Chromatec & Numedia

DIRAC CODEC

New forms of video compression

New forms of video compression

Why is there a need?

The move to increasingly higher definition and bigger displays means that we have increasingly large amounts of picture information to move, process and store. This applies not only for content to be consumed in the home but also in the emerging “D Cinemas”.
So these huge amounts of video information need to be moved through the production, post production, versioning, playout and multiple distribution (and no doubt, archive, repurposing and resale). Although the cost of storage is going down, this is at least balanced by the increasing demand for larger amounts of data storage.
Video compression has been used very successfully for many years as a way of reducing the amount of data being moved. But the number and complexity of “programme” paths to consumers via multiple distribution channels benefits from a rethink of how video compression is used.

Why Dirac?

Traditionally video compression technologies have been proprietary to specific manufacturers. Looked at in isolation these usually give excellent results. But concern has been expressed that when different systems are cascaded there can be concatenation effects, visible as a degradation of picture quality. This is an issue which particularly needs consideration in a production environment, where the signal may be coded and decoded many times across the programme path. The coding / decoding system used in Dirac and Dirac Pro has specifically been designed to minimise this type of effect. Indeed, in principle, the coder can be designed so there is no effect after the first generation. This system has been developed by BBC Research & Innovation using experience built up over tens of years and uses multi-level wavelet (as opposed to convention DCT) transforms to achieve optimum overall performance.

(a) original

(b) concatenation effect

The family of codecs has two parts,

- Dirac Pro (currently submitted to SMPTE for approval as VC-2) and
- Dirac.

Dirac Pro is intended for professional use in environments such as production and post production, where high quality (i.e. either visually or mathematically lossless) is needed and multiple coding / decoding processes are likely. This is an intra-frame codec, designed to accommodate normal production processes such as vision mixing, editing, etc.
The compression ratio can lie in the range 2:1 to 16:1. Latency for Dirac Pro through both coder and decoder is never greater than 2ms.
The Chromatec TR100,221 and 400 are all Dirac Pro products.
In contrast, Dirac is designed for distribution applications where there are no production requirements. Therefore this can also include motion compensated prediction across a number of frames, i.e. inter-frame compression, and therefore can achieve up to 100:1 compression and a much lower bit rate.

The benefits of Dirac Pro

Dirac Pro compresses each HD-SDI input by approximately 2:1 and multiplexes the compressed bit streams. The resultant bit stream is formatted as an HD-SDI (SMPTE 292M) signal that is fully compliant and compatible with existing HD-SDI equipment, such as switchers of routers. Any metadata, and up to 16 channels of embedded AES / EBU audio, from each source are kept without compression. Unique to the Dirac Pro implementation by NuMedia, the encoded bit stream includes a coarsely quantised interlaced version of each picture. This enables a recognisable image to be recognised as an HD-SDI signal and displayed.

So what is in the name, Dirac

Origin of the name "Dirac" is a reference to Paul A. M. Dirac, British physicist and winner of the 1933 Nobel Prize in Physics. The prize was shared with Erwin Schrödinger, for whom the Schroedinger implementation of Dirac was named.

The technology behind Dirac

This is a simple introduction to the technology of Dirac. It is hopefully of use to those who are new to the subject, or just want a simple overview.
Dirac has three main strands:

- a compression specification for the bytestream and the decoder
- software for compression and decompression
- algorithms designed to support simple and efficient hardware implementations, this is were NuMedia and Chromatec are involved.

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.

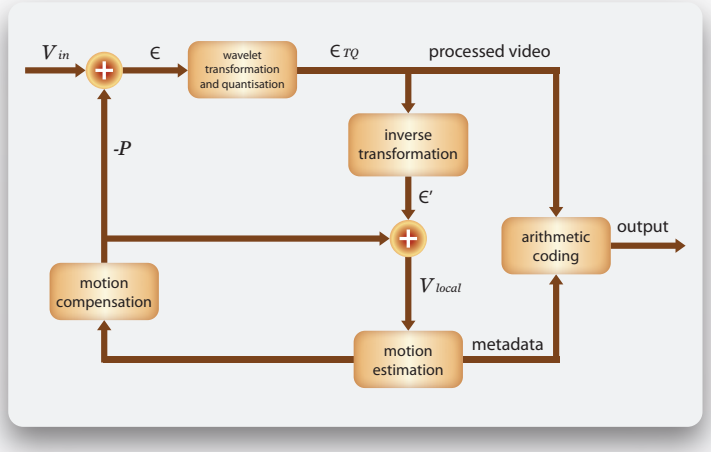

Motion Compensation

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 bandwidth. 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.

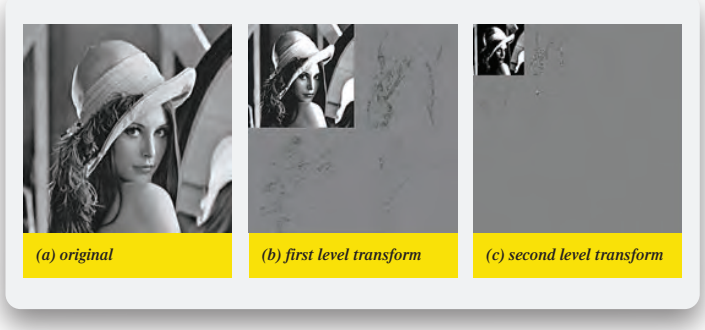

(a) original *(b) first level transform* *(c) second level transform*

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.

Dirac Applications

A few of the possible applications include:

- Moving HD over existing tech infrastructures
- Sending two camera channels over existing links
- Desktop production over IP networks
- Reducing storage in archive
- Reducing disk storage and bandwidth in D Cinema

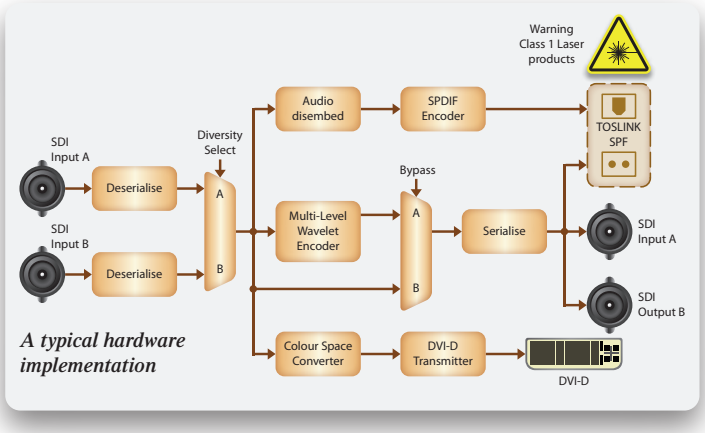

A typical hardware implementation

Using Dirac Pro

As mentioned above, Dirac Pro was designed with production and post production in mind. It supports a wide range of picture formats including up to 4K and 8K D Cinema images and beyond. It is not limited to 8 or 10 bit signals and can support the large word widths needed in the post production environment. All frame rates plus RGB, 4:4:4, 4:2:2 and 4:2:0 colour formats are supported.
The Dirac Pro stream includes essential metadata that describes how the coded signal should be interpreted.

Integration

The TR100, 221 and 400 codecs are part of the Chromatec range of openGear products. They can be housed in the Chromatec xxx mains powered rack frame (8310) or twinBOX along with any other module compliant with the openGear standard.

Applications

TR100 allows the transport of 3G HD-SDI over 1.5G infrastructures. Either via copper (Coax) and or fibre.

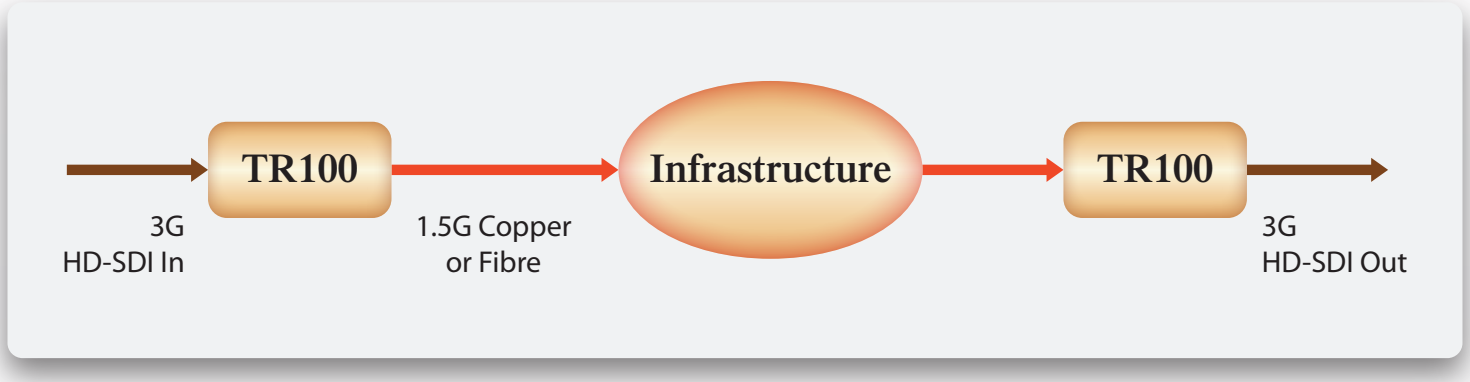

TR221 allows two 1.5G HD-SDI signals to be passed over single 1.5G infrastructures. Either via copper (Coax) and or fibre.

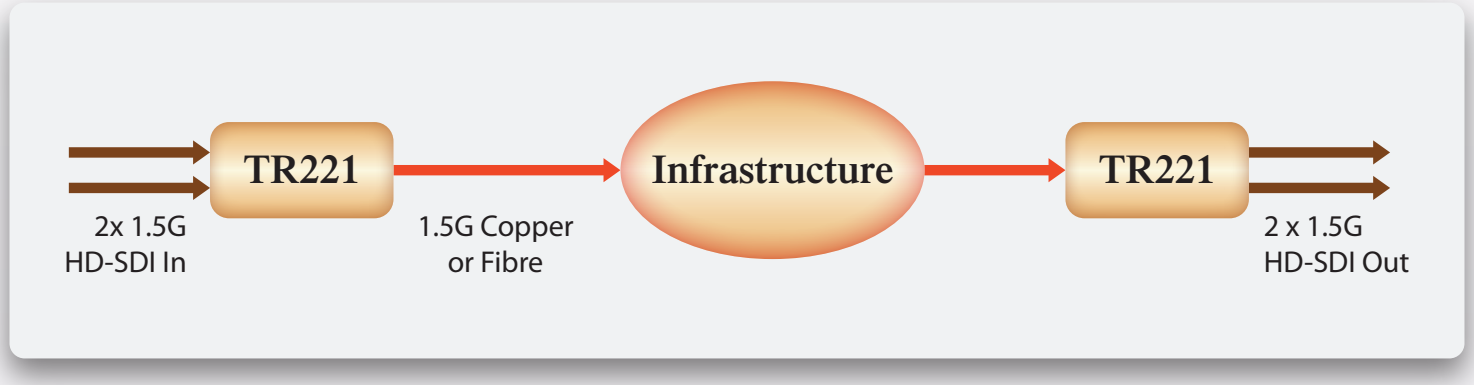

TR400 allows one 1.5G HD-SDI single to be passed over single 270Mb infrastructures. Either via copper (Coax) and or fibre.

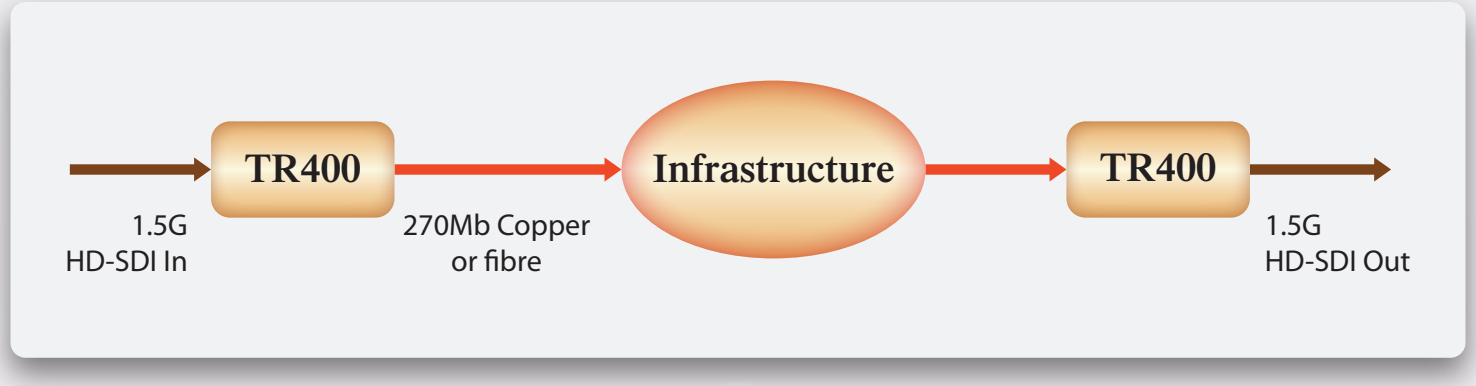

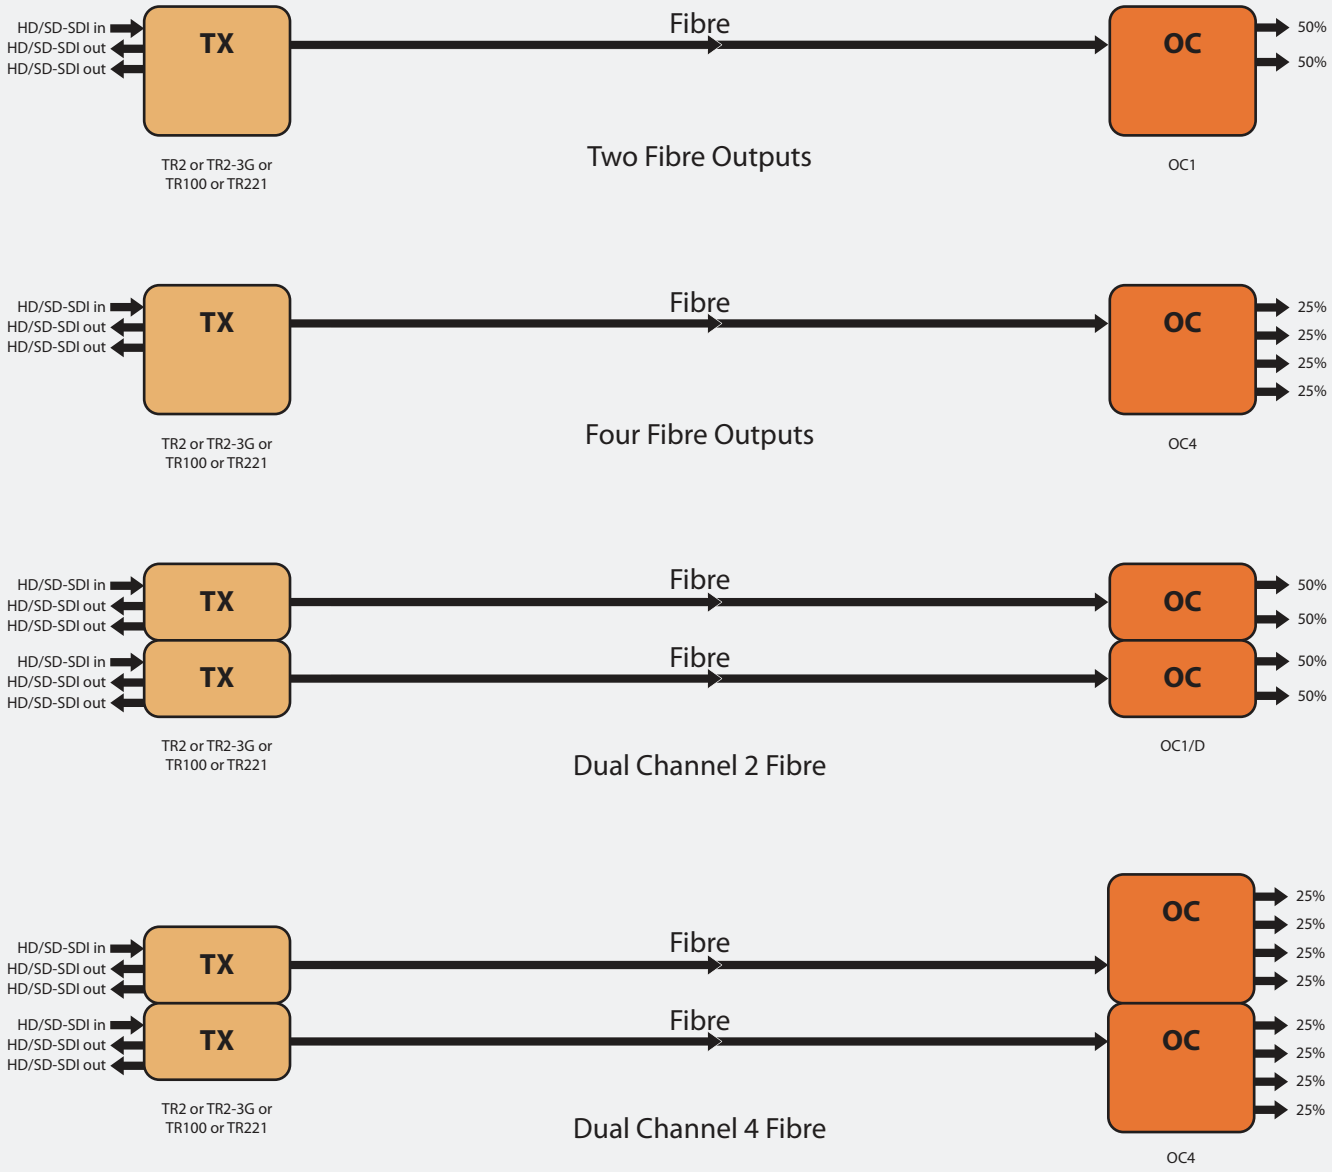
HD/SD-SDI in
HD/SD-SDI out
HD/SD-SDI out
TX
Fibre
OC
50%
50%
TR2 or TR2-3G or
TR100 or TR221
Two Fibre Outputs
OC1
HD/SD-SDI in
HD/SD-SDI out
HD/SD-SDI out
TX
Fibre
OC
25%
25%
25%
25%
TR2 or TR2-3G or
TR100 or TR221
Four Fibre Outputs
OC4
HD/SD-SDI in
HD/SD-SDI out
HD/SD-SDI out
TX
HD/SD-SDI in
HD/SD-SDI out
HD/SD-SDI out
TX
Fibre
Fibre
OC
OC
50%
50%
50%
50%
TR2 or TR2-3G or
TR100 or TR221
Dual Channel 2 Fibre
OC1/D
HD/SD-SDI in
HD/SD-SDI out
HD/SD-SDI out
TX
HD/SD-SDI in
HD/SD-SDI out
HD/SD-SDI out
TX
Fibre
Fibre
OC
OC
25%
25%
25%
25%
25%
25%
25%
25%
TR2 or TR2-3G or
TR100 or TR221
Dual Channel 4 Fibre
OC4

TXTX / Single Fibre with WDM / RXRX

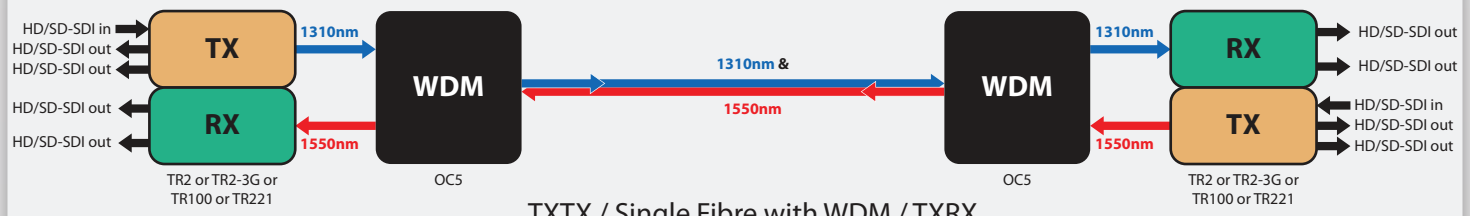

TXTX / Single Fibre with WDM / TXRX

2 x TXTX / Dual Fibre with WDM / 2 x RXRX

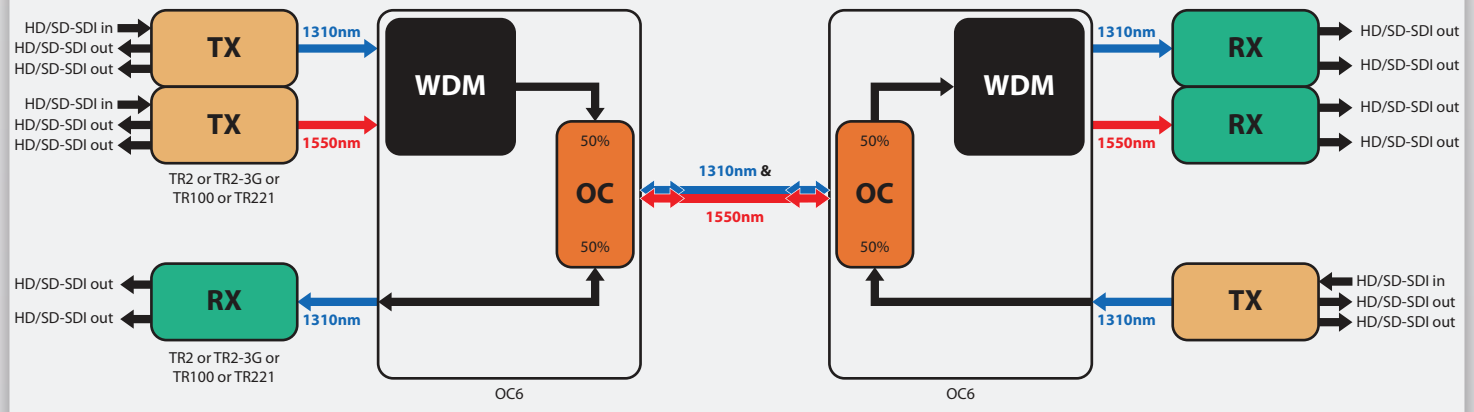

2 x TXTX / Dual Fibre with WDM / 2 x RXRX

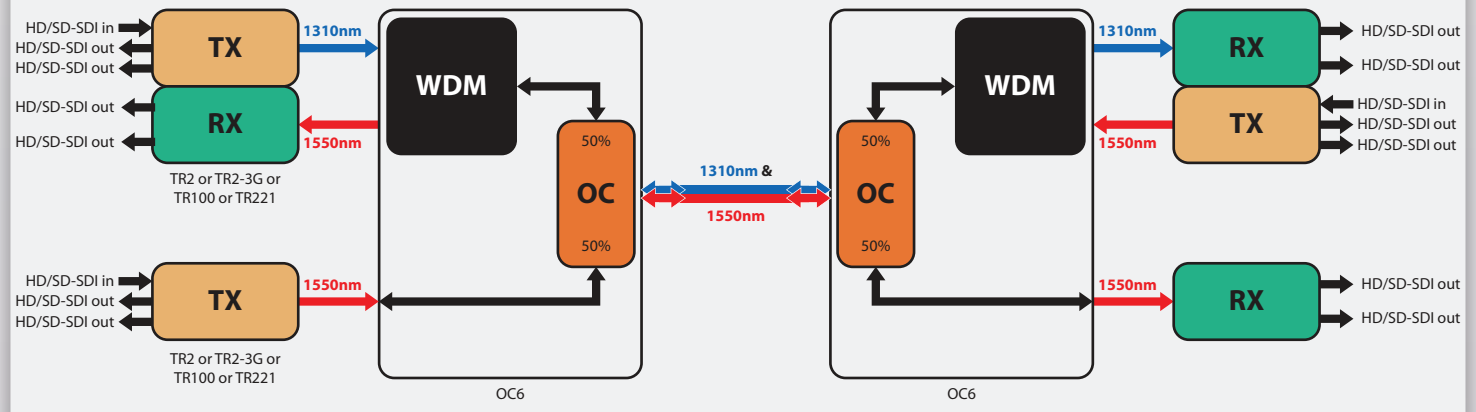

2 x TX1 x RX / Single Fibre with WDM & OC / 2 x RX x TX1

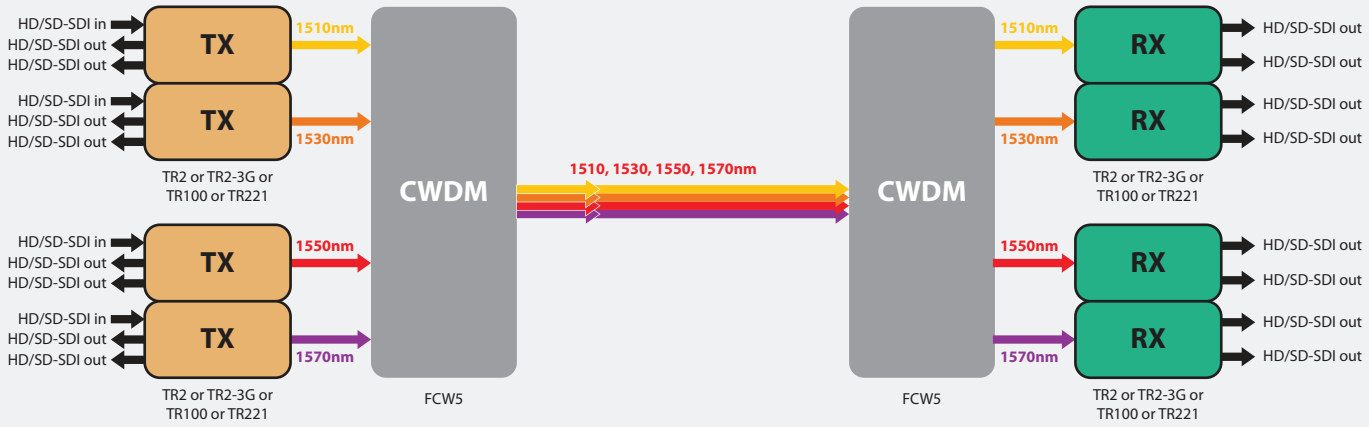

4 x TX/Single Fibre with CWDM/ 4 x RX

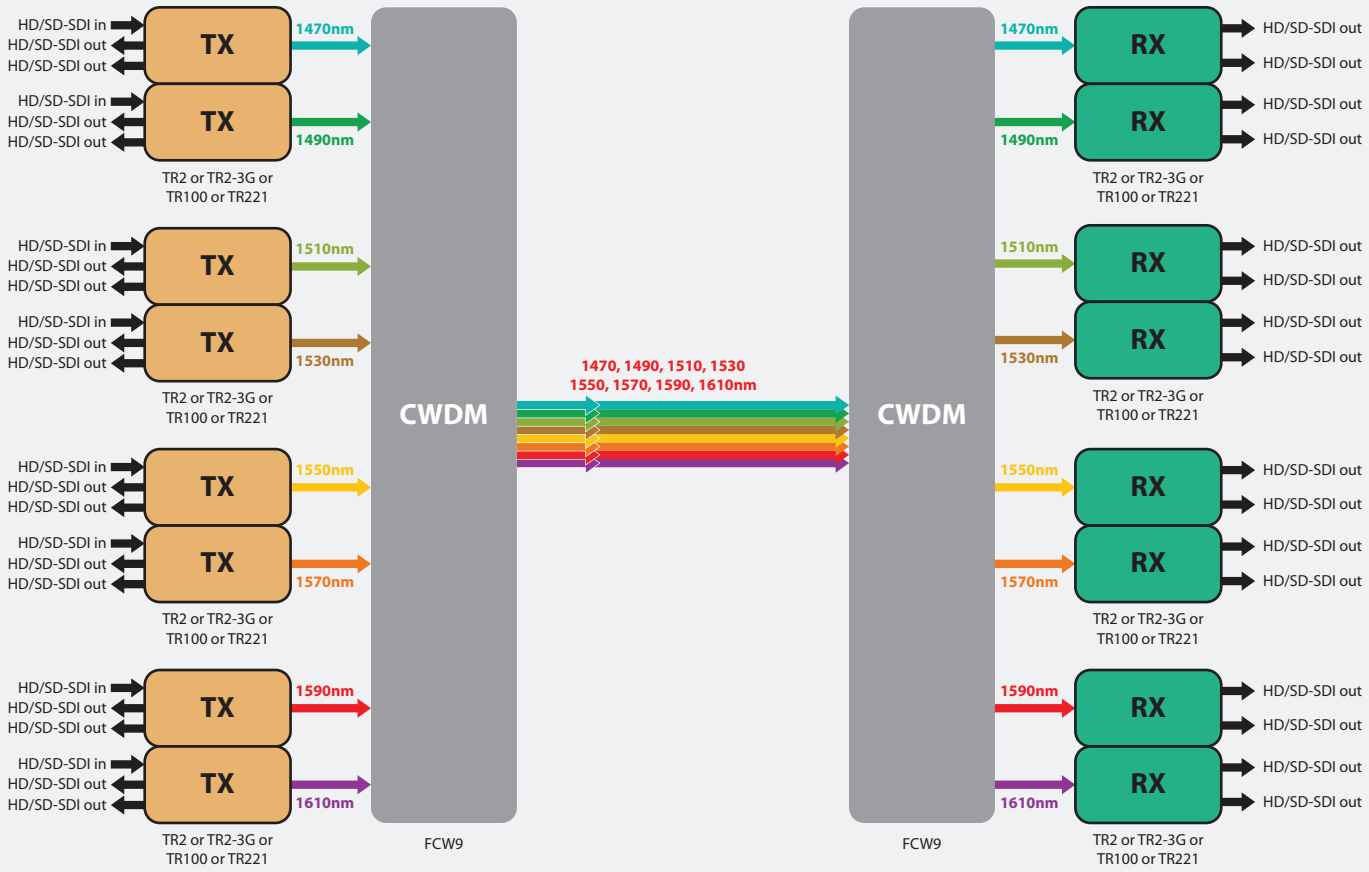

8 x TX / Single Fibre with CWDM / 8 x RX

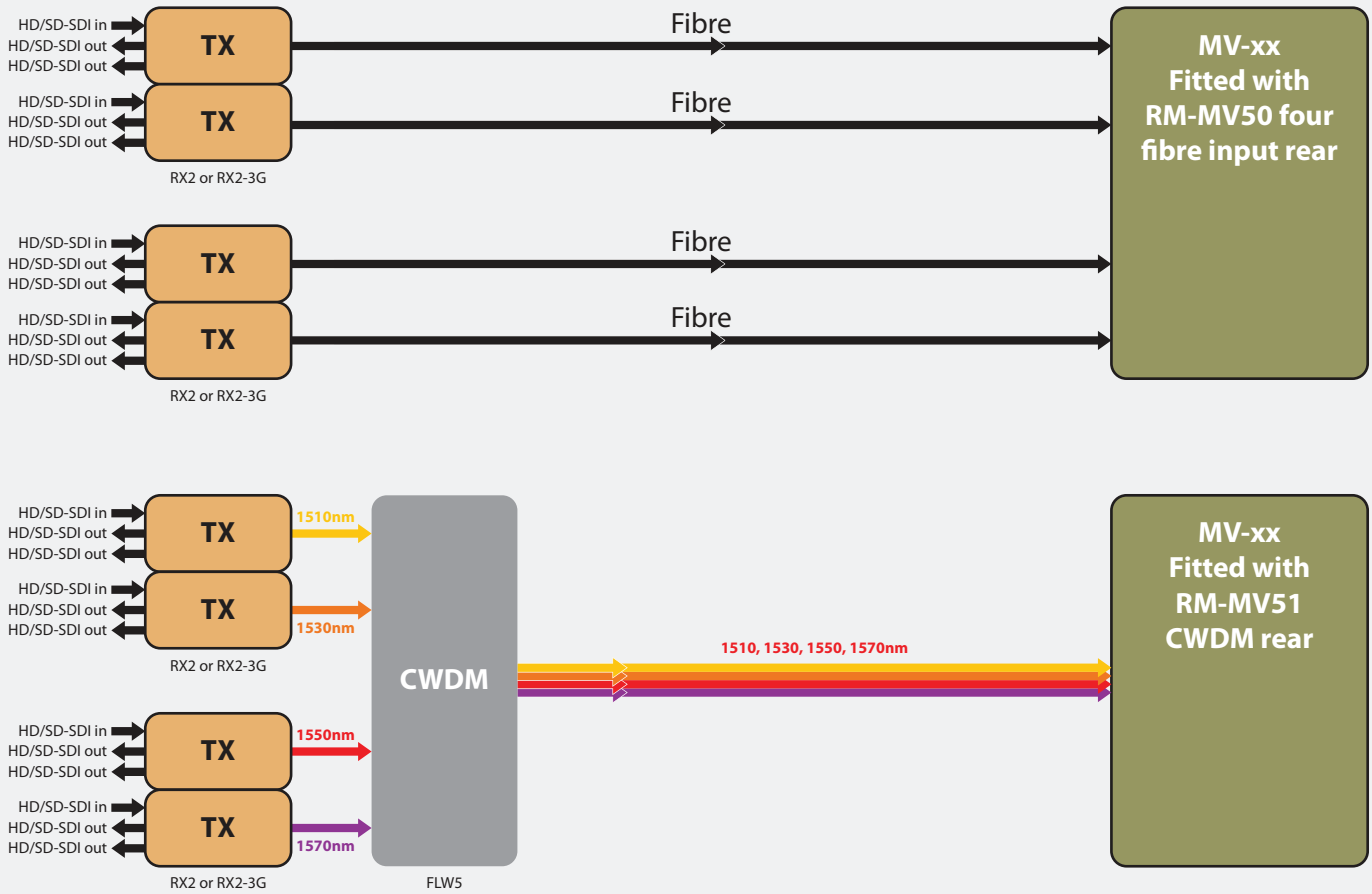
HD/SD-SDI in
HD/SD-SDI out
HD/SD-SDI out
TX
HD/SD-SDI in
HD/SD-SDI out
HD/SD-SDI out
TX
RX2 or RX2-3G
Fibre
Fibre
MV-xx
Fitted with
RM-MV50 four
fibre input rear
HD/SD-SDI in
HD/SD-SDI out
HD/SD-SDI out
TX
HD/SD-SDI in
HD/SD-SDI out
HD/SD-SDI out
TX
RX2 or RX2-3G
Fibre
Fibre
HD/SD-SDI in
HD/SD-SDI out
HD/SD-SDI out
TX
1510nm
HD/SD-SDI in
HD/SD-SDI out
HD/SD-SDI out
TX
1530nm
RX2 or RX2-3G
CWDM
1510, 1530, 1550, 1570nm
MV-xx
Fitted with
RM-MV51
CWDM rear
HD/SD-SDI in
HD/SD-SDI out
HD/SD-SDI out
TX
1550nm
HD/SD-SDI in
HD/SD-SDI out
HD/SD-SDI out
TX
1570nm
RX2 or RX2-3G
FLW5

Sample Dirac-Pro Applications

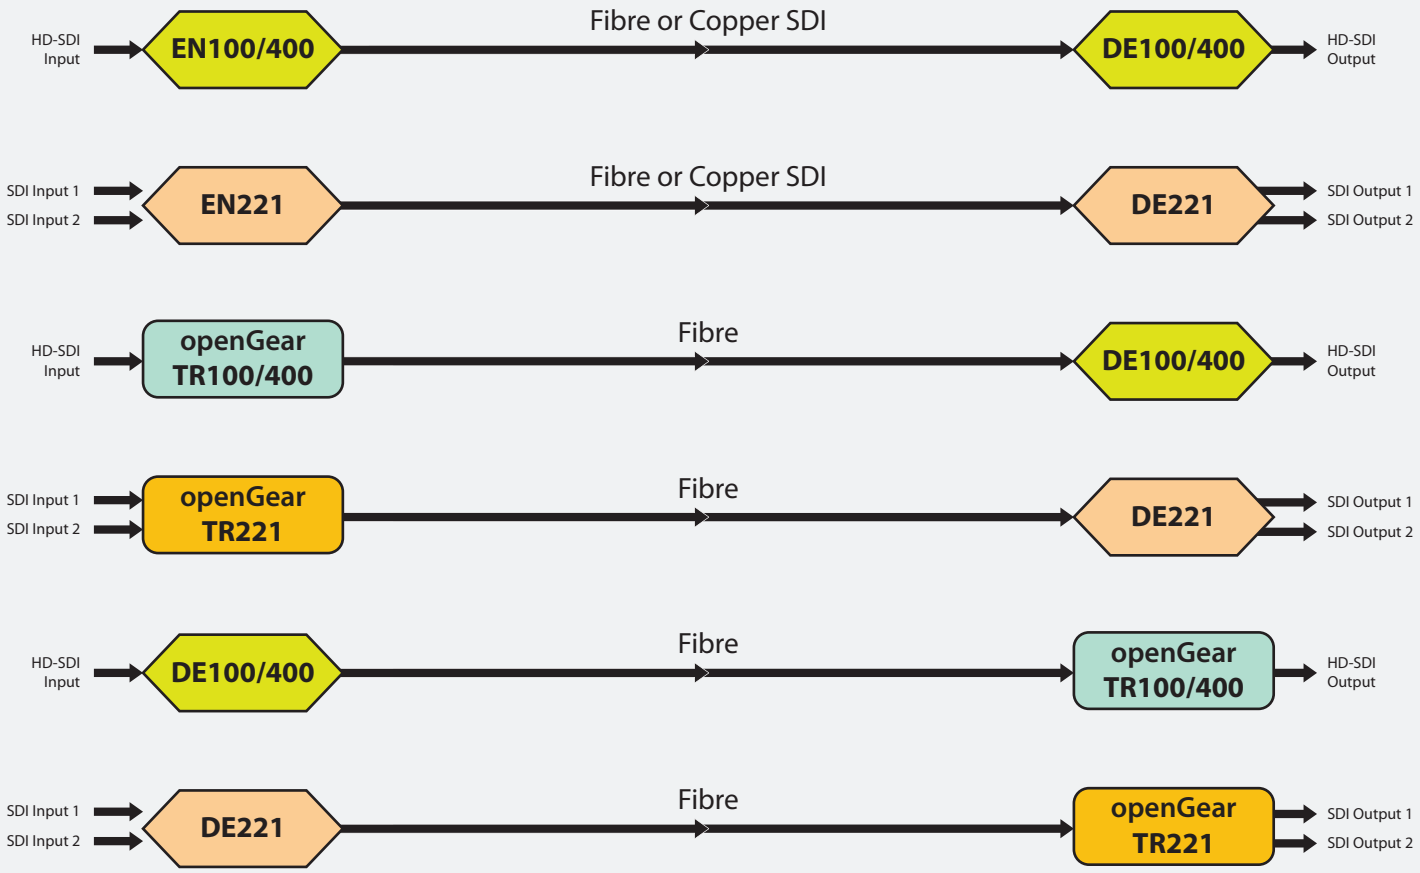

Four HD-SDI Signals with single fibre

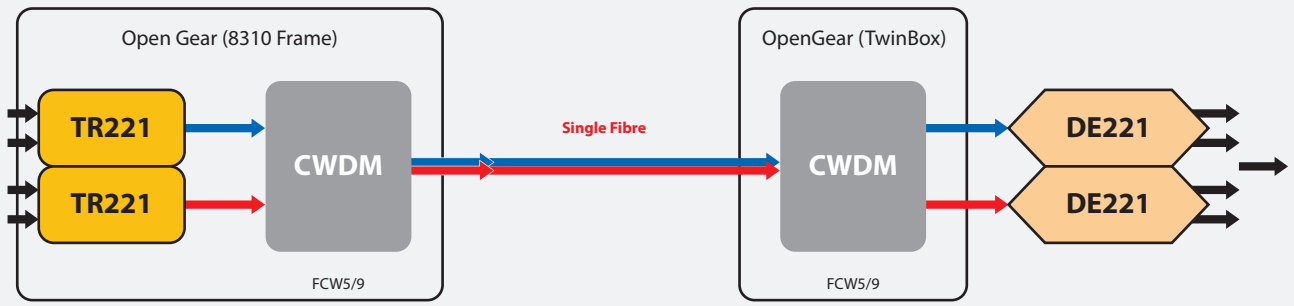

Note: *by combining TR221's and/or EN/DE221's and CWDM blocks up to 16 HD-SDI (1.5G) connections can be made down a single fibre.*

References and further reading

1. Gourier, G.G., 1957. Bandwidth Compression of a Television Signal. Proceedings I.E.E., Paper No. 2357 R, 104 B, pp 265-271

2. http://dirac.sourceforge.net/specification.html

3. Borer, T. & Davies, T. 2005. Dirac - Video Compression using Open Technology. EBU Technical Review No 303, July 2005

4. http://dirac.cvs.sourceforge.net/dirac/diracpro

5. http://www.numediatechnology.com

6. http://sourceforge.net/project/showfiles.php?group_id=102564

7. http://schrodinger.sourceforge.net

8. http://www.jpeg.org/jpeg2000

Acknowledgements

The author would like to thank his colleagues at the BBC and Numedia Technology for their contributions to the development of the VC-2 codec.
He would also like to thank the BBC for permission to publish this paper.

Michael Stevens & Partners Ltd
Invicta Works Elliott Road, Bromley, Kent, BR2 9NT UK
Tel: +44 (0)20 8460 7299 Fax: +44 (0)20 8460 0499
E-mail: sales@michael-stevens.com Web: www.chromatec.com Web: www.michael-stevens.com