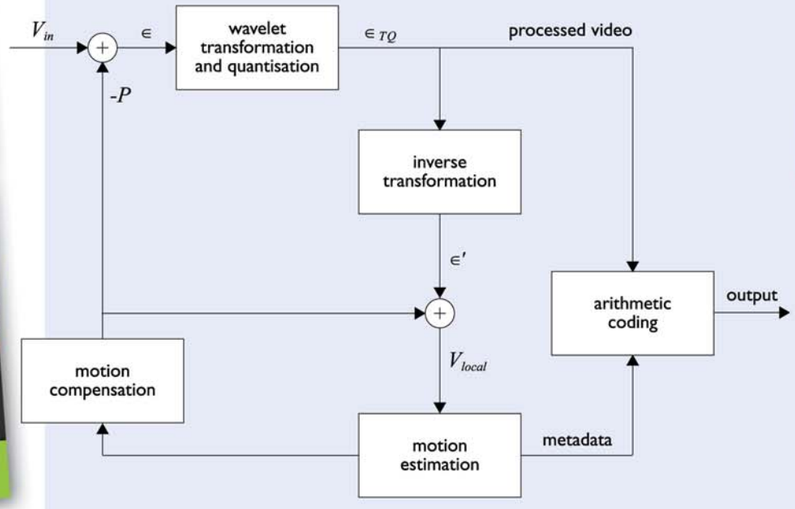

The open environment

So, Dirac is free – what does that mean in practice?

Given the surfeit of other codecs why is there a need for another codec, such as Dirac, and what does it provide to broadcasters that the others don't? The key feature of Dirac, in contrast to the other codecs, is that it is Open Technology. Dirac has been designed to avoid patent infringement. This means it may freely be used by anyone without the payment of royalties. This may seem like a minor issue but it could have a profound impact on the uptake of digital technology and, particularly the way it is used by public service broadcasters.

It is worth clarifying the difference between 'Open Standards' and 'Open Technology'. The two concepts are often confused. Public service broadcasters have always preferred open standards to vendor-specific technology. They provide interoperability and a competitive market place. Open standards are published and may be read by anyone, but implementers and users typically pay royalties to the owners of patents embodied in the standard. So open standards can be proprietary in the sense that the technology is owned and you may have to pay to use it. Open Technology, on the other hand, takes openness a stage further and gives you the right to use the technology, for any purpose, without royalty payment.

What are the license conditions?

Dirac is released under the Mozilla triple licence (MPL). This is an Open Source licence. It also allows for relicensing.

The licence gives completely open and non-discriminatory access to the technology to all comers, provided that they are prepared to respond in kind with their work.

Are you going to charge for Dirac?

No. The terms of the licence mean that as far as the BBC is concerned, there will be no charges or royalties for the Dirac software. As with all Open Source software, you can access the tools freely. However it is possible that vendors may put together specific packages with additional support as a commercial offering for those who may not have the full skills required to set up the basic code themselves.

Do the BBC have patents in Dirac?

Yes. We have patent applications in train for some of the techniques involved in Dirac, and others that we intend to put into Dirac in the future. There has been some criticism about this, so we'll be clear: this does not affect the Open Source status of Dirac, nor does it affect its royalty-free status. The conditions of the MPL mean that we're licensing these patents for use within the Dirac software for free.

Do you infringe any patents? Are there any royalties to pay?

The short answer is that we don't know for certain, but we don't think so. We haven't employed armies of lawyers to trawl through the tens of thousands of video compression techniques. That's not the way to invent a successful algorithm. Instead we have tried to use techniques of long standing in novel ways. Where we think we're novel, we're in the process of getting patent protection ourselves, which will invalidate others' claims of priority. There are some areas that are more heavily patented than others. Arithmetic coding is one such, even though the technique itself has been around for 25 years. We're keeping an eye on the situation, and we'll adopt alternative techniques if we have to.

What is the BBC's attitude towards Open Source?

The BBC is public body, whose remit is defined by its legal Charter and associated Framework Agreement with the British Government. One of the most relevant provisions is that the BBC 'must pay particular attention to the desirability of supporting actively in national and international forums the development of open standards'.

With this political guidance, many of the BBC's latest developments have been released as Open Source projects. This gives everyone an opportunity to experience the benefits of the technology, and has benefits as others have been joining in the development work, helping to create products which are much more functional, and less experimental.

Sources of information

The technology

The foundation of Dirac are the algorithms. These have evolved over the development phase of the project, as the team recognised that we could enhance performance by subtle changes. Now, with the release of Version 0.6 of the specification, we expect that there will be no more major changes.

The main source of information about the specification, and examples of code which can be used to try out the system is SourceForge.

Our co-workers are starting to produce products based on the specification. As this work expands we expect more information to be available from them.

There are also related projects whose work is directly relevant to the use of Dirac. Using the complementary MXF (Material Exchange Format) wrapper or the AAF (Advanced Authoring Format) wrapper very flexible very high quality systems can be constructed from a simple core.

Products

Schrödinger
sourceforge.net/projects/schrodinger
In parallel to the Dirac pages on SourceForge, there is also Schrödinger:

This is a project led by one of our collaborators, Fluendo, implementing the Dirac video codec in ANSI C code. It is meant to be highly optimized and portable to a wide range of platforms.

This software is therefore a good source for the technologist interested in practical experience with Dirac.

NuMedia Technology Ltd
www.numediatechnology.com
Numedia Technology Ltd is a manufacturer of high-technology electronic equipment for the broadcast and digital cinema markets. The Chameleon Image Processing Platform developed by Numedia is very well suited to the application of DIRAC algorithms that require real-time processing at high-definition and D-Cinema rates.

Numedia have also developed a new range of cost-effective modules to support the DIRAC PRO mezzanine profiles, which utilise existing infrastructure to transport high definition images with visually lossless compression.

High Definition & Digital Cinema Ltd
www.hddc.co.uk
HDDC are consultants in high end Digital Cinema, E-Cinema and the HDTV production and distribution process. Their role is assisting the BBC with extended roadmap planning and dissemination of information.

HDDC have over twenty year experience in HDTV and eight years experience of Digital and Electronic Cinema applications.

HDDC have recognised that Dirac provides solutions to Commercial and Technical problems which no other coding system can match.

The website for Dirac is **sourceforge.net/projects/dirac**

This site has the new specification, and working software to try out the system.

Many of the original ideas for Dirac came from the BBC Research & Development Department. This department has just undergone a reorganisation and is now being branded BBC Research within BBC Technology Group. The website **www.bbc.co.uk/rd** is a mine of information about broadcast technology and video processing in particular. As well as links to the Dirac project, you can find information about video coding and motion compensation in reports and white papers there.

Licences

The Dirac software and code are licensed under the Mozilla Public License, Version 1.1. A plain text version of this is available in **dirac.sourceforge.net/licenses.html**

I would really like to contact someone about Dirac

If you have any questions, the best route to ask questions is through the SourceForge pages. It may be that your question has already been answered there, or if not, others might be interested in knowing the answer too. Routing your query through SourceForge makes the process more inclusive and efficient.