

Background

The Curriculum Corporation Video Convergence in the Classroom Project consists of a series of whitepapers to be used as professional development resources. The whitepapers address the need for teachers to understand the principles within contemporary digital media; in particular, digital video. This will allow teachers to:

- make informed decisions when commissioning multimedia products and services
- effectively use current digital video production technologies, and incorporate digital video production into classroom activities
- understand the issues and opportunities associated with the online distribution of digital video
- consider the potential of digital video as a medium for the dissemination of their own professional development resources, and for communication with stakeholders including parents and school governing bodies.

The whitepapers have been developed for the Curriculum Corporation by convergent media specialist Michael Borthwick.

While each whitepaper can stand alone, the series as a whole explains:

- how computers can be used for the acquisition of analog material and its storage in digital form
- the problems associated with the storage and online distribution of large digital media files
- the characteristics of image-based digital media files that can dramatically reduce file sizes and mitigate distribution and storage problems
- techniques available for shooting and editing educational video content to maximise the quality of the compression process.

The Need for Compression

Useful background whitepapers

- *What is Digitisation*
- *Image Digitisation*

Further reading

- *Image Compression*
- *Video Compression*

Key concepts

- Multimedia files are large and consume lots of hard disk space. The files size makes it time-consuming to move them from place to place over school networks or to distribute over the Internet.
- Compression shrinks files, making them smaller and more practical to store and share. *Windows Media* and *QuickTime* are popular examples of compression formats for video clips.
- Compression works by removing repetitious or redundant information, effectively summarising the contents in a way that preserves as much of the original meaning as possible.
- Some compression processing may require a relatively fast computer to decompress footage, and can behave erratically on a slower system.

The mp3 phenomenon

Swapping *mp3* audio files has become an extremely popular, controversial activity for millions of Internet users. The *mp3* compression format squeezes digital audio files down to about a tenth of their original size making their exchange over the Internet practical, even for users with slow connections. *mp3* files are controversial because, unlike proprietary formats such as *RealAudio* or *Window Media* files, they can be easily written to an audio CD and distributed.

Introduction

Have you ever waited impatiently for a file to download from the web? This file may well have been a video clip of some sort. Video files contain a great deal of information and can be very large. Consequently, they are slow to download.

Consider this example:

Your students are participating in a project with another school and have shot a 5-minute video about the status of the project on a *DV camcorder*. When transferred to a computer for editing, this footage will consume around one *GB* (one *gigabyte* or 1000 *megabytes*) of hard drive space. How long would it take to send this clip using the fastest *dialup modem* available today? Five minutes? An hour? A day? In fact around two and a half days!

In contrast, files such as Microsoft Word documents, or scanned photos, can be sent quickly because they are relatively small and don't take too long to transfer. In comparison, multimedia files can be massive.

If only these files weren't so big then they wouldn't be such a headache. In fact files *can* be made more compact. This squeezing process is called, appropriately, *compression*. A compressed file requires a smaller amount of storage space and moves more quickly over a computer network, whether this is the school local area network (*LAN*) or the Internet.

Figure 1 shows the dramatic reduction in file size due to compression. On the left-hand side, we see that a single minute of full-screen *DVD* video weighs in at 37MB. *DVD* uses the *MPEG2* digital video compression format that was introduced in 1995. Compare this with the same full-screen video clip encoded with the latest *Windows Media* technology from Microsoft – it is about one-third of the size – revealing how far compression technology has advanced.

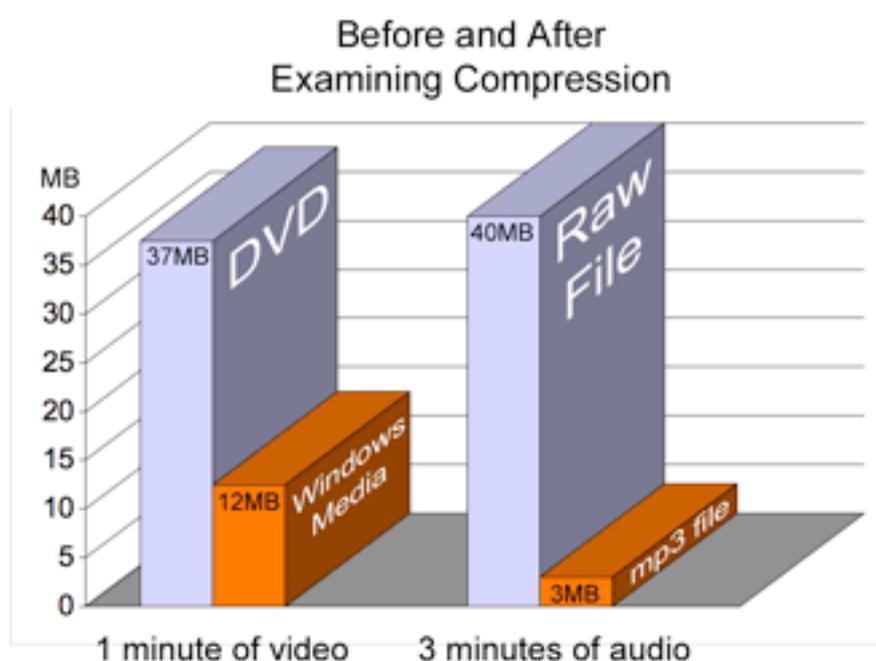


Figure 1 Reduction in file size due to compression

On the right, we can see that a 3-minute audio track, when compressed as an MP3 file, is only 3MB, less than one-tenth of the size of the original size.

Dehydrated files ... just add water

Compression is not that mysterious. If you've ever downloaded a file with the .zip suffix to your Windows PC or .sit on the Macintosh, then you've downloaded a file that has been subject to file compression in order to make the file more compact and minimise the download time. Software such as WinZip for Windows or Stuffit Expander for the Macintosh uncompress or reconstitute these compressed files back into *identical* copies of the original. Remember that you won't be able to open these files without these applications or other applications that support the compressed archives.

The fact that a digital multimedia file such as a video clip, sound file or image can be shrunk might make you wonder what was 'wrong' with the original file in the first place. Was it simply padded out with a lot of unnecessary information that can be thrown away?

In fact, compression *does* work by removing 'unnecessary' information (we're not talking here about cutting those wobbly shots of the grass when you forgot to turn the camera off – that's called *editing*).

Text documents can be compressed and yet retain their original essence; similarly, multimedia files are compressed by eliminating information whose absence will (ideally) go unnoticed. Typically, this is repetitive information or falls outside the range of human perception – in other words it is *redundant*.

More information about compression techniques, as they apply to still images, is contained in the *Image Compression* whitepaper. In the *Video Compression* whitepaper, this discussion is extended to *moving* images.

Key terms

Bandwidth
Camcorder
Compression
Dialup
DV
DVD
GB
Gigabyte
LAN
MB
megabyte
Modem
mp3
MPEG2
QuickTime
Redundancy
Windows Media

Further reading

- *Image Digitisation*
- *The Need for Compression*
- *Video Compression*

Key concepts

- In an analogue medium such as an LP record, information is recorded in the form of a direct physical correspondence between the medium and the event that is being captured.
- Digital representation involves measuring the original event at frequent intervals and storing those measurements as a string of digitally encoded numbers.
- The greater the range of numbers available, the greater is the accuracy and quality of the digitised result.
- Data imported into a computer takes up space; the more information used, the greater the storage space required.

Introduction

Picture in your mind a flower. In the centre, the petals are a rich purple. As you follow the petals towards their tips, the colour very gradually changes to a pale pink. This gradation of colour is incredibly subtle; at no stage is there an abrupt transition.

This flower is part of the inherently *analogue* world in which we live, a world composed of an infinitely varying palette of colours, sounds, tastes and textures. Our senses too are analogue, conveying sensation through electrical pulses and chemical reactions in sense organs such as the eyes and ears. These biological processes vary in accord with changes in the external world. There is an *analogous* relationship.

Similarly, the capacity of analogue media to capture an image of a flower or the ring of a bell relies on some form of direct physical change within the medium itself that reflects the character of the external environment.

A vinyl record is a good example of an *analogue* media format. The pattern of the grooves pressed into the surface of the record mimic the pattern of variation in air pressure associated with the original sound. When the record is played back, the sound wave is recreated for the enjoyment of the listener.

Digital media present a stark contrast. The development of digital electronic systems has profoundly changed the way information is captured, processed, stored and shared. Digital representation literally involves the use of *digits* or numbers in representation.

Where do these numbers come from?

The *digitisation* process produces a stream of digital information. This information comes from measurements of the original analogue phenomenon as shown in Figure 1.

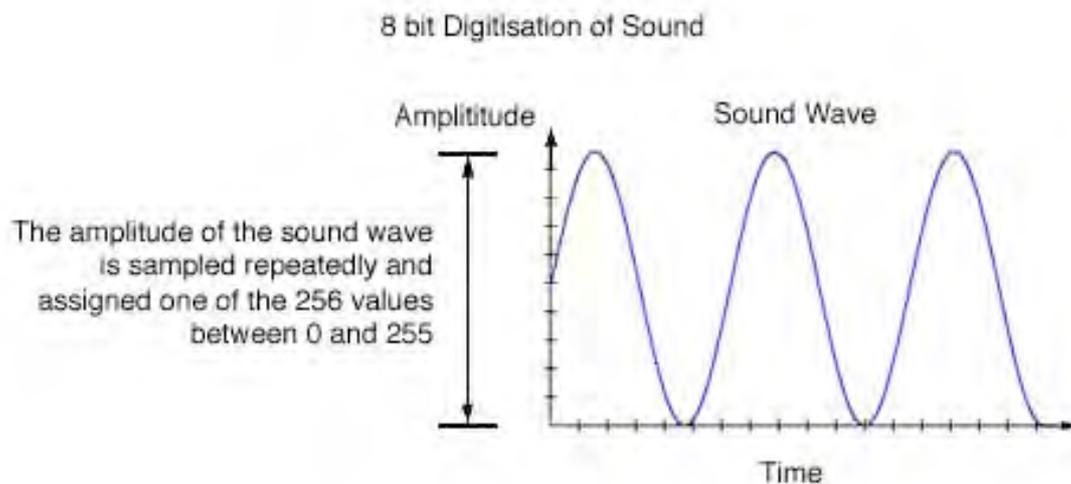


Figure 1 Measurements of the original analogue phenomenon

To represent this analogue sound wave digitally, we measure (or sample) the *amplitude* (height) of the sound wave at frequent intervals. In this case, we measure the amplitude on a scale that ranges from a minimum of zero to a maximum of 255. This gives a range that contains 256 separate intervals – a convenient number for digital storage because it corresponds to the information storage capacity of a single *byte* of computer memory.

Audio CDs store each sample using two bytes, yielding 16,384 potential values. This enables a much more fine-grained representation, and the reproduction of the original sound with a much higher fidelity. Since a byte consists of eight individual *binary digits* or *bits*, this is also described as *16-bit sampling*.

This correlation between the number of bits employed in every sample and the quality of the digitised result applies to all media, including images, as explored in the *Image Digitisation* whitepaper.

Once digitised, information can be manipulated and processed using computer programs. Digital content can also be distributed over networks – free from the costs and constraints associated with the creation and distribution of physical copies.

Remember that your computer has a limited amount of storage space! If you are capturing half an hour of high-quality footage, you will use up to 7GB of disk space. In the instance of editing and converting a DVD, you will be compressing this data to an appropriate DVD format, which will also use up disk space. Video editing can use up more disk space than editing other multimedia formats. If you are in doubt about your computer's capacity, talk to your IT support people.

Key terms

Amplitude

Analogue

Byte

Digital

Digitisation

Sampling

Distribution of Video

Useful background whitepapers

- *Video Compression*

Key concepts

- New video technologies have made the medium more accessible and enabled video to be more broadly incorporated into the curriculum.
- School and classroom video projects can be delivered on CD, DVD, online or videotape.
- The choice of delivery method involves trade-offs between cost and convenience, and requires consideration of the nature of the content and the audience.

Introduction

The historically high cost of video production equipment has tended to limit its use in schools.

The introduction of *Digital Video (DV)* cameras and computer-based editing systems has transformed both the *business* of video and television production and the accessibility of high-quality video production to individuals and households. Technological change has also introduced new distribution opportunities, ranging from the superb quality of DVD to the immediacy and reach of the web.

These technologically driven developments have the potential to positively impact schools, offering students and teachers the opportunity to work with this powerful medium in new ways and in new subject areas, and to share the results of their creative endeavours with unprecedented ease.

Creating DVD

The *DVD* format offers convenience and high-quality playback for those users with a DVD player or a computer with a DVD drive. The cost of blank DVD media continues to fall. Blank recordable *DVD-R* (DVD-Recordable) media is available for less than \$5. Commercial replication of DVDs is available and becomes cost effective if done in larger numbers.

A cheaper alternative to DVD that is still compatible with most DVD players is the *Video CD*. Video CD is a format popular in Asia for the distribution of movies on CD. It predates DVD and uses the MPEG1 video compression format. *MPEG1* was the first *MPEG* video format developed by the Motion Picture Experts Group. The higher quality (and higher *data rate*) *MPEG2* format is used on DVD.

If you create a Video CD, keep in mind that some older DVD players may be unable to playback recordable (*CD-R*) or rewriteable (*CD-RW*) discs. It pays to keep your audience in mind and undertake testing if practical.

CD-ROM video playback

The DVD format is highly standardised to ensure that DVD discs will play back on DVD players from any manufacturer. However, if the digital video files you are creating are destined for playback on a computer, perhaps as components of a larger interactive multimedia program, then there is less standardisation and a greater choice of formats.

These formats include *Windows Media*, *Real Video* and *QuickTime*, with each format requiring its own separate player.

While the older MPEG1 format is not particularly compact, it too can be placed on a CD-ROM and played back within both the QuickTime and Windows Media players, even on relatively old Macintosh and PC computers, an important consideration in schools.

Hardware and software

The cost of CD and DVD burners has declined rapidly and they are now built into most new computers.

DVD discs must be *authored* to conform to the strict international standards discussed above. Until recently, this has required the use of expensive DVD authoring software. Today, basic DVD authoring software is usually pre-installed on computers equipped with DVD burners or is included when a separate DVD writer is purchased. Your computer may have a dual, or double, layered DVD burner or even be able to burn DVD+R discs. Bear in mind that these may not play on all home players, so if you are authoring to this medium, always test your disc on a number of players.

Capacity of disc media

A CD has a capacity of over 600MB (megabytes) which is sufficient to store around one hour of low-quality full-screen MPEG1 video.

DVD-R and the rewriteable *DVD-RW* format discs currently have a capacity of 4.7GB (Gigabytes), allowing up to an hour of MPEG2 format video, depending on the quality level/data rate being used. Commercial mass-produced DVDs, which can be both double sided and double layered, provide sufficient space to store feature-length movies of several hours duration.

Online distribution



Figure 1 Screenshot of a QuickTime web video

Direct electronic transfer of digital video to end users via the Internet presents an exciting opportunity for teachers and students to distribute their video programs. Student work can be shown in virtual exhibitions, guest speakers and one-off PD can be captured and distributed, and distance students can experience classroom and school activities.

Uncompressed digital video files are large but they can be compressed (see the *Video Compression* whitepaper) for online distribution.

Web video clips can still be very compelling even when they occupy only a small part of the screen as shown in Figure 1. At this size (around 180 pixels wide by 120 pixels high), a clip can be downloaded surprisingly quickly, even by users on dialup Internet connections.

Video editing software is able to save clips in one or more of the popular web compression formats such as QuickTime or Windows Media (which need software that can read their formats installed on your computer). Specialised compression software gives even more control over the process for those seeking higher quality or who need to create lots of different clips. An example of needing software to view videos online is this website. If you don't have the QuickTime plugin or player, you will not be able to view the video content presented here.

Deciding which video format to select is often difficult because of the variety of formats available and their lack of compatibility with each other. In the school environment, users themselves are generally not permitted to install new software. This limits your choice of video formats to those that your end users are already set up to play.

The Motion Picture Experts Group has also developed a new standard called *MPEG4*. This standard is supported by some players, and has the potential to bring the benefits of standardisation to the web that came with the MPEG1 and MPEG2 formats.

If your video is available in several different formats and perhaps at different sizes, the likelihood that it will be viewed is maximised, regardless of the type of computer being used or the speed of the user's network connection. Users can then choose a clip of a size and format appropriate to their circumstances.

Figure 2 illustrates how these choices might be presented to users.

Download Annual Sports Day video here !				
Format	320x240 movies		160x120 movies	
QuickTime	sportsday.mov	2.3M	sportsdaysmall.mov	900K
Real Video	sportsday.rm	2.1M	sportsdaysmall.rm	800K
Windows Media	sportsday.wmv	1.9M	sportsdaysmall.wmv	750K
MPEG4	sportsday.mp4	3.0M	sportsdaysmall.mp4	1.1M

Figure 2 Table of various web video formats and their corresponding file sizes

Planning for online distribution of large and potentially popular files should always be discussed with the IT administrator.

Trade-offs

Your decision about how to distribute video not only involves a trade-off between cost to the school and convenience for the end user, but also a consideration of the nature of the content. DVD is appropriate for longer programs while web distribution suits shorter snippets of information that can be comfortably viewed by users seated at a computer.

Online distribution – a case study

This Week @ Hawthorne High is a 5-minute weekly video magazine program produced by Hawthorne High student to keep parents and the local community informed of and engaged in activities at the school.

The school has the equipment to author DVD. While an informal poll of students indicates that many homes have a DVD player or computer with a DVD-ROM drive, the production of multiple DVDs on a weekly basis is considered too costly.

CDs present a cheaper option but the school wants to avoid the environmental impact of hundreds of CDs being discarded every week. The potential for online distribution is therefore investigated.

The IT administrator confirms that the additional traffic on the school's website will not create any major technical or financial issues. The footage is shot by students on a *DV camcorder* that can be connected directly to a desktop or laptop computer via a *FireWire* cable as shown in Figure 3. By transferring the digital video data from tape directly onto the computer's hard disk there is no loss and the entire process of digitisation is eliminated. This process is called *Capturing*.



Figure 3: *FireWire* enables digital video data to be transferred directly from a camcorder to a computer for editing

Once the video data has been transferred, it can be edited and then the finished program can be compressed for Internet delivery. The school decides to invest in specialised video compression software so the compression process can be finetuned. Even a small reduction in the average size of the files will improve the download experience for parents, and reduce the volume of data served by the school's web server over the school year.

The video newsletter is divided up into smaller sections by topic area, allowing parents to choose what they would like to download, based on their interests. There are also several variations of file size from which to choose. Some viewers will have slower *dialup* connections while others will have faster *broadband* connections such as cable or *ADSL*.

Dealing with data

As discussed in *The Need for Compression* whitepaper, the manner in which you capture and compress your footage will have an effect on the size of the video or audio clip on your computer's hard drive. When capturing directly from your camcorder, most applications will assume you want the best possible quality and capture your footage at DVD quality. All of this information is automatically stored on your computer – sometimes even in places where you may not consider looking! Always pay attention to where your footage is being captured to, and if you want your footage at a smaller size (ie, you're only interested in publishing your work on the web), consider setting your program to capture the footage at the size you plan to output it to. Note that if you ever need to re-use the footage for a higher resolution output, you will need to re-capture it.

The end of the tape?

Video programs are traditionally distributed on videotape. Programs shot and edited digitally can still be distributed using this analogue format. Small numbers of *VHS* cassettes can be duplicated at school. Also, you can find local video production companies with bulk duplication facilities in most communities.

Key terms

ADSL

Authoring

Broadband

Camcorder

CD-ROM

CD-R

CD-RW

DV

DVD

DVD-R

DVD-RW

FireWire

MPEG

MPEG1

MPEG2

MPEG4

QuickTime

Real Video

VHS

Video CD

Windows Media

Image Compression

Useful background whitepapers

- *What is Digitisation*
- *Image Digitisation*
- *The Need for Compression*

Further reading

- *Video Compression*

Key concepts

- Many photographic images contain regions that are relatively uniform; for example, the brilliant blue of a sky unbroken by cloud.
- The repetitive nature of these uniform regions allows them to be very efficiently compressed.

Introduction

An *uncompressed* digital image file can be many *megabytes* in size. The science of compressing images involves finding ways to describe the rich detail of a complex image using less space.

Run length encoding

Representing a scene digitally, as discussed in the *Image Digitisation* whitepaper, leaves us with an original uncompressed or *raw* image file in which each pixel is stored independently. If network *bandwidth* and storage were not obstacles, then these uncompressed images could be used directly; however, a web page would take a very long time to download if every image was several megabytes in size!



Figure 1 Bonsai

Consider the image shown in Figure 1 of a bonsai shot against a solid purple background.

Notice that the image isn't uniform in complexity. There are large areas in the background where there is little if any variation. As noted in *The Need for Compression* whitepaper, complexity and variation are difficult to compress. Therefore, areas like the purple background of Figure 1, that contain less 'information', are an ideal starting point in the search for opportunities to economise on the storage requirements for images.

In Figure 2 we have zoomed in to show the area indicated by the arrow in Figure 1.

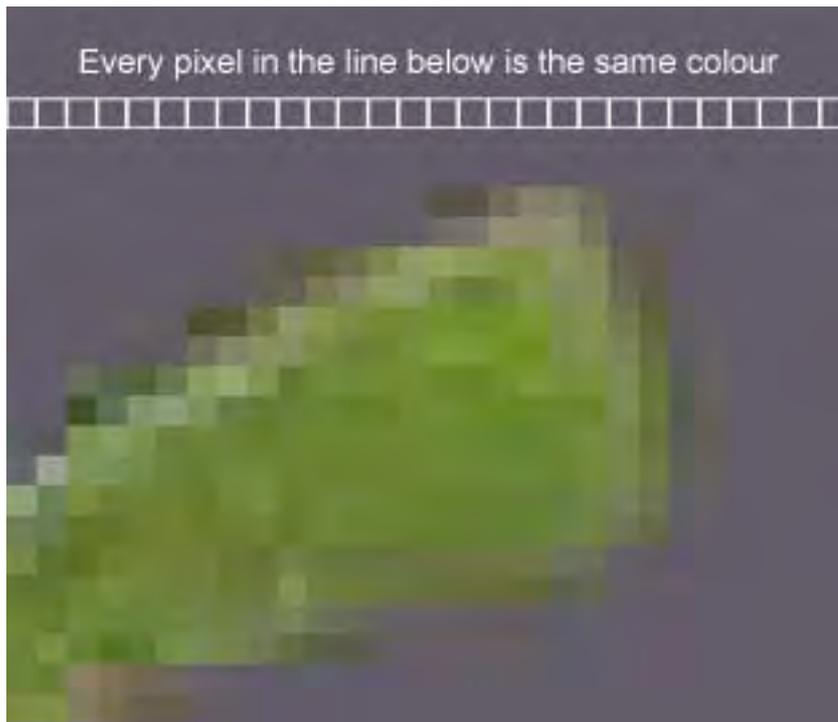


Figure 2 Close up of the bonsai shown in Figure 1

We can see a line of identical purple coloured pixels in the background of the bonsai image just behind the topmost leaf. The *Image Digitisation* whitepaper also explains how the colour of a pixel in a digital image is represented using a *binary* number to indicate the proportions of red, green and blue required to create the colour. Why store this number for each and every pixel within the line when they never change?

Instead of repeating the same numbers over and over again, a more efficient strategy is to use a simple code consisting of only two numbers. The first number stores the colour of the pixels (which does not vary across the line), and the second number simply indicates how many pixels there are.

This method of compression is known as run length encoding or *RLE*, because a *run* of pixels of a certain *length* is being encoded. Run length encoding is incorporated into the popular Graphic Interchange Format, known as GIF used for web graphics. GIF is a CompuServe trademark.

GIF works well for simpler graphics containing areas of solid colour, but is less appropriate for more complex real world images such as photographs. Generally, the *JPEG* format is more effective for the compression of natural images, and relies upon more advanced mathematically based techniques to do this.

Although run length encoding is perhaps the simplest of the many techniques available in the compression arsenal, it serves as effective illustration of the principle of removing redundant or surplus information in order to compress a digital image.

If you want to compress an image (or video), it's always a good idea to retain the original, uncompressed file rather than trying to make do with a compressed version. If you ever want to go back and edit your material and you only have a web delivery version, you will no longer have the breadth of information in the file as discussed at the start of this chapter. Therefore, to *present* or *display* footage or an image, you will compress it, but to *edit* it you will not.

Key terms

GIF – Graphics Interchange Format

JPEG – Joint Photographic Experts Group

Run Length Encoding (RLE)

Curriculum Corporation – Video Research Project		
		Version 1.0
Image Digitisation: (Representing an Image Digitally)		

Useful background whitepapers

- *What is Digitisation?*

Further reading

- *The Need for Compression*

Key concepts

- Just as a tapestry is composed of many individual stitches, a digital image is composed of many separate *pixels*.
- The colour and brightness of each pixel can be represented by a number, and these numbers can be combined into a digital image file.
- Higher levels of image quality can be achieved by increasing the range of numbers (colours) available, but this brings with it a commensurate increase in file size.

Introduction

Digitisation involves taking a snapshot of our inherently *analogue* world and breaking this snapshot down into tiny digital fragments, which can be stored and manipulated by computer. When digitising images, each chunk or element has a colour and brightness value corresponding to that of the original scene.

With enough picture elements, we can create an acceptable facsimile of the original scene. The process is akin to creating a mosaic from individual ceramic tiles. Creating a highly realistic mosaic requires the use of many small tiles.

Rather than tiles, digital images are composed of individual elements called *pixels* (an abbreviation of the term *picture elements*). Information about the colour and brightness of each pixel is represented numerically (or *digitally*) and combined into a digital image file. These files will be large if the images are highly detailed. This proportional relationship between quality and file size comes into play whenever analogue media is converted to digital.

To help understand the process of image digitisation, imagine a set of library pigeonholes as depicted in Figure 1.

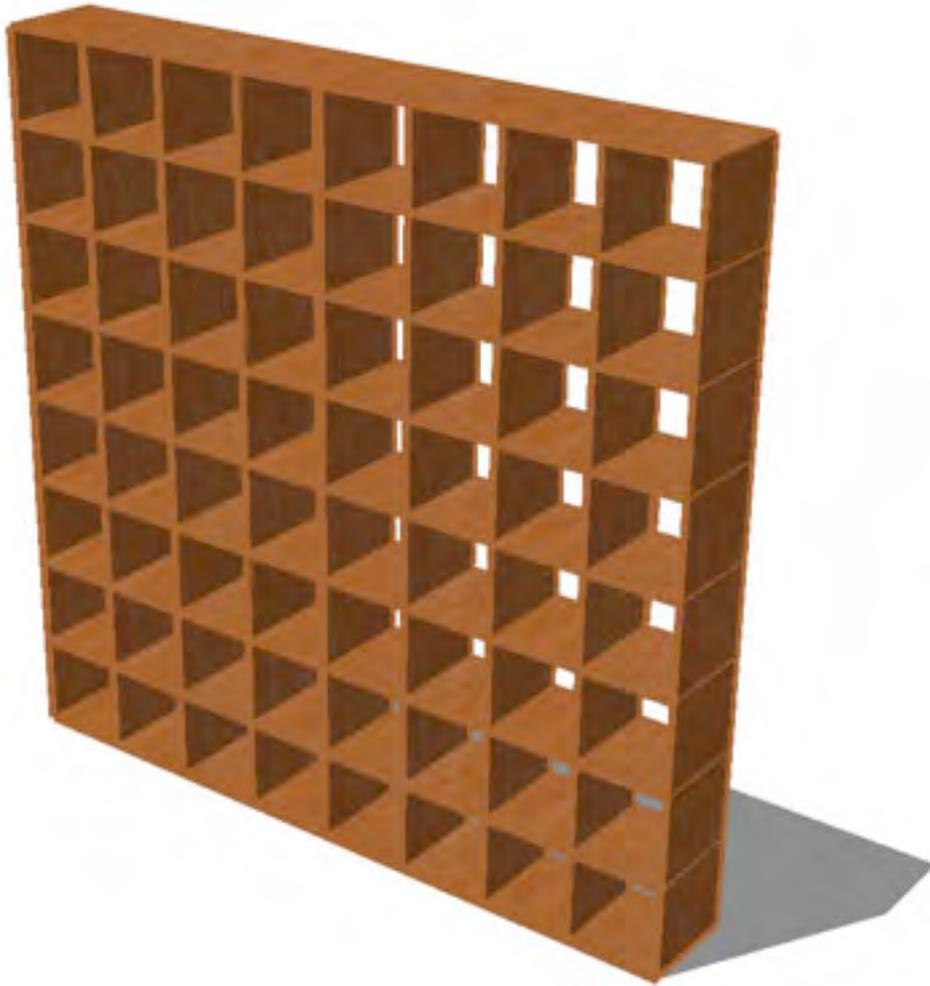


Figure 1 Library pigeonholes

To show how a simple image can be converted into digital form, imagine that a group of students have placed a sketch of a smiley face over the front of the pigeonholes as shown in Figure 2.

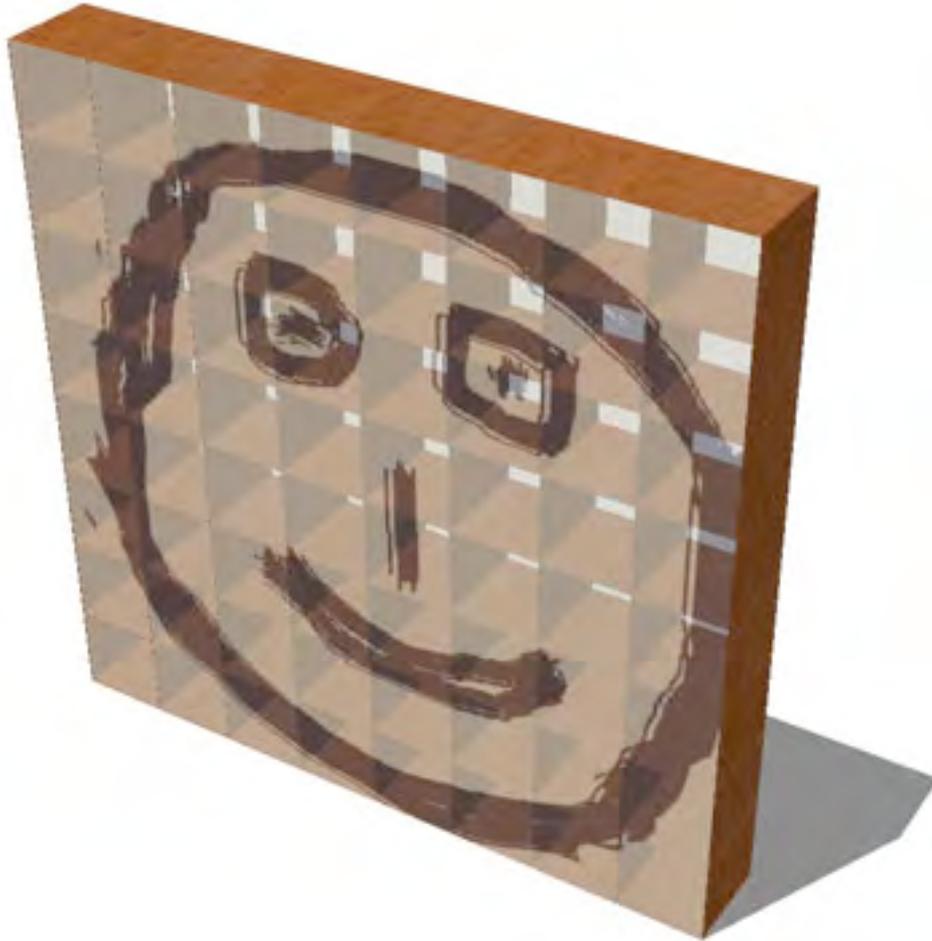


Figure 2 Pigeonholes with smiley face

Each one of the 64 pigeonholes can be thought of as a container holding a single pixel. The students' next step is to determine whether each pigeonhole should contain a pixel that is on (white) or off (black). A black card can be used to represent the off pixels and a white card for the on pixels (see Figure 3).

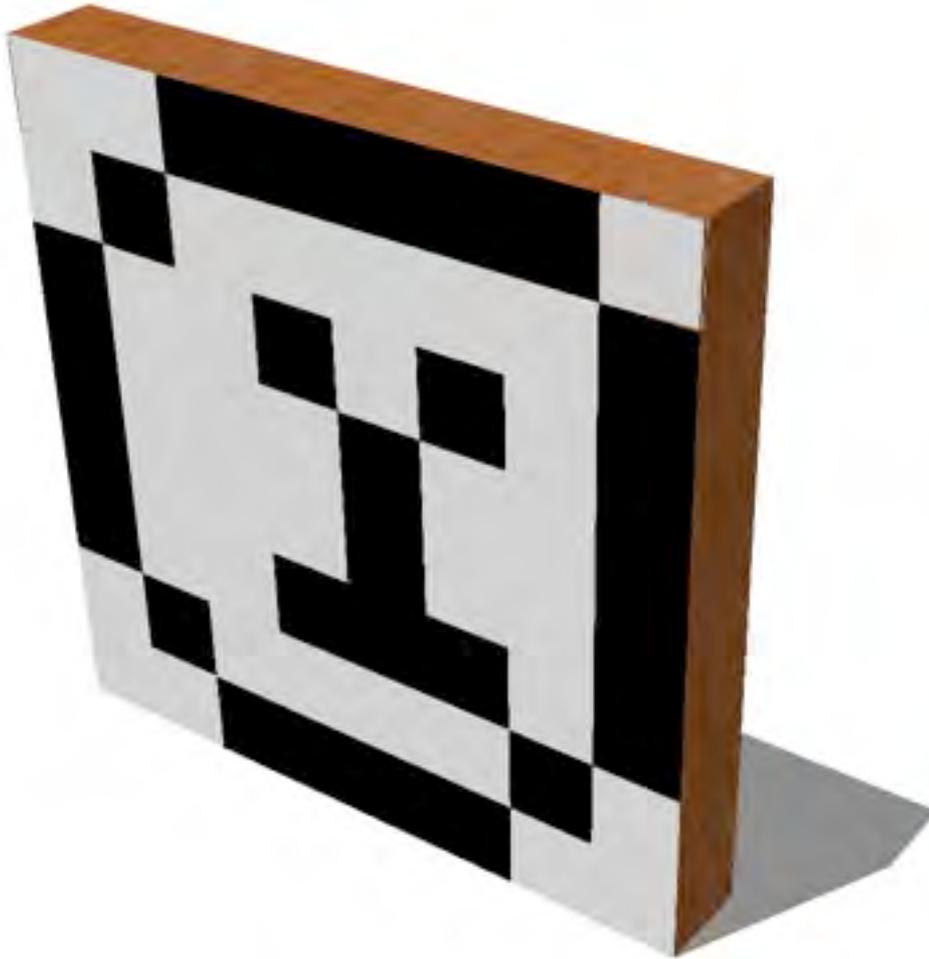


Figure 3 Pigeonholes with on and off pixels

This is obviously a crude facsimile of the original image – but its very simplicity enables it to be represented very economically in digital (number) form.

A single number in the binary system can be used to represent the colour of each pixel. In contrast to the normal decimal number system of ten digits from zero to nine, the binary system allows only two digits: zero and one.

Let's use the binary number *zero* to represent black (or those pixels that are *off* when the image is finally displayed on a monitor), and the binary digit *one* to represent the white areas of the image (or the *on* pixels).

The term *bit* is an abbreviation of *binary digit*. Therefore, we are digitising our image using a *bit depth* of one bit-per-pixel as illustrated in Figure 4.

Any levels of brightness that fall between completely black (*zero*) and completely white (*one*) must be assigned absolutely to one or the other value. Obviously, apart from those few areas of the poster that are actually fully black or fully white, this will be an approximation and is an (extreme) example of the approximations that are always associated with digitisation.

Designers of contemporary digital systems such as CD players or digital cameras try to *sample* at a sufficient depth as to make these approximations imperceptible to most users.



Figure 4 Pigeonholes with binary data

In Figure 4, each row of eight pigeonholes stores eight bits. Computer memory is measured in bytes and there are eight bits in a byte. Therefore, the state of the pixels in each row of the smiley face can be captured with a single byte and the entire image, consisting of eight rows, will require only eight bytes.

Although an eight-byte file is very small, this image is tiny at only eight pixels by eight pixels. Figure 5 shows the actual size of the smiley face.



Figure 5 Actual size of the smiley face

The relationship between bit depth and image quality is better illustrated using a larger image that contains photographic rather than pictorial material.

Figure 6 is a scene of trees and buildings shot from an apartment balcony. This image is larger and has dimensions of 320 pixels by 240 pixels. At a bit depth of one bit-per-pixel, the buildings and tree are just recognisable.



Figure 6 One bit image

To achieve a more realistic image we need to select pixels that are closer in intensity to those of the original image.

Digitising with a bit depth of one bit-per-pixel limits us to only black or white. By using *two* bits-per-pixel the number of shades doubles to four as shown in Figure 7.

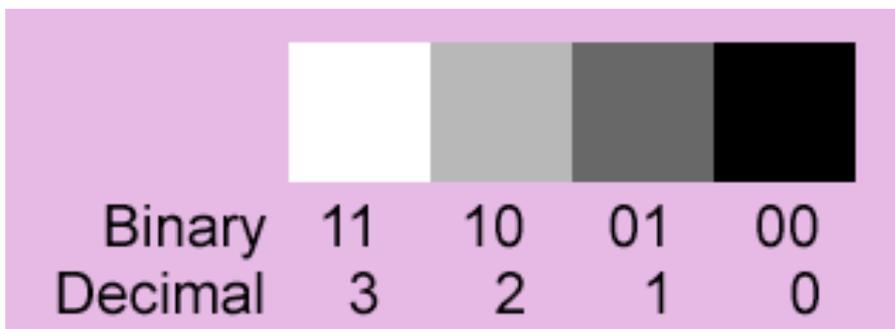


Figure 7 Two bit swatch

With a bit depth of two bits-per-pixel the quality improves noticeably (see Figure 8). It is important to note that with this improvement in quality, twice as much memory will be required to store the file.



Figure 8 Two bit image

With four bits-per-pixel, there are 16 ($2 \times 2 \times 2 \times 2$) levels available. In Figure 9, we can see subtle details beginning to emerge.



Figure 9 Four bit image

To achieve high-quality reproduction, a bit depth of at least eight bits-per-pixel would typically be used.

Colour images

The concepts introduced above describe how the brightness of the pixels in a black and white digital image are represented numerically and stored digitally. Digitisation of colour images is very similar. However, instead of one binary number being used to describe the brightness of each pixel, we use three

binary numbers to describe the amount of red, green and blue components that must be mixed together to create the desired colour. Red, green and blue is often abbreviated to RGB. You may well have heard colour computer monitors referred to as RGB monitors. Black and white images are considerably smaller than their colour counterparts

A black pixel is off and emits no light at all. Consequently, the red, green or blue RGB value is 0,0,0. A white pixel contains the equal and maximum amounts of red, green and blue. Assuming an eight-bit *colour* image in which eight bits or 256 different levels are available for *each* colour channel (normally referred to as 24-bit colour), then the red, green and blue colour values for full white can be represented in the decimal system as 255,255,255.

The three component colours themselves are described in Table 1.

Eight-bit colour display table			
	Red value	Green value	Blue value
Red	255	0	0
Green	0	255	0
Blue	0	0	255

Table 1: Eight-bit RGB values of the red, green and blue components

Any arbitrary colour can be created with an appropriate blend of red, green and blue. For example, by combining the maximum amount of red (255), a moderate amount of green (139) and a touch of blue (45), we can create the orange at the centre of the flower depicted in the eight-bit image below (Figure 10).



Figure 10 Eight bit image and RGB values of selected colour

Key terms

Binary

Bit

Bit Depth

Byte

Pixel

RGB

Sampling

Useful background whitepapers

- *Video Compression*

Key concepts

- Applying *practical* video shooting and editing techniques that are derived from an appreciation of basic video compression *theory* can visibly improve the quality of video shot for the web.

Introduction

When shooting video destined for the web or for a CD-ROM project, are there things to keep in mind to ensure the best possible compressed result? Absolutely!

Keeping the final output in mind during editing can also enhance your web video clips.

Guidelines to maximise compression

These guidelines for shooting video for the web derive from an understanding of how video compression *codecs* squeeze uncompressed video into clips compact enough for delivery over the web or on CD as *QuickTime*, *Windows Media*, *Real* or *MPEG4* formatted files (as discussed in the *Distribution of Video* whitepaper).

(Note that because the *DVD* format provides such a high *data rate*, it is not necessary to adopt special shooting and editing techniques.)

Video can be compressed by information that is repeated both within a single frame and from one frame to the next. (See the *Image Compression* and *Video Compression* whitepapers respectively.)

Therefore the more *repetitive* the video sequence is from one frame to another, the easier it will be to compress. Expressed another way, the more one frame has *in common* with its predecessor, the less *new* information needs to be sent down the line to describe it.

From this basic understanding of compression, we can derive four guidelines for maximising the 'compressibility' of your footage:

1. Place the video camera on a tripod. A large camcorder that sits on the shoulder is quite stable but holding a small video camera steady by hand is difficult. The results look unprofessional *and* make for difficult-to-compress video because the entire video frame is constantly in motion.

If objects in the scene such as buildings are stationary, they should also be stationary in your footage. A tripod stabilises the camera resulting in a more professional looking video and enabling any intentional movement of the camera to be performed smoothly.

2. Shoot your subject against a static background. Data must be used to describe the motion of objects in your video; therefore, the data should describe the movement of your *subject*, rather than superfluous background detail such as leaves blowing about or traffic passing by.

3. Avoid pans and tilts. *Panning* is horizontal movement of the camera and is used to follow action or reveal more of a scene. In Figure 1 the camera is being panned left to follow the van as it negotiates the corner. Panning the camera creates motion within the video frame that can starve the crucial foreground subject of data.



Figure 1 Panning left

Ideally, you should plan your web video production so that the story can be told using *locked-off* shots with fixed framing. If panning is necessary in order to follow the action, it should be done with a high-quality *fluid head* tripod to facilitate smooth and consistent panning.

While panning involves horizontal movement of the camera, vertical movement is known as *tilting* and is often used to keep a seated subject in the video frame as they stand. Again, if vertical camera movement is necessary, use a fluid head tripod.

4. Avoid the use of the zoom lens. Excessive use of the zoom lens on a camcorder looks unprofessional and results in video that is particularly difficult to compress. Panning, tilting and even the camera shake that is associated with handheld operation merely result in the movement of pixels horizontally and vertically from frame-to-frame. Most codecs can handle this reasonably well because the pixels themselves are unaltered as they move. However, zooming causes *every pixel* of the image to change as objects in the frame alter their size.

In compressed clips that contain zooms, a noticeable pixelisation of the image can occur, as the codec is unable to keep up with this rapid change. As soon as the zoom ends the image recovers. Instead of zooming, try shooting a wide *establishing shot* followed by a *close-up* of the subject in question, or cut the zoom out later during editing.

Editing suggestions

Ideally, your video material will have been shot according to the guidelines outlined above. However, most of the time you will be reusing material shot without the bandwidth constraints of the web in mind. Even here though, decisions made when editing can affect the final quality of the compressed output.

When editing two shots together, a hard cut is best in terms of compression. Hard cutting joins the final frame of the first clip directly to the first frame of the next shot (see the top filmstrip in Figure 2) without any complex *transitions*.

The worst transition from a compression standpoint is the dissolve. This transition should be avoided. During a dissolve, one shot is blended or gradually mixed into the next (see the bottom filmstrip in Figure 2). During the period of the dissolve every pixel changes as the relative proportions of the two shots is progressively altered.

Figure 2 compares the data rate impact of hard cuts and dissolves.

Editing for Compression Impact on Data Rate of Cut vs Dissolve

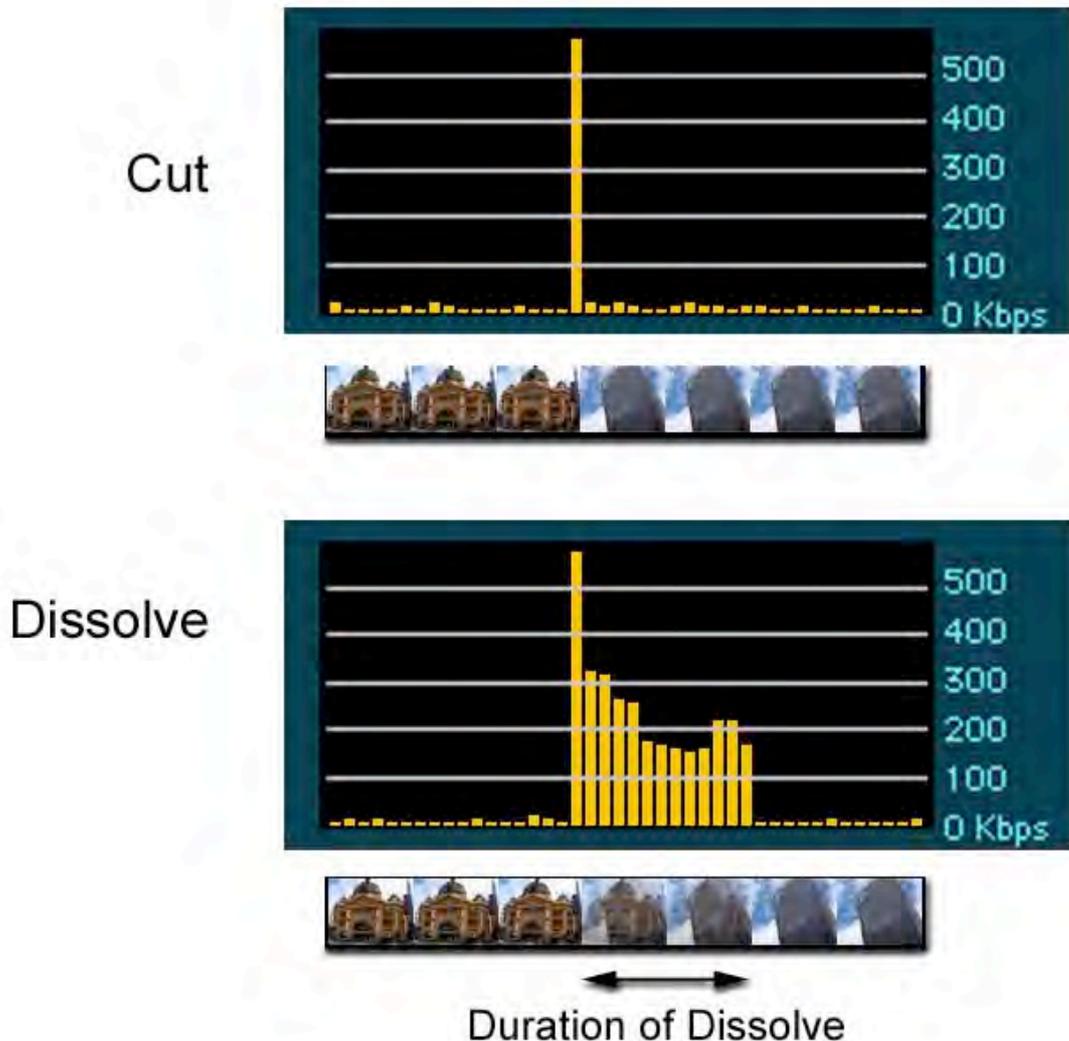


Figure 2 Data rate impact of hard cuts and dissolves

In the top graph, the orange bars indicate a low data rate associated with a locked off shot of a building with light traffic passing by in the foreground. Then an instantaneous spike of data rate as the first frame of the new scene is introduced. In the bottom graph, instead of using a cut, a prolonged dissolve has been employed resulting in one video sequence slowly blending into the other. It is obvious that during this period the data rate must be increased because over the duration of the dissolve, each frame is effectively a *new frame*, created from ever shifting proportions of pixels from the two separate scenes. Once the dissolve is complete, the data rate again settles down once again to the same low level associated with a static shot.

Aesthetically the impact of a dissolve, which can be quite compelling on film or television, is often lost on the computer screen, deteriorating into a mess of chunky pixels. It is much better to make a clean break between scenes. In this way we take all the pain in one hit, rather than prolonging it over a second or more. It's the video compression equivalent of quickly ripping off a bandaid rather than peeling it off slowly.

If a transition other than a cut is required, then a 'mechanical' transition such as a *wipe* will be more readily compressed.

Key terms

Close-Up

Dissolve

Establishing Shot

Fluid Head

Locked Off

Pan

Tilt

Transition

Wipe

Video Compression

Useful background whitepapers

- *What is Digitisation*
- *Image Digitisation*
- *The Need for Compression*
- *Image Compression*

Further reading

- *Shooting and Editing Video for the Web*

Key concepts

- Compression of moving images, in other words video, involves the application of the same techniques that are used to remove repetitive or redundant information within *still* images.
- However, with moving images as well as achieving compression through the reuse of information *within* an image, we have the additional opportunity to reuse information that is repeated *across* adjacent images or video frames.

Introduction

Video creation by students and teachers is increasingly common in the classroom, and yet video files are the largest multimedia files commonly encountered in the educational context and their size makes them the most problematic to distribute and store.

Compression encoding

If video is conceived of as merely a sequence of still images, could we not simply string together a sequence of (say) *JPEG* compressed *still images* to create compressed *moving image* or *video* sequences?

In fact, this is exactly the case with the *Motion JPEG* format. Figure 1 shows a user selecting Motion JPEG as the format in which to save a video file within the QuickTime Player Pro application from Apple.



Figure 1 User selecting Motion JPEG as the format in which to save a video file within the QuickTime Player Pro application from Apple

However, treating video as nothing more than a series of stills *in isolation*, as is done within Motion JPEG, ignores the opportunity to remove a major source of *redundancy* that exists within video. That opportunity is the removal of information that is repeated *across* multiple frames.

Just as compressing images, as discussed in the *Image Compression* whitepaper, removes repeated information within different parts of the *same* frame, *video* compression can also benefit dramatically from removing regions of pixels in one frame that also appear within subsequent frames of the sequence. Note that video files also contain audio data, another large format media element. This part of your video will also need to be compressed to contribute to a small file size. You can see the audio *waveform* in the screen grab below.

Consider the brief video fragment in Figure 2 that shows the entrance to Flinders Street Station in Melbourne.



Figure 2 Video fragment

If each one of the eight frames is compressed in isolation, as say individual *JPEG* images, then the file size of this short video sequence will be the same as that of eight separate *JPEG* compressed photos of the station. However, because the initial frame contains a full description of how to draw the station, we can simply re-apply that description in the subsequent frames.

Rather than sending all the information required to create an image of Flinders Street Station 25 times every second, we send it once, in the first frame or *key frame*, and effectively paste a copy of the station into the subsequent frames in which it appears. In this way, we get these subsequent frames more or less 'for free'. These subsequent frames aren't actually free, but they are much smaller than the key frames. This is particularly applicable to the type of footage shown in Figure 1. Here, the camera itself is not moving and most of the detail within the shot, such as the station and sky in the background, change little if at all from one frame to the next.

We can continue to borrow elements of the first frame, perhaps with slight adjustments, until such time as we strike a frame so totally different (perhaps a new scene), that it is more economical in data terms to insert a new fully *independent* self-contained key frame and begin the sequence again.

This approach, known as *interframe* (between frame) compression is used in the very popular *MPEG* format, and contrasts with the *intraframe* (within frame) approach used in *Motion JPEG* and, incidentally, the *Mini-DV* format that is in widespread use in digital *camcorders*.

Obviously, compression will be easier if there is minimal variation from one frame to the next. In fact it is possible to exploit this knowledge when shooting original footage so that it compresses well (see the *Shooting and Editing Video for the Web* whitepaper).

When compressing existing video material, it is also important to work from the best, cleanest, raw material. If your source video is a well-used *VHS* tape, you will notice random 'sparkles' on the screen during playback which is known technically as random noise. Noise is what you hear on a radio when no station is tuned in or see on a television screen when no station is being received. The random nature of this 'noise' plays havoc with compression techniques that, as we have seen, rely upon the very non-randomness or

predictability of video content in order to work. So try to obtain the highest quality source tapes, such as *Betacam SP* or *DV*, rather than a *VHS* dub of the original.

It is also important, if offered digital video as files that have already been captured to disk, to avoid those that have themselves already been compressed. Trying to compress from already heavily compressed source material often gives poor results.

When working with video a good general rule of thumb is to get the highest quality original source material as possible and do the compression and conversion yourself. However, professionally shot footage may be on a high-end device that may be costly or difficult for you to convert in your school. Try to familiarise yourself with the formats listed above, and compare them with what your school can handle or work with. Also, remember that high-quality footage takes up disk space, so look at *The Need for Compression* whitepaper to help you understand how much footage can fit on to a basic PC.

Key terms

Dub

DV (Digital Video)

Key Frame

Mini-DV

Motion JPEG

MPEG

Noise

Redundancy

Curriculum Corporation – Video Research Project		
		Version 1.0
Glossary		

3G	
ADSL	Abbreviation for Asynchronous Digital Subscriber Line. A type of <i>broadband</i> high-speed Internet connection technology that enables ‘real-world’ connection speeds of 25KB/sec to 100KB/sec. That is, five to 20 times faster than standard <i>dial-up connection</i> . ADSL piggy-backs on a normal phone line, allowing the phone to be used for voice calls while connected to the Internet. ADSL is not available in all areas.
Authoring	Simply placing MPEG2 video files on a DVD-ROM is insufficient to enable the disc to be played in a conventional DVD player. The disc must be <i>authored</i> . This means the information must be laid out in a particular way to conform to the official DVD standard. In this way playback on any domestic DVD can be assured. Basic DVD authoring software is generally included with DVD burners.
Bandwidth	The bandwidth of a communications connection refers to the speed at which data can be moved over it. <i>Dial-up modem</i> speeds reach 56Kbps or 56,000 bits per second. <i>ADSL</i> services begin at 256Kbps. Typical <i>Ethernet</i> LAN (local area network) connections are 11Mbps or 11,000,000 bit per second or greater.
Betacam SP	A professional broadcast quality analogue videotape format introduced by Sony in 1986. Betacam SP uses half-inch tape and is popular for the acquisition of news footage and in professional video production.
Binary	Our normal decimal numbering system has 10 digits from zero to nine. The binary system has only two digits, zero and one. This is convenient for computers as a number can be represented simply as the absence (zero) or presence (one) of an electric charge in the computer’s circuitry. See also <i>bit</i> , <i>Byte</i> , <i>Gigabyte</i> and <i>Megabyte</i> .
Bit	A single digit in the <i>binary</i> number system, either a zero or a one. See also the related larger units of <i>Byte</i> , <i>Gigabyte</i> and <i>Megabyte</i> .
Broadband	A faster than <i>dial-up</i> Internet connection such as <i>ADSL</i> or cable <i>modem</i> .
Byte	A sequence of eight <i>binary</i> digits (<i>bits</i>). For example, 11111110 (which in the normal <i>decimal</i> system is 254). There are 2 ⁸ (2 to the power of 8 or 2x2x2x2x2x2x2x2) different values that a byte can represent, which is equivalent to the 256 numbers in the range 0 to 255 in conventional decimal notation. Computer memory is generally measured in multiples of a byte such as <i>Gigabyte</i> and <i>Megabyte</i> . See also <i>bit</i> , <i>Byte</i> , <i>Gigabyte</i> and <i>Megabyte</i> .

Camcorder	A contraction of the words 'camera-recorder' meaning a video camera and videotape recorder integrated into a single housing. At one time separate video cameras and videotape recorders were the norm.
Capturing	The process of getting footage from an input (usually a camera) and bringing it into your computer via a FireWire or IEEE 1394 cable or capture card.
CD-R	Compact Disc Recordable. A compact disc that can have data written to it once only using a CD burner.
CD-ROM	Compact Disc-Read Only Memory. A variation of the digital audio CD format that facilitates the storage of general computer files.
CD-RW	Compact Disc Re-Writeable. A compact disc that can be erased and re-used in the same way as a floppy disc by using an appropriate CD-RW burner.
CinePak	A video compression <i>codec</i> supported in both QuickTime and Windows Media that was popular in the 1990s for the development of interactive multimedia <i>CD-ROM</i> .
Close-Up	Traditionally a video shot in which the head and shoulders of the subject fill the screen.
Codec	A contraction of Compressor-Decompressor. A set of software rules which define how to compress and decompress video. Examples include <i>Cinepak</i> , <i>Sorenson Video</i> and <i>MPEG</i> .
Compression	Reducing the size of a file. Compression can either be 'lossless' or 'lossy'. Documents such as Word files are compressed 'losslessly' so that no information is lost. Multimedia files such as video can have lossy compression applied, as generally our compressed video file doesn't need to be identical, or <i>bit-for-bit</i> , to the original.
Data Rate	The rate at which information moves over a network connection or from a storage medium like CD or DVD. Measured in units of Kilobits-per-second (Kbps), Kilobytes-per-second (KB/sec) or related larger units.
Decimal	The conventional (base 10) number system comprising 10 digits from zero to nine. Contrasts with the base two or <i>binary</i> system employed in computing containing only the two digits zero and one.
Deck	The video production industry term for a video recorder or VCR (video cassette recorder).
Dial-Up	Dial-up Users connect to the Internet via a <i>modem</i> , a device that plugs into a normal phone connection. Modem speeds have gradually improved over time to a maximum of 56Kbps; however, this is still a very slow link over which to transfer large multimedia files. The phone line becomes engaged when using a dial-up modem. See also <i>broadband</i>
Dissolve	A type of video <i>transition</i> involving the gradual mixing or blending of one video shot with another over time.

DV	Digital Video. A manufacturer standard for consumer and professional recording of video in digital form. The popular consumer Mini-DV format uses 6.25mm (quarter inch) videotape in compact cassettes that store up to an hour of video at a <i>data rate</i> of 25Mbps. DV can be transferred directly from the camcorder to a computer for editing via a <i>FireWire</i> .
DVD	Digital Versatile Disc (or Digital Video Disc). The same size as a compact disc, the DVD packs information much more densely with data capacities from 4.7GB (Gigabytes) to almost 20GB.
DVD-R	A recordable (once-only) <i>DVD</i> that can be written to a DVD burner. The cost of DVD burners (or writers) has plummeted and now they are a standard feature in many new computers. The DVD-R is the DVD equivalent of the <i>CD-R</i> .
DVD-RW	A rewriteable version of the <i>DVD</i> that can be written to and erased, and re-used.
Dub	A copy of a video tape.
Establishing Shot	Often the opening shot of a video sequence that establishes the context for the viewer. Typically, an establishing shot is wide and is usually followed by increasingly tighter shots such as the <i>close-up</i> .
Ethernet	A computer network standard that is common for the connection of computers and printers in the workplace. Speeds available generally range from 11 Megabits per second upward. The ubiquitous blue cables snaking around offices are Ethernet cables.
FireWire	A data transfer standard developed by Apple Computer and ratified as international standard IEEE1394. FireWire enables high-speed transfer of digital information (often video information) between computers and hard drives, scanners, digital cameras, camcorders and external drives. FireWire cables are slim and flexible. Sony refers to IEEE1394 as Sony iLink.
Fluid Head	A fluid head tripod enables extremely smooth <i>pans</i> and <i>tilts</i> to be performed because the jerky motion of the camera is damped or smoothed by fluid contained within the head of the tripod.
GIF	Graphics Interchange Format (tm). An image compression format developed by online service provider CompuServe that is commonly used for graphics on web pages.
Gigabyte	One thousand million bytes or 1000 <i>Megabytes</i> of digital information. See also <i>bit</i> , <i>Byte</i> and <i>Megabyte</i> .
Interframe	Interframe compression works <i>across</i> frames, taking advantage of information that is common across multiple frames. Interframe codecs contain complete independent <i>keyframes</i> in combination with other frames that are <i>derived</i> from the keyframe. MPEG1, 2 and 4 are popular interframe codecs. Interframe compression contrasts with <i>intraframe</i> compression.
Intraframe	Intraframe compression occurs only within a given frame and without reference to adjacent frames. <i>Motion JPEG</i> and <i>DV</i> are common intraframe codecs. See also <i>Interframe</i> compression.

JPEG	Developed by the Joint Photographic Experts Group, a compression format popular for digital still images such as those taken with digital cameras.
Key Frame	A fully self-contained, independent frame of video that can exist without reference to any other frames.
Kilobyte	1000 Bytes of digital information (strictly 1024 Bytes). See also <i>bit</i> , <i>Gigabyte</i> and <i>Megabyte</i> .
Locked Off	When the head of a video camera tripod is 'locked off', the camera remains fixed in a stationary position even if unattended by the camera operator.
Lossless Compression	In general when we compress a file we want that file, when subsequently uncompressed, to be identical to the original, with no loss of data. You want your Word document to be identical to the original once it is uncompressed. A software developer distributing software in a zip archive wants to know that you receive an identical copy of the file. With multimedia files such as video it is not critical that every bit be identical to the original source file. See also <i>Lossy Compression</i> .
Lossy Compression	In contrast to <i>Lossless Compression</i> , many multimedia files can be compressed using Lossy Compression formats so that even though the file, when it is uncompressed during playback, is not identical at the <i>binary</i> level to the original source file, we still find the results (hopefully) acceptable.
Megabyte	One million bytes (strictly 1048576 bytes) or 1024 K. See also <i>bit</i> , <i>Byte</i> , and <i>Gigabyte</i> .
Mini-DV	A high-quality digital consumer videotape format using a data rate of 25Mbps/sec popularised in Mini-DV <i>Camcorders</i> . The Mini-DV format, in conjunction with <i>FireWire</i> , has revolutionised video production by dramatically improving quality and convenience.
Modem	A contraction of Modulator-Demodulator, a computer peripheral that converts computer data into audio tones that can be transferred over a standard phone line. The maximum speed of a dial-up <i>modem</i> is 56K bits/sec. <i>Broadband</i> modems are capable of much higher <i>data rates</i> .
Motion JPEG	A variation of the <i>JPEG</i> still image compression format adapted for use as a video compression <i>codec</i> . Motion JPEG uses <i>intraframe</i> compression.
mp3	A popular compression format for music. <i>mp3</i> is an abbreviation for <i>MPEG1</i> Layer 3 Audio and was developed as one of the options available for compressing the audio tracks in an <i>MPEG1</i> video file.
MPEG	Motion Picture Experts Group. An international body that develops standards for video compression; for example, <i>MPEG1</i> , <i>MPEG2</i> and <i>MPEG4</i> .
MPEG1	The first <i>MPEG</i> standard for digital video compression. Released in 1992, it uses <i>data rates</i> of less than 1.8 Mbps and frame sizes of 352x240. Used in the popular <i>Video CD</i> format. <i>MPEG1</i> was designed to be played back successfully even on relatively low-powered computers.

MPEG2	Released in 1995, a higher quality successor to <i>MPEG1</i> allowing larger frames sizes of up to 720 x 576 pixels (in <i>PAL</i>) and higher <i>data rates</i> of up to 15 Mbps. Has proved highly successful as the video compression standard for <i>DVD</i> and is also used in the transmission of Digital TV.
MPEG4	Released in 1999 as a video compression standard for low <i>bandwidth</i> devices, it supports <i>data rates</i> as low as 5Kbps and frame sizes down to 144 x 176. It is the first <i>MPEG</i> video compression standard to address the Internet delivery of video and is being adopted for the delivery of video to wireless devices such as 3G mobile phones.
Noise	Random information in a communications information such as static heard on a radio when station is tuned in or the 'snow' seen on a television when no signal is being broadcast. Noise is very difficult to compress because it contains no <i>redundant</i> information.
NTSC	Abbreviation for National Television Systems Committee. The television and video standard used in the USA and Japan. Incompatible with the <i>PAL</i> standard used in Australia and Europe.
PAL	Acronym for Phase Alternative Line. The television and video standard used in Australia and Europe. Incompatible with the <i>NTSC</i> standard used in the USA and Japan.
Pan	Horizontal movement of a video camera; for example, to follow a running figure. See also <i>tilt</i> .
Pixel	A contraction of <i>picture element</i> . The smallest dot on a computer screen.
QuickTime	The multimedia software architecture for Macintosh and Windows developed by Apple Computer. QuickTime supports a number of video compression <i>codecs</i> including <i>Sorenson</i> , <i>Cinepak</i> and <i>MPEG4</i> .
Real Video	A video compression <i>codec</i> developed by Real Networks to enable video to be <i>streamed</i> over the Internet.
Redundancy	Redundant information is unnecessary information. Information may be redundant because it is already available elsewhere; for example, within another part of an image, or from an earlier frame of video. Redundant information can be removed from a file to reduce its size. See also <i>noise</i> .
RGB Monitor	Stands for red, green, blue. An RGB monitor displays colour by combining different amounts of red, green and blue light.
Run Length Encoding	A method of compressing still images by compacting 'runs' of identically coloured <i>pixels</i> within a single row.
Sampling	Used to describe the regular, periodic measurement of a signal in order to digitise it. Sampling is usually done very frequently in order to ensure that rapid changes, such as those in a musical performance, are captured with sufficient detail.

Sorenson Video	A video <i>codec</i> incorporated into QuickTime from Apple. Sorenson Video is developed by the company Sorenson Vision and was the codec used in 1999 to compress the trailer for <i>Star Wars Episode 1 – The Phantom Menace</i> . When made available from Apple's website, this was the biggest download event in Internet history with six million users downloading the clip in the first three weeks of its availability.
Source File	The uncompressed input file to a video compression program.
Streaming	Enables the playback of multimedia files such as video without the file ever 'touching' the hard drive. Data is transferred straight from the server to the media player (which might be the <i>Windows Media</i> player, <i>QuickTime</i> Player or <i>Real Video</i> player). Quality is limited by the bandwidth of the connection and can be quite variable. A specialised streaming server is required in contrast to a conventional web server, which treats video files in the same way as any other file.
Tilt	Vertical movement of a video camera. For example to follow the motion of a seated individual who then stands. See also <i>pan</i> .
Transition	A visual effect used to join two video clips together. Common transitions include <i>dissolves</i> and <i>wipes</i> .
Video CD	An early but long-lasting format for the storage of video data on Compact Disc. Theoretically offers 'VHS quality' and uses the <i>MPEG1</i> video compression format. Video CD is most popular in Asia.
VHS	VHS or 'Video Home System' is the most popular consumer analogue videotape format. VHS is a half-inch tape format and was developed by JVC and released in 1996, a year after Sony's Beta format which it ultimately succeeded in the marketplace.
Windows Media	Microsoft's video compression architecture for Windows. Players are also available for Macintosh.
Wipe	A type of video <i>transition</i> known as a mechanical transition in which the two video clips being joined 'slide' over each other without any of the underlying clip showing through. See also <u>dissolve</u> .