32

Mass Storage, Display, and Hard Copy

Guy Cox

INTRODUCTION

Confocal microscopes commonly generate their images not as real or virtual patterns of light, but as pixel values in the memory of a computer (Cox, 1993). This gives the image a measure of permanence — unlike a visual image, once acquired it will not fade — but it will be lost if the computer is turned off, or if that area of memory is overwritten. To store that image with all its information intact we must write it in digital form — a copy on paper or film, however good, cannot contain all the information of the original. However, a copy on disk or tape is not directly accessible to human senses. For publication or presentation of the image, or even just to access it, we must have a display or a hard copy, a picture which can be viewed by the human eye.

This chapter reviews the range of possible solutions to these two problems. Because this is a rapidly moving area, new alternatives will doubtless become available almost as soon as this is printed. A measure of the rate at which this happens is that many of the technologies reviewed in the previous edition are now obsolete, leaving users with the task of copying images to new media if they are to retain access to their data. As well as assessing currently available technologies, therefore, I will try to provide enough background information to enable users to assess the latest high technology advances in a rational way. It is always worth considering the scale of the adoption of a technique as well as its technical efficiency because most of us will still want to be able to use our data in 10 or even 20 years' time, and only mass-market solutions are likely to survive on that timescale.

MASS STORAGE

The major problem in storing confocal images is their sheer size. The smallest image we are likely to acquire would be 512 × 512 pixels, at one plane only and containing only one detector channel. Assuming that we store only 1 byte per pixel (that is, each point in the image can have one of 256 possible gray levels) this will require one quarter of a megabyte (MB) to store. We will require a little more space to store basic information about how the picture was acquired, either in a header at the start of the file, or at the end, or even in a separate file. Most confocal microscopes will capture larger images than this, and most will capture more than one channel. A three-channel, 2048 × 2048 pixel image (routine on any current system) will require 12 MB to store one plane. A three-dimensional (3D) image data set could easily contain 100 or more planes, thus requiring 1200 MB (1.2 gigabytes, GB) or more to store. At the time of this writing, current personal computers typically have 80 to 200 GB hard disks, a 200-fold increase on the norm when the last edition of this chapter was written, but still not enough to be regarded as a permanent store. To provide archival storage, we must have some form of removable storage media.

Data Compression

Before considering the contending bulk storage devices, is there any way we can reduce the size of the problem? Can we compress the image data to make it smaller? Lossless data compression systems, which preserve the integrity of our original image data, generally work on some variation or combination of three well-known algorithms. Run-length encoding (RLE) looks for sequences of identical values and replaces them with one copy of the value and a multiplier. It works very well with binary (black/white) images or simple graphics using a few bold colors, and is used, for example, for all the splash screens in Microsoft Windows.

Lempel–Ziv–Welch (LZW) and Huffman encoding look for repeated sequences, assign each a token, then replace each occurrence by its token. Neither of these works well with real images (though they do an excellent job with computer-generated graphics). Thus, if you save a confocal image as a GIF (graphics interchange format) file, or as a compressed TIFF (tagged image file format) file, both of which use LZW compression, you will be lucky to get even a 10% to 20% decrease in size, and sometimes your file size will actually get larger. You will not do much better with the popular archiving systems PKzip, gzip, or WinZip, which (to avoid patent problems) use LZ77, an earlier Lempel–Ziv algorithm, and Huffman encoding (Deutsch, 1996), though these systems do at least recognize if compression is not working and insert the uncompressed data instead, so your file should not get larger.

There are a couple of exceptions to this generalization. First, some confocal microscopes store 12 bits of data at each pixel (4096 gray levels), but they store this as 16-bit numerical values. Clearly these images have redundant space — a quarter of the file contains no information — and they will therefore at least compress to 75% of the original size. The file will nevertheless become even smaller, often with little or no real loss, if it is converted to 8-bit data. Second, even though it may not immediately be obvious, a three-channel image of moderate size, saved as a 24-bit RGB file, must always have redundant information. Twenty-four bits of data can specify 16.7 million colors, but a 512 × 512 image with only a quarter of a million pixels can contain at most a quarter of a million colors. Efficient algorithms will automatically find this redundancy and yield effective compression (how this is done is explained in the description of the PNG format, below).

Guy Cox • University of Sydney, New South Wales 2006, Australia

PNG, which stands for portable network graphic, but is pronounced "ping," is a lossless compression system (Roelofs, 2003). It will usually offer the highest lossless compression currently attainable for confocal images. The formal compression system is identical to that of the "zip" systems; the deflate algorithm (Deutsch, 1996) a combination of Huffman and Lempel–Ziv algorithms that in essence looks for repeated patterns. The secret of PNG's improved performance lies in its prefiltering of the image to establish the best way to represent the data. In a real-world image of any kind, the difference between adjacent pixels will rarely be extreme so often the data can be reduced substantially by storing only the difference. The different filters vary essentially in the pixels used for comparison (no filtering, pixel before, or before and after, or before and above, etc). Any implementation contains all filters and so will decode any image, but the better implementations will offer improved compression by careful choice of which filter to apply. (The standard allows different filters to be used on each line of the image if required.) So if lossless compression is important it may be worth experimenting with different vendors' implementations of PNG (see below). It tends to be much slower than LZW to compress, partly because it is a two-pass process but mainly because, to get the best results, the program should test which algorithm will give best results. Decompression is fast (see Table 32.1).

The demands of computer multi-media have led to the development of compression techniques specifically aimed at real-world images, both still and moving. Unlike the compression techniques mentioned above, which are completely reversible, these approaches discard information from the image. The picture created after compression and decompression will not be the same as the original. However, very large file compressions can often be achieved with losses which are barely detectable to the eye, though they may affect numerical properties of the image.

TABLE 32.1. Time to Compress and Read Back an Image Using Different Techniques

Compression	Save Time (s)	Read Time (s)	File Size (KB)
Uncompressed TIFF	5	3	9220
LZW TIFF	7	5	6014
PNG	40	5	2861
Wavelet (lossless)	9	12	1934
Wavelet (high-quality)[a]	16	10	841
Wavelet (high-quality)[b]	8	10	821
Wavelet (low-quality)[a]	15	6	9
Wavelet (low-quality)[b]	6	6	10
Lossless JPEG	8	7	3877
DCT JPEG (high-quality)	4	4	840
DCT JPEG (low-quality)	4	4	164

[a] Specifying required quality.
[b] Specifying required file size.
The image used was that seen in Figure 32.1, but scaled up (using bicubic interpolation) 6-fold to 3072 × 3072 pixels in order to make the times measurable. All conversions were done using Paint Shop Pro version 8 (Jasc Software); the results should only be taken as relative and will vary greatly with processor speed. Scaling the image means that it contains substantial redundancy and therefore the compression levels achieved are unrealistic; the file sizes are given mainly to illustrate the trade-off between processing time and disk access. PNG was by far the slowest in compressing the image, but was rapid to read back. The processing requirements of DCT JPEG compression were more than compensated for by the reduction in disk access, so that it was very fast, but lossless JPEG was slower and its compression did not match lossless wavelet or PNG. Wavelet compression (JPEG 2000) showed the curious result that selecting a "compression quality" gave much longer save times than selecting the "desired output file size." At equivalent final sizes, the resulting images seemed similar. This is probably a quirk of the implementation of what is, at the time of this writing, a very new standard. Wavelet images were the slowest to read back, particularly at high image qualities.

The most common still image format is the Joint Photographic Experts' Group (JPEG) compression protocol (Redfern, 1989; Anson, 1993; Pennebaker and Mitchell, 1993), which is supported by many paint and image manipulation programs. This breaks the image into blocks of 8 × 8 pixels, each of which is then processed through a discrete cosine transform (DCT). This is similar to a Fourier transform, but much faster to implement, and gives an 8 × 8 array in frequency space. The frequency components with the lowest information content are then eliminated, after which high-frequency information (fine detail) will be selectively discarded to give the desired degree of compression. The remaining components are stored (using Huffman encoding) in the compressed image. The amount to be discarded in frequency space can be specified, which gives the user control over the trade-off between image quality and degree of compression. Typically, monochrome images can be compressed down to one fifth or less of their original size with no visible loss of quality (Avinash, 1993). Compression and decompression are similar operations, and require similar amounts of computer time. Ten years ago, when the standard was first published (Pennebaker and Mitchell, 1993), the time required was quite noticeable but with a modern processor the reduced amount of disk access will more than compensate for the processing time (Table 32.1).

Color images can be compressed further than monochrome because luminance (brightness) and chrominance (color) are treated separately. The eye can tolerate a greater loss of information in the chrominance signal, so this is normally handled at half the resolution. (The standard allows many different options here but specific implementations usually do not make these evident to the user.) This has certain consequences in confocal microscopy because a three-channel confocal image is *not* a real-color, real-world image but three images which are largely independent of each other. A three-channel confocal image compressed as a color image will look quite adequate but should not be used reliably for numerical analysis; for example, the lower resolution of the color information would make many pixels show colocalization when in fact there is none.

The JPEG standard itself specifies a compression technique, not a file format. As such it is used in many different situations (including one of the compression options in the TIFF standard and in programs such as Microsoft PowerPoint). However, it is most familiar to the end user in the form of files conforming to the JFIF (JPEG file interchange format) standard, which typically use the suffix .jpg. JPEG compression is designed for photographic images so that it only manipulates gray-scale or true color (RGB) images. Adding a false-color palette to a gray-scale image will make it less suitable for JPEG compression because the JPEG algorithm would convert it to a full color image, tripling its size, before compression. Lossless JPEG compression also exists; there have been two distinct lossless compression modes specified in the JPEG standard over the years, but these do *not* use DCT to compress the image and typically do not perform very well, so they have not become popular. The current version, JPEG-LS, uses a predictive algorithm formerly called LoCo, and is designed to be both fast and easy to implement.

Other specific image compression techniques show considerable potential but have yet to achieve the popularity of JPEG (DCT). Fractal compression, a proprietary technique developed by Iterative Systems Inc. (Anson, 1993; Barnsley and Hurd, 1993), creates mathematical expressions which, when iterated, recreate the original image. It can give spectacular levels of compression. Unlike JPEG compression, creating the compressed image is a very time-consuming process but decompression is very quick.

This has made it most useful for such items as CD-ROM encyclopedias but its initial promise has not led to widespread adoption.

Wavelet compression is currently the hot topic in image compression and will undoubtedly be in common use throughout the lifetime of the current edition of this book, though at the time of this writing it is only just appearing in the latest releases of mainstream implementations. It is, in a sense, mathematically comparable to JPEG in that it separates the frequency components in an image, but it works in real space rather than reciprocal space. The basic idea of separating an image into components of different resolution and discarding the lowest information content and highest frequencies first is similar, but it is achieved by passing a series of filters over the image at a range of different scales. The filters — wavelet filters — are the key to this, and are designed to be reversible. The claim is that wavelets can offer useful compression without loss, and much greater compression with losses that are not obvious to the eye. Other advantages include the ability to rapidly generate a low-resolution image (using the coarsest wavelets) and fill in the detail afterwards.

Wavelet compression can treat an image as a whole or break it down into blocks which are compressed individually. The JPEG has introduced wavelet compression into a new version of the JPEG standard (JPEG 2000), and it is in this format that most mainstream applications will offer wavelet compression. In the interests of speed and portability (wavelet compression is intrinsically slower than DCT), the JPEG 2000 implementation uses only two wavelet filters, one for lossless compression and one for lossy compression. Even so, the time required is quite noticeable even on a fast computer (Table 32.1). While a wider range could offer better performance by finding the best wavelet for each image, the practical difficulties involved were deemed to make it not worthwhile. Also, in the JPEG implementation the image is broken into blocks before compression. A major criticism of the DCT JPEG standard was that the 8×8 blocks could often become visible at high levels of compression and JPEG 2000 therefore offers variable sized blocks within a single image, so that one compression level can be applied to featureless regions (such as sky, or the background in a confocal image) and another to regions containing fine detail.

In practice, however, wavelet compression does not seem to offer superior performance over DCT for confocal images, as Figure 32.1 shows. Figure 32.1(A) shows a cultured He-La cell labeled with fluoroscein isothiocyanate (FITC) tagged to an antibody against β-tubulin. It is an average projection from 16 confocal optical sections — a 512×512 pixel 8-bit image. Using an average projection rather than a maximum brightness projection improves the signal-to-noise ratio, but it also reduces the total intensity (because so much of the image is dark) and this therefore reduces the number of gray values present (there are only 120 values in this image). Both factors make the image a better candidate for compression. To preserve the visual quality the contrast has been scaled and the gamma changed (see below); these operations simply change the values assigned to each of the 120 tones, they do not change the number of tones and should not affect how it will compress. Figure 32.1(B) is one of the original slice images with no modifications to gray values. It shows more noise than the projection, but contains 248 gray levels, showing that the gain and black level controls had been used optimally to make use of the full dynamic range without overflow or underflow.

The raw image size in each case is 256 KB, and tif and bmp files are 257 KB. An LZW-compressed tif file of Figure 32.1(A) offered a reasonably useful reduction to 170 KB, while a PNG file created with the well-known program Paint Shop Pro (JASC Software) did rather better at 143 KB. The PNG optimizing program Pngcrush (freeware; see Roelofs, 2003) made an insignificant improvement to 142 KB. This is 55% of the original file size and shows that with a restricted gray range and dark noise-free background reasonable compression can be achieved without loss. Lossless wavelet compression (JPEG 2000) was less effective, giving a file size of 168 KB, scarcely better than LZW-compressed tiff but taking very much longer to compress and decompress. Lossless JPEG was comparable, at 169 KB.

As predicted, the original single-slice image [Fig. 32.1(B)] did not compress nearly so well; the LZW version, at 256 KB, was hardly changed from the original size. PNG did better, at 195 KB (204 KB before optimization). But at 76% of the original size it hardly seems worth the effort. It does, though, reinforce the point that PNG is the only format worth considering for lossless compression of confocal images.

DCT (JPEG) compression of the projection [Fig. 32.1(A)] to two different levels is seen in Figure 32.1(C,D). Figure 32.1(C) shows the image compressed to 26.4 KB, around 10% of its original size. While some loss of quality is evident, the image remains perfectly usable and the compression is very substantial. In Figure 32.1(D), compression has been increased to the point where the image is visibly degraded but still recognizable and even informative, though the file size is only 7.7KB, a mere 3% of the original! Figure 32.1(E,F) shows the same levels of compression but using wavelet compression with JPEG 2000. Both are substantially worse than equivalent DCT images. A specialist wavelet compression program (not using JPEG 2000) was also tried, and gave worse results at equivalent compression levels. It seems probable that the relative failure of wavelets to compete with DCT lies in the rather limited range of resolution levels which contain substantial information in these confocal images. The interest lies primarily in the microtubules, all of which are the same size. In reciprocal space, regions with no information will automatically compress to nothing, whereas the wavelet function may perhaps be chosen to treat all frequencies more or less equally because this may be the best strategy for conventional photographic images. There may therefore be scope for a wavelet implementation dedicated particularly to confocal images.

Figure 32.2 shows the histograms of the images in Figure 32.1. In Figure 32.2(A) the missing gray values are obvious, whereas the single optical section [Fig. 32.2(B)] shows a continuous spectrum. At 10% compression the DCT image [Fig. 32.2(C)] shows a similar spectrum, but smoothed and with the gaps in the gray levels now filled. The wavelet version [Fig. 32.2(E)] also preserves the same shape, but is rather more smoothed at the same compression. At 3% of the original size the DCT histogram [Fig. 32.2(D)] is very much changed, while the wavelet one [Fig. 32.2(F)] shows little change from the 10% compression. In each case, the mean value remains unchanged. These figures show that photometric parameters are surprisingly well conserved even at levels of compression that would seldom be used in practice. While wavelet compression affects the histogram more than DCT at 10% compression, it is more accurate than DCT at 3% compression so that even though the image looks worse, its photometric parameters remain closer to the original.

In practice these compression levels would only be used for such purposes as Internet transmission of images. Compression to between 25% and 50% of the original size would give images of more general usefulness, with little visible change from the original. Even essential photometric parameters are preserved. In spite of the current interest in wavelet compression, DCT still seems a

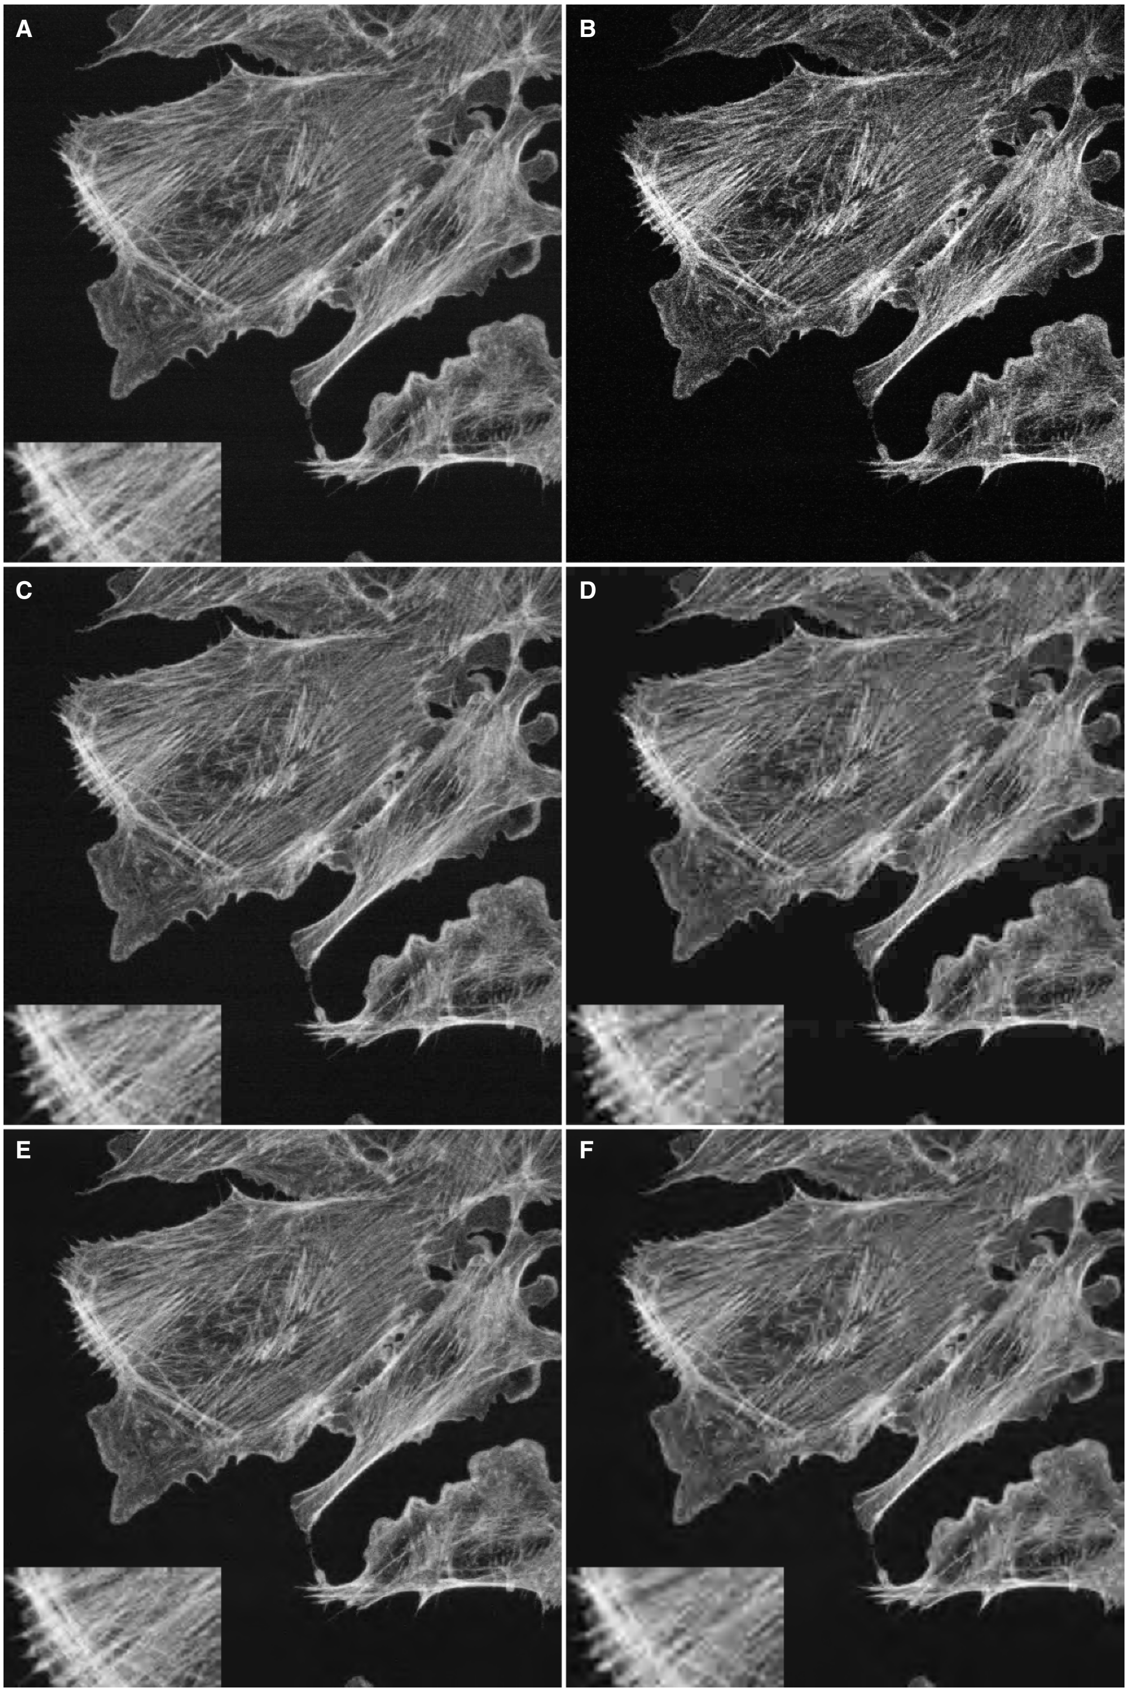

FIGURE 32.1. Effects of image compression on a confocal fluorescence image of a cultured He-La cell immunostained with FITC against β-tubulin. (A) Average projection of the original dataset of 16 optical sections, with contrast scaled and gamma subsequently corrected; original uncompressed image. (B) One optical section from the stack, with no subsequent processing. (C) JPEG compressed (DCT) to ~10% of the original size. (D) JPEG compressed (DCT) to ~3% of the original size. (E) Wavelet compressed (JPEG 2000) to ~10% of the original size. (F) Wavelet compressed (JPEG 2000) to ~3% of the original size. Insets in (A, C–F) are part of the image at 2× magnification to show the losses in compression more clearly.

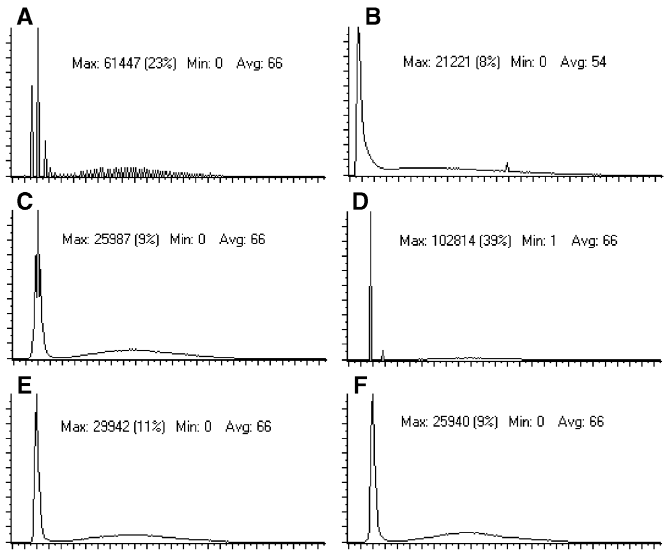

FIGURE 32.2. (A–F) Histograms of pixel intensities in Figure 32.1(A–F), respectively.

better choice for confocal images in cell biology. Not only is it more effective, it is much faster then wavelet compression (Table 32.1). Lossless compression only gives useful results on images with large amounts of uniform background and low noise but in these cases it can be effective. The most likely use would be for storing the output of 3D reconstructions, as in Figure 32.1(A). Although generating a complex 3D movie sequence can take as long as acquiring the original confocal data, and the output files can be just as large, we typically do not have the same concerns about preserving data integrity. It is therefore sensible to use JPEG compression for storing the output.

Some confocal datasets contain only very sparse information. Figure 32.3(A) provides an example, a frame (pre-calcium wave) from a time series of calcium transients induced by testosterone. There were 193 images in the series and without compression this dataset occupies close to 50 MB. However, as only 12% of the pixels lie above the background noise level, the dataset even in its original form compresses without loss to below 100 kB per frame — 40% of the original — with LZW or PNG. If we remove background by setting pixels with a gray value of 14 or below to zero [Fig. 32.3(B)], we have a virtually unchanged image which is now highly compressible without further loss. PNG compression gave a file size of only 49.3 KB, less than 20% of the original. Our original 50 MB dataset will now only be 10 MB. Lossy JPEG compression makes no sense with such a dataset — using a typical setting for reasonable image quality the resultant file size was actually larger (59.3 KB) than the lossless one. What is more, the compression process brought background back into the dark areas. So the message is either use a lossy compression on the original data or compress it by background subtraction and then save it without further loss — do not do both.

Other image manipulations will also affect the compressibility of images. Smoothing, to remove noise, will reduce the high-frequency content and therefore make images more compressible. Deconvolution, on the other hand, aims to restore high-frequency content. This will make images less compressible, or will mean that more is lost in lossy compression. Figure 32.4 illustrates this

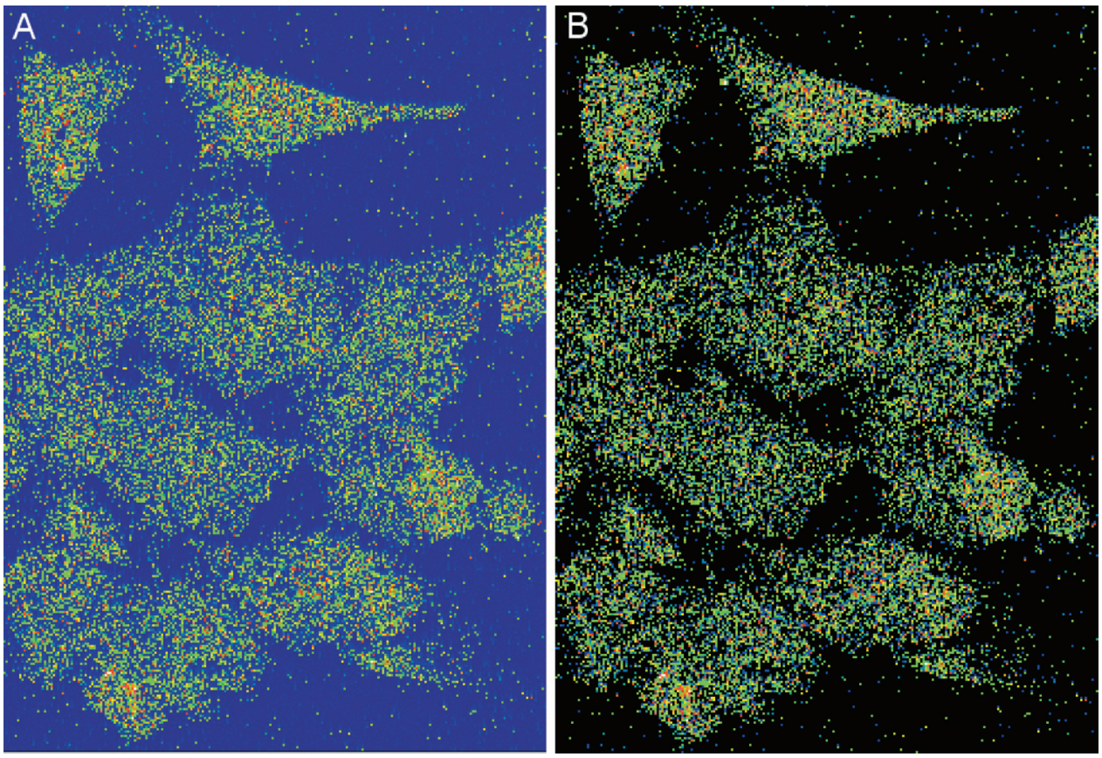

FIGURE 32.3. Calcium imaging (non-ratiometric) of transients induced by testosterone in cultured cells. Pre-stimulation, time point 42 from 193 images taken at 1-second intervals. (A) original image (B) background removed by setting all pixels below a value of 15 to zero, a process that permits no-loss compression to reduce file size of (B) by a factor of two compared to (A). A false-color palette has been added to show how the background has been set to black but none of the "data" pixels have changed. Taken using a 63x/NA 1.2 water-immersion lens. Image width as shown (only part of the original image) is 154μm. Image courtesy Dr. Alison Death.

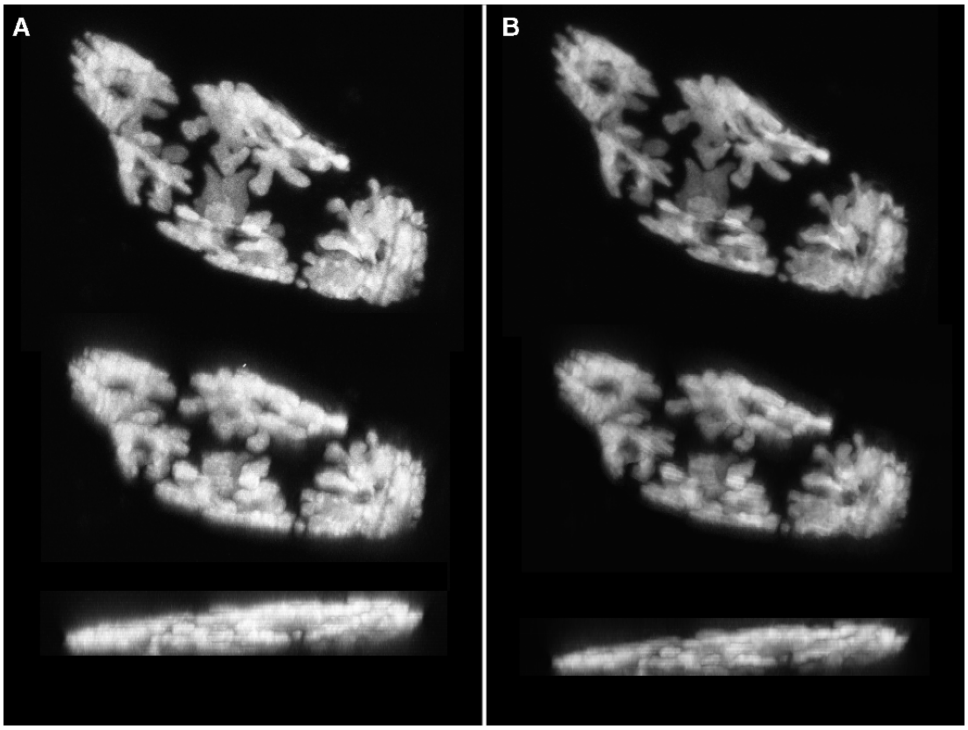

FIGURE 32.4. Three different angle projections (0°, 45°, and 90°) from a 3D dataset of the dinoflagellate alga *Dinophysis*. (A) Maximum intensity projection from the original dataset. (B) Maximum intensity projection after smoothing (3D median filter) and one-dimensional deconvolution.

point. This is a 3D dataset of the dinoflagellate alga *Dinophysis*, which was collected at 3 pixels per resel and therefore is slightly oversampled. This provides the opportunity to smooth the data down to the Nyquist limit and thereby reduce noise without adversely affecting resolution. Figure 32.4(A) shows maximum intensity projections, at three angles, from the original dataset. This type of projection leaves noise unchanged so the view shows an accurate impression of the noise content of the original set. Compressed with LZ77 the original 4.5 MB dataset reduces to 1.9 MB, a useful saving, reflecting the large proportion of background in the set. When the entire dataset is smoothed with a median filter, acting in three dimensions (Cox and Sheppard, 1999), it becomes more compressible, now reducing to 1.38 MB. If we deconvolve this dataset we can restore some of the resolution lost in the *z*-(depth) direction by the transfer function of the microscope (Cox and Sheppard, 1993, 1999). As expected, it is now rather less compressible, at 1.47 MB, but this is still a useful saving on the original. The smoothed, deconvolved dataset is shown in the same projections in Figure 32.4(B).

In any image compression strategy, it is important to bear in mind that confocal images can become virtually meaningless if the information about the acquisition is lost. Some confocal systems (e.g., Bio-Rad) store this data in a header within the same, single file as a series of optical sections. Even if the slice images are exported by the Bio-Rad software, the acquisition data is not exported and the images cannot be re-imported for subsequent processing. Other systems (e.g., Leica, Zeiss) store a database of information about the images — exported images generated within the acquisition software will still retain some of this information but typically 3D reconstructions can only be done from the original images. In either case it is important to ensure that the all-important image acquisition data are preserved, and if possible that the images can be restored to their original file name and type.

A final point: the most common waste of disk space consists of storing completely featureless areas! If your sample is rectangular, select a rectangular window to image it rather than collecting a strip of nothing on each optical section. And do not collect three channels if you have only two labels! Modern systems make it all too easy to accept the default method, or configuration; it will save a lot of time in the long run if you spend a minute or two changing settings to collect only what you want.

Removable Storage Media

Storage media can be divided into those which are sequential (records are written and read from one end only) and those which are random access (it is possible to move directly to any record, whenever it was written).

Sequential Devices

Sequential devices are tapes of various formats and sizes storing up to 200 GB on a single cassette. Tape is still the largest-capacity bulk storage medium available, but is no longer competitive in cost with optical storage. As an image storage system, it also suffers from the time taken to locate and recover any one file. A single file cannot be erased and replaced by another; one must erase either the whole tape or a large group of files, depending upon the recording system. Also, although it is rewritable it will not stand an infinite number of uses. The tape surface has a much harder life than the surface of a disk — it comes into direct contact with the recording heads and capstans, and is coiled and uncoiled each time. Even reading files repeatedly wears the tape, and its long-term archival potential is dubious. Once tape drives were regularly used for data storage and transfer but now their use is almost exclusively for backup purposes — making a copy of a complete file system or subsystem which will typically be read only once, in the event of a hard disk failure.

Modern tape systems are very specifically designed for this task; their purchase cost is high but cost per megabyte stored can be low compared to other rewritable media. This gives them some attraction for long-term archival storage of images that will not need to be accessed regularly, and for very large collections of images. Dumping a 40 or 100 GB hard disk full of images on to a single tape will be much quicker and simpler than writing to dozens of compact disks (CDs) or digital video disks (DVDs). However, most tape systems now rely on specific software to handle them and both this software and suitable hardware will need to be available for the tape to be read in the future — past experience suggests that this will limit effective use to 5 years or so, and this is probably the realistic limit for tape life also.

Transfer rates up to 24 MB s^{-1} are available on expensive high-end systems, although systems designed for small computer use will offer no better than 3 MB s^{-1}. At 24 MB s^{-1} writing one CD worth of data will take only 30 s, but it will take a quarter of an hour to copy a 20 GB hard disk. At 3 MB s^{-1} that same disk will take almost 2 h to copy.

Manufacturers typically quote compressed capacities for their tape drives, based on a notional 2-fold compression ratio that they expect to achieve with their archiving software. This is unrealistic when dealing with image files, and when evaluating competing systems, it is important to compare actual, uncompressed, storage capacities; this is much closer to the figure achievable with microscope images.

Random-Access Devices

Random-access devices comprise a range of disk media, either magnetic or optical, and solid-state devices.

Magnetic Disks

The oldest and simplest of removable media, rewritable, random-access systems is the humble floppy diskette. These are now virtually obsolete, limited by their small capacity — 1.4 MB in the only (marginally) surviving 3.5″ version. As many will have found out, finding a drive to read the once ubiquitous 5.25″ disks is already difficult. In any case, they are too small to be relevant for confocal images.

Various types of super-floppy have had a vogue in the past, but the only current survivor seems to be the Iomega Zip disk, which originally held 100 MB but now comes in capacities up to 750 MB. These are robust and durable but seem unlikely to be current for very much longer, driven out by far cheaper optical technology. They are also too limited in space to meet most modern needs for confocal image storage. Cost per gigabyte is around US$20–100.

Other removable platter magnetic devices have been current, and suffer from the same limitation that in the course of time there may no longer be hardware available to read them. One of the most successful at the time of writing is the Orb drive, available in capacities from 2 to 5 GB. Like many other portable devices they connect to the host computer by the USB (universal serial bus) port, or the parallel printer port. Parallel port connection is relatively slow and USB is by far the preferable option. Cost per gigabyte is of the order of US$10 to US$20, so it is a reasonably affordable option.

There remains the option of just using conventional hard disks. Mounting kits are available to fit a conventional disk in a pull-out mount; disks are also available in cases for connecting to USB, Fire Wire, or SCSI (small computer systems interface) ports, and there are microsized ones which fit the PC card (PCMCIA) slot in notebook computers. The recent fall in price and increase in capacity of hard disks has made this a surprisingly affordable option (below US$1.00 per gigabyte for IDE disks, more for SCSI). Data transfer is as fast as the disk — certainly faster than most other options — and rewriting capacity is effectively unlimited. The long-term potential is less certain because the durability of the system depends not only on the longevity of the magnetic medium, but also on the lifespan of the motor and heads.

Optical Disks

In the previous edition, devices such as WORM (write once, read many) and MO (magneto-optical) disks were discussed. These, like so many technologies, are not only dead but virtually forgotten except by those laboratories which have a huge stock of the disks! However, optical technology is certainly the current preferred option because there is good reason to have faith in the archival durability of the media. Furthermore, mass-market devices now have sufficient capacity to meet many users' demands so that one can have some confidence in the longevity of the technology.

Compact Disks

Compact disks (CDs) have already been with us for over 20 years, and writable CDs for 10. The cost, high when the previous edition was written, is now very low both in first cost and media (around US$0.70 per gigabyte). Speed, though it has increased about 12-fold since then, is still the major problem. The rate of data transfer for an audio CD is a rather pedestrian 150 KB s^{-1}, and this is referred to as single speed. Read and write speeds up to 52× this base value are now available. A complete 700 MB CD can thus be written in 5 min or so, and modern software will adjust the writing speed on the fly so that the need to maintain a constant data stream is less of an issue. This means that CDs can now even be written across a network, though this will inevitably carry a speed penalty. It may still be preferable to carry the additional overhead of first copying files to the writing computer.

Rewritable CDs are also widely available at a cost only a little higher than conventional single-use CDs. Erasing data for re-use is, however, a relatively slow process. They may be useful when images are to be stored for a short time only, but for long-term archival use it would seem wiser to use single-use disks. Many manufacturers have conducted accelerated-aging tests on their single-use CDs and their security as archival storage seems to be the best of all mass-market computer media. It seems inevitable that rewritable disks could not offer equal security, and the risk of accidental deletion is always present with any rewritable medium. In fact, the time spent trying to decide which files can be overwritten is usually worth much more than the disk space saved.

Various formatting options now allow multiple use, either by writing multiple sessions (which does carry an overhead of about 15 MB per session) or by using the packet CD format, which allows a CD to be treated almost exactly like a conventional mounted drive. Multi-session CDs can be read on most systems but the packet CD format cannot. Because it reduces compatibility with other systems and has little point when the content of a CD is relatively small compared to a modern hard disk, packet CD has not become widely popular.

One of the limitations of the CD format is its handling of file names. The standard laid down by the International Standards Organization (ISO) requires file names to fit an 8 + 3 character format similar (but not identical) to that of MS-DOS. ISO-compatible CDs are readable on Apple, PC, and Unix computers, which is very convenient for data exchange. Unfortunately, most confocal microscopes give files and directories (folders) much longer names. Extensions to the standard allow for longer files names in both Macintosh and Windows computers, but these are unfortunately not cross-platform compatible. Because most confocal microscopes use Windows it is important to use the Joliet extension which caters for these file names, otherwise the disk will contain a useless collection of truncated names, particularly with microscopes such as current Leica models, which save each plane and channel as a separate file, and rely on a database program to identify these images. On an ISO disk, it will be impossible to identify which plane and channel are which, and the data becomes completely useless.

Because CDs are likely to be a major archival medium, in the medium term at least, the question obviously arises as to how permanent they are. Pressed CDs have a polycarbonate blank into which the pits are pressed to carry the information. This surface is then coated with an evaporated metal layer, and then a coat of varnish and the printed label (Fig. 32.5, upper). The CD is read through the thickness of the blank. If the clear side of the blank gets scratched, it will hinder reading but it can often be repolished. The label side is more vulnerable because only a layer of varnish and the printed label lie between the data and the outside world. Recordable CDs have a dye layer between the polycarbonate blank and the metal film and it is this which is modified by the writing laser beam (Fig. 32.5, lower).

In terms of pressed CDs, excluding physical damage, the key issues are the aluminum reflective coating (which can get oxidized, particularly if there are any flaws in the varnish) and the polycarbonate blank. So far, no plastic seems to last forever and I doubt if polycarbonate will stay clear and flexible indefinitely. However, as polycarbonate is vulnerable to almost all organic solvents, excluding light and solvent fumes will doubtless help.

Archival quality recordable CDs usually use something better than aluminum. Several manufacturers offer silver, silver + gold, or pure gold. Obviously pure gold should be highly stable, but it is less reflective, which increases the risk of read errors. Whatever metal is used, archival life still depends on the dye layer in front of it remaining stable. The claims made by the manufacturer for their different dyes (typically cyanine, phthalocyanine, or azo) are difficult to evaluate. The cyanine dyes used in the earliest recordable CDs were rather vulnerable to bright ambient light. Some manufacturers have chosen to concentrate on extending the durability of these dyes, whereas others have turned to alternative dyes such as azo or phthalocyanine. All archival tests depend on accelerated aging (typically at higher temperature) and, while this is valid up to a point, it is unwise to trust it too implicitly (Nugent, 1989; Stinson *et al.*, 1995).

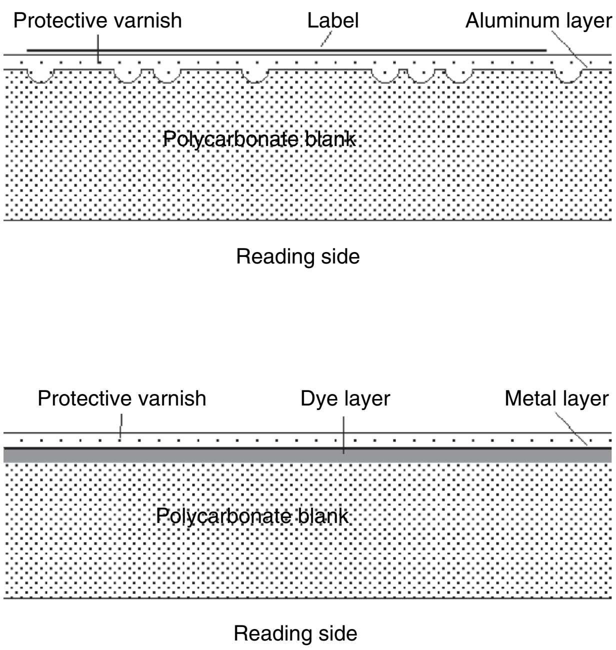

FIGURE 32.5. Structure of a pressed (above) and recordable (below) CDs (not to scale).

The speed at which drives will write CDs has increased enormously over the years, with 52× now routine. This has placed pressure on manufacturers to increase the response time of the dye layer but it would seem logical that a dye which can be bleached at 50 times the original speed is unlikely to be as archivally stable as the older disks. Often the layer is made much thinner to enable the high-speed writing. Of course, dye technology is also evolving but, if an archival-quality disk will not support the latest writing speeds, there may be a good reason.

Because one is likely to write or put a label on the back, the varnish is important. Most makers object to labels even though labeling kits are widely sold. The varnish is typically water based because the polycarbonate of the disks is very vulnerable to solvents. This leaves one in a cleft stick as to how to label it because water-based inks may loosen the varnish but solvent ones may attack the disk! Different manufacturers vary in their recommendations and the safest approach is to follow the recommendation for each particular brand. Quality disks will have an extra writable protective layer over the base varnish giving you a bit of extra security, and this is well worth having.

Kodak recommends that CDs not be stacked adjacent to each other or to any other surface. They should therefore be stored in "jewel cases" or in a custom storage box which separates the disks, and not kept in envelopes or stacked on a spindle.

Blank CDs are now so cheap that the cost of storage is below US$1 per gigabyte, depending on the brand and quality of the media. Common sense suggests that, however reasonable it may be to choose cheap disks when just sending data through the post or taking it from laboratory to laboratory, saving a few cents by choosing unknown brands is a false economy if the intention is archival storage. Writers are very cheap and quite fast (48 speed corresponds to 7 MB/s, comparable with modern tape systems). Best of all, every computer can read the disk without extra hardware. The huge range of commercial CD-ROMs ensures that readers will remain available for many years, so that archival material will be accessible as well as secure.

Digital Video Disk (DVD)

DVD (digital video disks) represent the next stage of optical disk technology. Using similar technology, but shorter wavelength lasers so that resolution is better, 4.7 GB can be stored on one side of a disk with the same size as a CD. Because the optics that read the disk are confocal, a DVD can carry two separate layers of information, thus storing over 9 GB, but recordable two-layer disks are only just coming on to the market at the time of this writing.

DVD-R write-once DVD disks — are now reasonable in price, at around US$1.00 per disk from the cheapest sources, and the writers are now reasonable at US$200 to US$300 (a mere 1% of their price 4 years ago!). Two standards (DVD+R and DVD-R) exist for these disks. This has hindered their general acceptance, but newer players will handle both (Nathans, 2003). At 4.7 GB, it is clear that DVD is now both cheaper and more convenient than CD-R for storage of confocal images, though the question of the long-term stability and longevity of the format is still not so well known as that of the CD format. Nevertheless, DVD players are now common domestic appliances, so it seems likely that the standard will be durable.

Rewritable DVD disks also suffer from incompatible standards and because they are of lower reflectivity are sometimes difficult to read in DVD players. The older standards (DVD-ROM and

DVD-RAM) were also lower capacity than 4.7 GB. As with CD-RW, they are probably not the best alternative for archival storage, but could have their place for data transfer. Current drives mostly handle RW and R disks in both + and − options. As with CDs, rewritable disks are always slower to write.

It is only in the past couple of years that the DVD market has really showed signs of maturity. Because most computer drives will read and write the CD format as well, it would seem to be the logical choice when purchasing a new system, and DVD writers are now routine on new confocal microscopes.

Solid State Devices

A development which was not foreseen in the last edition of this chapter has been the proliferation of ultra-compact solid state memory devices which retain data even without a source of power. While small in capacity compared to a hard disk, these range up to more than the capacity of a CD in a tiny fraction of the space. Much of this development has been driven by the explosive growth of the digital camera market.

Compact flash cards (Compact Flash Association, 2003) are used by many such cameras, making the computer accessories to read them an essential. Typically these use either the PC-card (PCMCIA) slots in notebook computers or else the USB or Fire Wire ports found on both desktop and notebook systems. A key feature is that both are designed to appear as hard disks to the computer without the need to install any drivers. As the card is the size of a postage stamp, and about 3 mm thick, it represents a highly portable data store, and many people use them for convenient portable storage or transfer between computers without reference to digital cameras. Flash drives currently can hold up to 2 GB. Data transfer rates of the compact flash (confocal) chips are currently 5 to 7 MB/s, but the latest revision of the interface is designed to cope with rates of up to 16 MB/s to allow for advances in chip technology in the future. Practical test speeds achieved in computers (Digital Photography Review, 2003) are around 3 to 4 MB/s with writing being slower than reading; performance in digital cameras will always be much slower. The cost is still around US$100 a gigabyte so it will not compete with CDs for archival storage, but as a fast, rewritable method of transporting relatively large files — whether images, documents, or digital presentations — compact flash has an important place.

Memory Stick (Sony) and Smart Media (Samsung) are similar, more proprietary flash memory devices which fulfill similar functions, but so far offer a smaller range of useful options than the more open standard Compact Flash. They tend to be more popular in the portable music player market, showing again how several once-different technologies are converging. The Sony Micro Vault is a dedicated USB-only version that comes in capacities from 32 MB to 256 MB, and requires no further accessories; it even has a cover for the plug when removed from the computer. Similar "keychain" memory devices are available from other manufacturers.

These flash memory devices have established a quite different market niche from other removable storage devices, but as photography becomes increasingly a digital process, this convergence seems likely to continue.

DISPLAY

Before looking in detail at how the image is displayed and printed, we should consider the nature of the confocal image (see Chapter 4, *this volume*). The image in a conventional optical microscope has an infinite gradation of tones within it, whereas the confocal image typically has just 256, if it is monochrome or false color. Merged two- or three-channel images may have up to 256^3 colors, but often have considerably fewer. Confocal images have a finite number of pixels, whereas photographic images have limited resolution, but a smooth transition from point to point. In more general terms, a confocal image is quantized in both spatial (x, y, z) and intensity dimensions (see Chapter 4, *this volume*).

What the microscopist actually sees is not the image itself, but a display on a monitor. Both the monitor and the way it is driven will have a major effect on the appearance of the image. This in turn is interpreted by the human eye when we see it directly, or by a camera if we record the image photographically. As a preliminary, we should therefore look at how monitors display confocal images.

Monitors

Monochrome cathode ray tube (CRT) monitors simply have a layer of phosphor coated on the inside of the glass, so that an illuminated spot will be produced wherever the electron beam hits. The resolution of the monitor, therefore, depends solely on how small the electron beam hitting the screen can be. Color monitors, on the other hand, have red, green, and blue phosphors arranged either in dots (shadow-mask tube) or stripes (Trinitron tube). The image on a color monitor will always be made up of a mosaic of the three primary colors; the finer this mosaic, the better the image will be. This is specified by the dot pitch of the tube in millimeters — 0.28 mm would be a typical value for a good quality modern PC monitor, though pitches as small as 0.18 mm are available, and cheaper or older monitors will have pitches up to 0.4 mm. These are absolute values, so a larger monitor will have more dots in the total width of the image.

The number of pixels that may be displayed on the monitor is a function of the speed at which the electron beam can respond to a changing signal, and is not related to the actual dot pitch of the CRT. Thus, a confocal image displayed on a color monitor will have each pixel subdivided into a pattern of red, green, and blue dots. If the pixel spacing of the data comes close to the dot pitch on the screen, *aliasing* (below) may occur, creating undesirable effects. Many color monitors can be set up to display more pixels than there are dot triplets available, so the full resolution of the image cannot actually be shown. Thus, a large monitor is essential on a confocal microscope if we are to be able to see the detail in a high-resolution image.

A CRT-based monitor is intrinsically capable of displaying an almost infinite number of colors.[1] However, the video board inside the computer imposes its own restrictions. Display boards suitable for a confocal microscope will permit 256 gradations in each primary color, so that a 24-bit (three 8-bit channels) confocal image can be displayed without compromising the intensity range. (12-bit or 16-bit images will, however, need to be reduced to 8-bit for display.)

The other factor determining the appearance of the image is the frequency with which the display is scanned. The more rapidly the screen is refreshed, the less the image will flicker, and a suitable monitor should redraw the entire image at least 70 times per second. Low cost boards will often compromise one or other of these attributes at their highest resolution and are therefore inher-

[1] Limited only by Poisson noise as it affects the charge deposited by the electron beam during the pixel dwell time and the number of phosphor grains in each dot, etc.).

ently unsuitable for confocal use, but because high-end display boards are now very cheap compared to confocal microscopes, this is not likely to be a problem so long as it is understood that just any computer will not do.

Displaying large numbers of pixels on low-priced monitors can reduce the refresh rate to 60 per second, which is about the lowest tolerable value. Interlaced scanning is a strategy used to reduce flicker when it is not possible to scan the entire image at a sufficient rate. First, the odd lines of the image are drawn, then the next scan draws the even lines. This technique is primarily used for broadcast video signals to enable the signal to fit into the available bandwidth, but it has in the past also been used to obtain higher resolution on low-cost computer monitors.

International television standards use 625 scan lines per frame, with each interlace drawn 50 times a second (and thus a full-frame rate of 25/s). The system used in the Americas and Japan uses only 525 lines, but a faster refresh rate of 60 interlaces per second. (Not all scan lines are available for display; standard video can display only 512 pixels vertically, US video only 483.) Video displays once played a significant part in confocal microscopy (they were standard, e.g., on the widely popular Bio-Rad MRC 500 and 600) but they are not used now. Apart from the low resolution and refresh rate, the problem of displaying multi-channel images adequately led to their demise. In broadcast television, the color signal is encoded as a chrominance signal at much lower resolution than the monochrome luminance signal, and this does not give adequate quality for a multi-channel confocal image. The alternative is to generate a three-channel video signal that will give a much higher quality display (on the same monitor so long as the appropriate inputs are provided). However, this signal cannot now be recorded on a standard video cassette recorder (VCR) or printed on a low-cost video printer. The utility of video in microscopy is primarily in recording fast-moving items and it may still have a part to play with direct-vision confocal systems (Nipkow disk or slit scanners), but it is no longer relevant to point-scan systems.

In a confocal image, pixel intensity values are linearly related to the numbers of photons captured from the specimen. However, this linearity may not be preserved when the image is displayed. In the simplest case, if the value of the pixel is converted directly to the voltage at the control grid of the CRT, the actual brightness of the pixel on the screen will be proportional to the three-halves power of the pixel value. This may not be all bad because the human eye responds logarithmically, not linearly, to light (Mortimer and Sowerby, 1941). However, confocal images often look excessively contrasty on an uncorrected display, and fine detail in the mid-tones will be lost.

This relationship may be modified by more sophisticated display electronics, and high-quality display cards typically come with software to allow the user to set up the display optimally. These are all too often ignored by researchers who "don't have time" (and then waste much more time struggling with pictures which fail to show details which "were there when they took the picture"). Failing such an option, image manipulation programs such as Photoshop, Paint Shop Pro, and Corel PhotoPaint place the display gamma under user control; but there again the user must make the effort to use that control (for more discussion, see Chapter 4, *this volume*).

Liquid Crystal Displays

Flat screen (liquid crystal) monitors are inherently different from CRTs. They use liquid crystal devices between crossed polarizers to display the image, and as each pixel is addressed independently, the question of display resolution not matching the pixel resolution does not arise. Typically different resolutions are not available on LCD monitors; in the rare cases that they are, the pixels are remapped in software (see below and Fig. 32.6). Therefore, LCD monitors will often give a crisper display than CRTs, though on the other hand the pixels may be more visible simply because their edges are more clearly defined. Some older displays (typically used on lower cost notebook computers) used passive supertwisted nematic (STN) displays instead of active thin-film transistor (TFT) technology. These are both cheaper and far less demanding of power. However, they may not offer full 24-bit color, and the image may be less bright and have a smaller viewing angle. Large freestanding monitors are always TFT.

Large LCD monitors have many advantages in the confocal laboratory. Because the display is not continually redrawn as in a

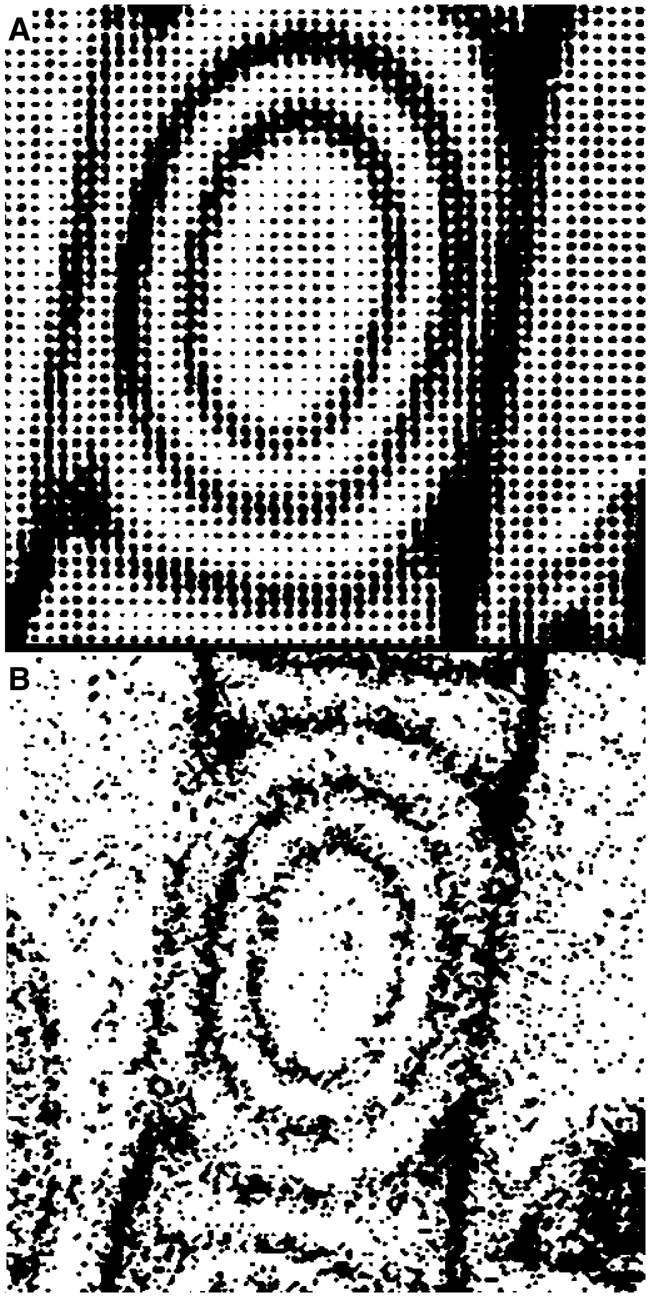

FIGURE 32.6. Halftoning (A) versus dithering (B). Highly enlarged view of part of an image of an integrated circuit chip; above printed by halftoning using a 4 × 4 matrix of laser dots, below printed by dithering using a 3 × 3 matrix.

CRT, flicker is not an issue. Consequently, neither is refresh rate (it can become an issue with video images). The screen is flat and compact, and as confocal systems often include two monitors and three lasers in a small room, minimizing heat is worthwhile. However, LCDs are costly, and the cheaper models often sacrifice color quality, viewing angle, or both. These are sacrifices which are not worth making. Good saturation, wide control over contrast and gamma, and a wide viewing angle are all essential. If you can afford it, buy a pair of top-quality LCD monitors (and don't just take price as a criterion of excellence, check them out carefully yourself). If cost is a major issue, buy high-quality CRT monitors rather than low-quality LCDs.

Data Projectors

Data projectors are now a very common display format for confocal images, and naturally they often represent the occasion when high-quality display is most important. However, they typically have lower resolution and often a poorer gray-scale rendition than computer monitors. In terms of resolution, 800 × 600 pixels (SVGA) is common on projectors intended for the home market and 1024 × 768 (XGA) on ones intended for academic and teaching use, though higher resolutions are available at correspondingly higher prices.

Two different technologies are used in these projectors. A very good description is given by Powell (2004). LCD projectors use liquid crystal screens, as in flat panel monitors, except these do not have a color mosaic. Instead, three panels are used, one for each of the primary colors. Digital light processor (DLP) projectors use a micromirror array, where the pixels are tiny mirrors and these are tilted to send more or less light to the image. Very expensive projectors use three DLP chips, one for each primary color, but these are rare. The projectors one is likely to encounter in a lecture room or conference have one DLP chip, with a rapidly spinning filter wheel in front of it. The colors are therefore generated sequentially and merged by persistence of vision.

The two types have their own strong and weak points, and typically DLP projectors are favored for home theatre use and LCD for data projection. As a comparison, for this review the signal from a notebook computer was sent simultaneously to two moderately high-end projectors, projecting on to adjacent screens. One was a DLP projector, the other an LCD, and price and luminous output were comparable. Contrast, brightness, and other display parameters were set to midpoint values on both projectors.

The native resolution of both projectors was 1024 × 768, and the computer was set to the same value. Both projectors were able to handle higher resolutions and scale them down, but the quality suffered very markedly when this was done. The first lesson, therefore, is to set your screen display to the resolution of the projector, if possible. Even at the native projector resolution there may still be some pixel re-mapping taking place because projectors correct for keystone distortion caused by a non-horizontal projection angle. This means that either the top or bottom of the image cannot use all the displayable pixels.

The LCD projector gave a much sharper image, which was obviously preferable for fine text. However, its color rendition, particularly on real-world photographic images, was inferior, having a slight color cast and excessive saturation. The DLP projector gave images with a very accurate color rendition, free of any cast and natural in appearance. The two projectors differed markedly in gamma. The DLP projector had a gamma of 1 (measured with the gamma test function of an imaging package), while the LCD projector was around 1.6 (slightly higher in blue and red than in the green). This means that the LCD projector was very comparable to both the screen of the laptop and to a CRT monitor, both of which checked out with similar values, but the DLP is more accurate for confocal images in which pixel value is typically linear with number of photons.

To test the displayable gray scales, a test image with intensity scaling from 0 (black) on one side to 255 (white) on the other was used. All 256 values were present, and on both CRT and notebook monitors the change seemed totally smooth. On the LCD projector it also seemed smooth, though with some minor streaks, which may have been aliasing rather than posterizing (see Digital Printers, below). However, there were noticeable bands with the DLP projector. This implies it was incapable of reproducing a full 8 bits in each color, and in fact posterizing was noticeable in large pale areas of scanned pictures. Also relevant in this context is the contrast range of the projector: the difference between its whitest white and darkest black. This is an important figure of merit for a digital projector and is always quoted by manufacturers. The number of tones which can be reproduced has little relevance if they are squeezed into such a small range that the eye cannot distinguish them. In the past this has been a major concern when projecting confocal images, with detail disappearing in both highlights and shadows. This is an area in which digital projectors (of either technology) have made huge strides in recent years. DLP is normally regarded as leading in contrast ratio but in this test both projectors seemed comparable, with good rich blacks.

Both projectors seemed evenly balanced in response time, with rapid mouse movements appearing equally (and acceptably) jerky in both (at 60 Hz refresh rate). At very close quarters some misalignment of the different color images was visible with the LCD projector. This was invisible at normal viewing distance and may be inevitable in a projector with three different LCD arrays (especially one which is regularly transported). This would not be expected in a DLP projector because there is only one display element, but in fact some color fringing could still be seen at the edges of the screen, though not in the center.

The verdict on this test was that the LCD projector was way ahead for text, diagrams, and other computer-generated graphics, but the DLP had the edge for micrographs and other real-world images. This is in line with the commonly accepted merits of the two technologies. The question of different gamma is likely to be significant when projecting confocal images, and in the rare case where one knows in advance which type of projector will be used, the images in a presentation could be adjusted to suit. But the most useful point to remember when giving a presentation at a conference is still to set your screen resolution to the native resolution of the projector.

HARD COPY

When it comes to recording images, the confocal microscopist has to make a choice between two fundamentally different technologies. One option, **photography**, was in the past familiar ground to most microscopists. The other option, computer **printers** of one sort or another, are more likely to be relevant in the 21st century.

Photographic Systems

In the 10 years since the previous edition of this book, photography has almost completely disappeared from the cell biology laboratory. So far as confocal images are concerned, this is all to the good because there is a fundamental mismatch between film and the digital image. A photograph can reproduce far more tones than the 256 that an 8-bit image possesses, but the interposition of lenses and film means that pixel positions are not reproduced sharply and with complete accuracy. Blurring below the level of microscope resolution will not be noticeable in a conventional micrograph, but if pixels are not rendered clearly in a confocal image it will look soft — especially if there is any text superimposed on the image. At one time, screen-shooting devices were standard equipment with confocal microscopes but now they are no more than historical curiosities.

Digital Printers

Printers typically work by putting dots on a page of paper. As printer technology has evolved, the size and resolution of these dots has become smaller and more precise, but in general the dots are still either present or absent, which limits the ability of a printer to represent images with a range of tones. However, these pixels are placed with extreme accuracy, so providing the data is handled properly (below), it is possible for each pixel in a confocal image to be printed sharply and in its correct place. There are two ways in which we can break up a gray-scale image into a pattern of black dots for printing: *halftoning* or *dithering*.

Halftoning is the way images are reproduced in printed books and newspapers. The image is broken into a series of black dots of varying size, darker grays being represented by larger dots. Halftoning is unarguably the method of choice if the resolution of the output medium is sufficient. However, it will be clear that to produce halftone dots of varying sizes, each dot must be a multiple of the basic dot pattern of the printer. If the halftone dots are to be small, the printer's basic dots must be *very* small. The halftone screen in a printed book is typically 133 or 150 dots per inch (Cox, 1987). A 1200 dpi printer can give is 8 × 8 dots — 65 gray shades — within that resolution. To get 256 gray shades at 150 dpi, we need a 2400 dpi printer.

Dithering uses a probabilistic method to decide whether a printer dot should be present or not. If the pixel is dark, there is a high probability that a black dot will be printed; if it is light that probability is low. The effect is of a grainy image without any regularly repeating pattern. When using a low-resolution output device, dithering is the only option; halftoning would result in an impossibly coarse screen. Figure 32.6 shows a magnified view of a confocal image of an integrated circuit, (IC) device, printed by halftoning and dithering. In the dithered image each pixel of the original micrograph is represented nine times, by a 3 × 3 matrix of laser dots, the probabilistic dithering calculation being applied independently each time. Thus, on average, a 50% gray would have either four or five dots black, the others being white, but which of the nine dots were black would vary each time.

The eye can perceive about 64 shades of gray reflected from a solid surface. If fewer shades than this are used to reproduce an image, areas that should show a smooth transition in tone will reproduce as a series of bands. Because the human visual system is extremely sensitive to edges, this can be extremely distracting. This effect is termed *posterizing*. In printing it occurs when the printer is unable to reproduce at least 64 shades of gray. Figure 32.7 shows an example of this. In Figure 32.7(A), using 16 gray tones (roughly what an old 300 dpi laser printer can give) the smoothly-graded gray appears as a series of discrete bands. To some extent, the problem can be reduced by combining dithering with halftoning — values which bridge the boundaries between the levels the printer can produce are randomized to decide which value they should have [Fig. 32.7(B)]. Ultimately, though, good reproduction of a confocal image will require at least 64 gray levels to be reproduced on paper [Fig. 32.7(C)].

Proper reproduction of a confocal image by a printer will generally require the image to be reproduced either pixel for pixel, or with an integer multiple (or fraction) of printer pixels or half-tone dots to each dot of the confocal image. If this is not the case, aliasing (Chapter 4, *this volume*) will generate artifactual patterning in the image. This is shown in Figure 32.8; scaling the original 465-pixel wide image [Fig. 32.8(A)] to fit the common 512-pixel size has given the diagonal lines and the circle a very jagged appearance [Fig. 32.8(B)]. This introduces a considerable constraint on printing confocal images; photography can reproduce them with equal accuracy at any size, but printing works best with integer multiples or fractions of the image pixels. The only ways out of this are either to use such a high-resolution print device that it

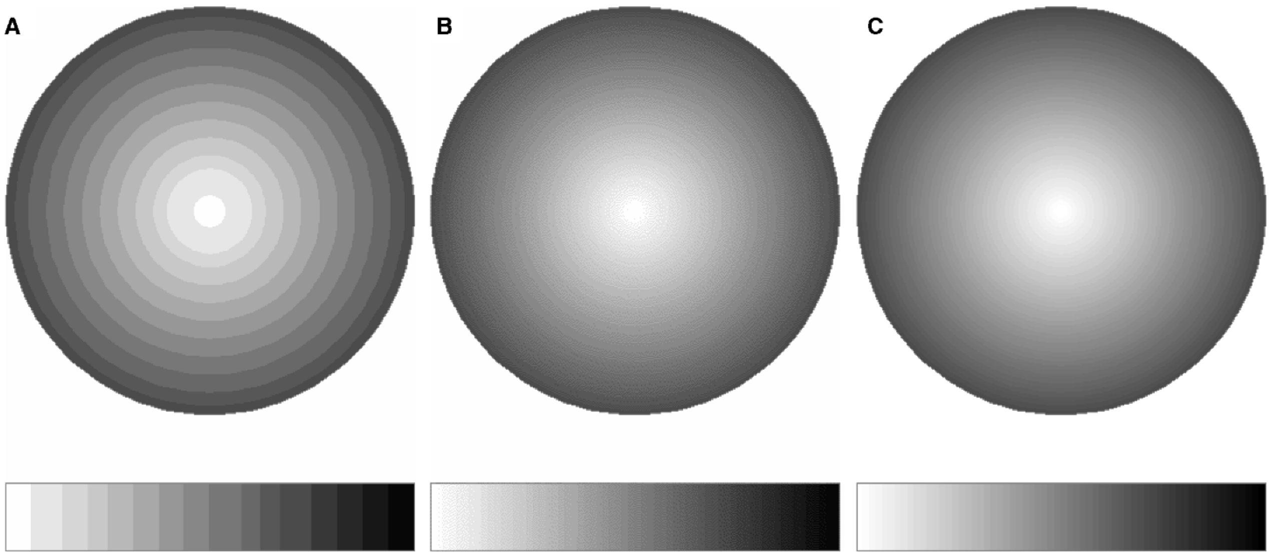

FIGURE 32.7. Posterizing. The smooth ramp of shade in the original shows banding when reproduced with only 16 gray levels (A). This can be partly disguised by dithering, still only using 16 levels (B), but using 80 gray levels is enough to give a smooth-looking result (C).

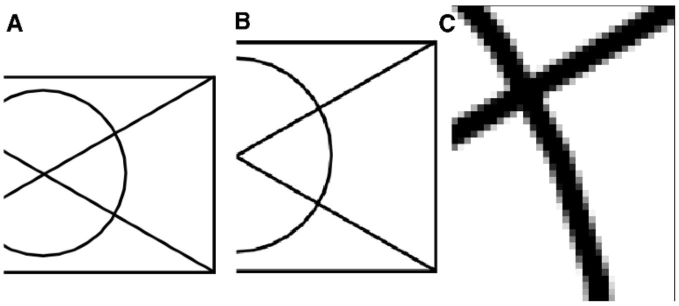

FIGURE 32.8. Aliasing. The circle and the diagonal lines appear as smooth as the horizontal and vertical ones in the original image (A), but when it is scaled up slightly they become jagged (B). In the original image the curved and slanting lines are actually made up of black and various shades of gray, shown much enlarged in (C).

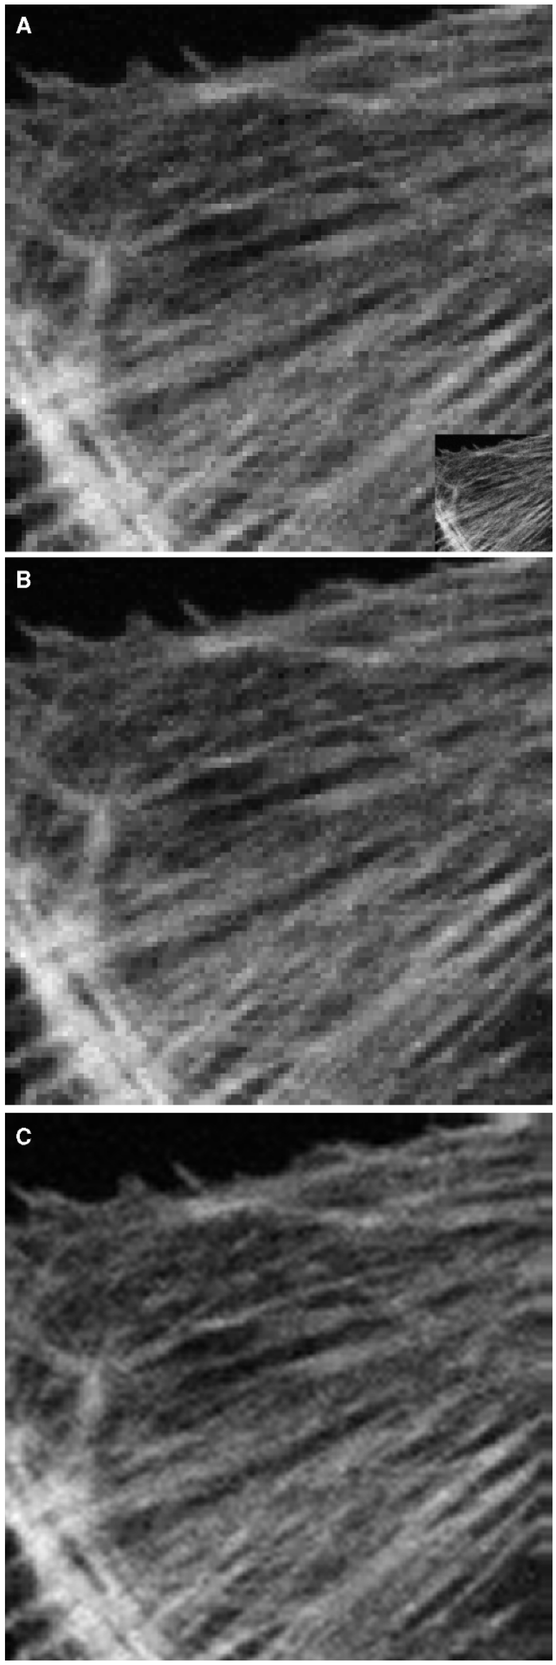

FIGURE 32.9. Scaling techniques. A small part of the image shown in Figure 32.1(A) enlarged by the odd amount of 467%, which cannot be achieved by simple scaling of pixel size. (A) Pixel enlargement, as well as looking blocky, severe aliasing is obvious. (Inset) The original image. (B) Bilinear resampling: aliasing is much reduced but the image is still blocky. Little use is made of the extra pixels now available. (C) Bicubic resampling gives a hugely better result; it cannot produce more resolution than the original data contains, but it does the optimum job of mapping it to the output resolution.

exceeds the Nyquist criterion at the output resolution or else to remap the image with sophisticated software (several algorithms are in general use, and bilinear or bicubic resampling are probably the commonest) to the output resolution. The first option is becoming more common as printer resolution improves but cannot yet be guaranteed. Figure 32.9 shows the effect of different remapping algorithms. The image (shown in the inset) is a tiny part of Figure 32.1(A), enlarged by the odd amount of 467%. Direct pixel scaling [Fig. 32.9(A)] gives, as expected, very poor results, and bilinear resampling [Fig. 32.9(B)] is very little better. Bicubic resampling [Fig. 32.9(C)] is hugely better and obviously the only useful choice in this case.

Aliasing can also appear if the printer's gray levels do not match the intensity levels of the image. When a black diagonal, or curved, line is reproduced in a pixelated image it will appear as a mixture of gray and black pixels [Fig. 32.8(C)]. Pixels which lie wholly within the line are black; those which were partially intersected by the line are gray. So long as they are reproduced at their original intensity the line will retain the illusion of smoothness, but as soon as they are made lighter or darker the line will appear jagged.

Color images present all the above problems, as well as some of their own. Most laser-scanning confocal microscopes (CLSM) do not produce color images in the sense that a conventional optical microscope does. Color images produced on a CLSM are either pseudo-color images in which a false-color palette is applied to a gray-scale image, or multi-channel images in which two or three different signals are each assigned to a different primary color. An image with a fluorescein signal in the green channel and a rhodamine image in the red channel might look very similar to a real color photomicrograph of the same slide, but the way the image is made up is very different.

A further problem arises because multi-channel images, and some false-color palettes, tend to use fully saturated colors. These almost never occur in nature. On a monitor, which emits light, these can look very effective, not least because confocal microscopes commonly operate in dimly lit rooms. They can also make good slides. However, when printed on paper, where the image is created by light reflected from the paper through the ink, the image will look very dark. This problem is exacerbated by the different color models used to form the image.

A computer image is usually stored, and always displayed, using an RGB (red/green/blue) or additive color model. Adjacent points on the monitor screen emit light of the three primary colors, and the different colors are created by adding these together. A printed page uses the CMY (cyan/magenta/yellow) or subtractive model. A red point is created by putting on the paper both magenta (which subtracts green from the reflected light) and yellow (which subtracts blue) so that only red will be reflected. [Commercial printing, and many computer printers, also add a black ink to compensate for the fact that the three color inks may not add up to a perfect black. This is then a CMYK (cyan/magenta/yellow/black) color model.]

Thus, reproducing each of the primary colors that look so brilliant when formed by a single phosphor on a monitor requires the light to pass through two separate color dyes before being reflected by the paper. To put back brilliance into such an image, it is necessary to unsaturate the colors — to add some "white" into them. Many computer image-manipulation programs provide this facility. Often, some experimentation is required before a good screen image can be turned into a good print and it can be very useful to have a program that permits one to print an array of small test images, each made using slightly different settings.

Laser Printers

Laser printers are the workhorse of the modern office. Their crisp black type and graphics are unrivaled for most computer output. For reproduction of confocal (or other) images their abilities are more limited, though they can produce quite reasonable proof prints. Laser printers use a low power laser to write dots directly onto the charged drum of a photocopier.

The limitation of a laser printer is that it is difficult to get the pattern of dots fine enough to reproduce a full range of tones by halftoning (above). However, laser printers with 1200 dpi resolution are now common, and many also have the capability of modulating the size of the spot to some extent. Output from such a printer is adequate for many purposes, and the cost is far lower than either photography or higher-class printing. Their weakest point is that large expanses of black are still not rendered as uniformly as in a photograph.

Color laser printers have only recently started to make a significant impact on the marketplace. This is partly because of a huge decrease in price, and partly because of improvements in resolution (600 dpi is common) and in the tricks used by the built-in firmware, which at last make near-photographic quality routinely attainable. The quality does not yet match that offered by inkjet and dye sublimation printers, but on a cost-versus-quality basis it hugely exceeds it, so that a color laser printer is ideal for such purposes as printing preprints of journal articles in quantity.

Ink Jet Printers

These have long ago taken over from dot matrix as the everyday printer for home and single-user office use. They operate by squirting small jets of ink on the paper (for best results, a slightly absorbent paper). They typically offer three- or four-color printing at much finer resolution than any other printers. Some use seven inks (high and low intensities of the three subtractive primaries) to give more realistic results. Though the tendency of the ink to spread limits their ultimate performance, it also helps improve the perceived realism of the image by making individual pixels less visible. Printing to photographic quality requires special paper and also uses large amounts of the expensive ink, so it is not cheap, but the results bear comparison with those from expensive dye sublimation printers. Because (unlike other printer technologies) there is no inherent limitation on the size of the paper, A3 and larger printers are readily available for such tasks as printing conference posters. Unlike dye sublimation printers there is always the option with an inkjet of printing at lower quality and hugely lower cost for proofing and layout purposes.

Dye Sublimation Printers

These printers have changed little since they the previous edition of this book. They are still an excellent output medium for routine production of photographic quality output, and although still not cheap, the price has not increased in line with inflation so they are not much more expensive per page than the inkjet for photo-quality output, though the purchase price is much greater. They use a full-page sheet of ink for each color they print, but in this case the ink sublimes when heated and is absorbed by a specially coated sheet of paper. The vaporized ink will tend to diffuse laterally to a limited extent, making pixelation less obvious. More or less ink can be transferred, depending upon the amount of heat applied, so that true gray scales can be produced. The output from the best dye sublimation printers can be comparable to photography, but the cost is also similar — up to US$5.00 for an A4 size color print. The cost per page is fixed, unlike a laser or inkjet printer, where the cost per page depends upon the degree of coverage.

CONCLUSION

The big change since the previous edition of this book 10 years ago is that now mass-market media are effectively equal to the demands of the confocal user. It is a truism that for several years now developments in the personal computer market have been driven not by business or scientific usage, but by the domestic market. The requirements of games, music, digital photography, and video have been the driving factors for processor speed, interfaces, display quality, print output, and data storage. We no longer need specialist image manipulation hardware, custom video cards, non-standard monitors, dedicated data buses, or expensive storage devices. Computers are fast enough for the necessary image manipulation, everyday video cards handle 24-bit images at high resolutions and fast refresh rates, as well as providing hard-wired image manipulation functions. Monitors offer megapixel displays at high bit depths and refresh rates. Fire Wire and USB-2 will carry data faster than any point-scanning confocal can scan, and domestic video disks are big enough to handle huge data sets. Just about every home has a printer giving photographic quality output — and these are printing higher resolution images than most microscope users generate.

The major computer magazines generally run annual surveys of color printers and it is always worth seeking out the latest of these before making a purchase. A final word of warning: If you are evaluating a hard copy system, of whatever sort, insist on testing it on real confocal images from your own work. Every manufacturer has a gallery of images which reproduce superbly on his own hardware, and if you try to judge a system on the basis of such pictures you will be disappointed once you start using it yourself.

SUMMARY

Bulk Storage

Image compression has made huge strides but still, as always, needs to be used with care. The new wavelet compression system seems to offer little to the confocal user but the PNG format has at last given us a lossless technique that works. Recordable CDs are currently the most common and most reliable choice for mass storage, though recordable DVDs seem likely to take over during the lifetime of this edition. For archival purposes, there seems little merit in selecting the rewritable version of either of these media. Tiny solid-state FLASH memory devices have become a very effective way to carry presentations from place to place.

Display

Monitors are no longer a problem; 24-bit displays of adequate size and refresh rate, as well as the display boards to drive them, are now the norm rather than expensive exceptions. Flat-screen LCD displays are still expensive, but worth the cost for their extra sharpness and for the complete elimination of flicker.

Hard Copy

Inkjet printers have now essentially replaced dye sublimation printers for optimal photographic output. They have much lower purchase cost but the cost per glossy print is still high. Color laser printers are now more than adequate for proofing and preprints. Finally, if you are evaluating a hard copy system, insist on testing it on real confocal images from your own work.

ACKNOWLEDGMENTS

I am very grateful to Mike Speak and Dennis Dwarte, who have been generous with their time and effort in helping me prepare this chapter. I also owe a debt of gratitude to Todd Brelje, who wrote the original outline but who was unable to proceed with the chapter due to other commitments. In spite of this, he gave much assistance in producing the original edition, for which Martin Wessendorf, Nick White, Andrew Dixon, and Jim Pawley also provided a tremendous amount of information and help. The revision for the 3rd edition was carried out while I was at the Instituto Gulbenkian de Ciência, Oeiras, Portugal, and I thank them for their hospitality; Nuno Moreno and José Feíjo provided much help and assistance, and critical reading of the manuscript.

REFERENCES

Anson, L.F., 1993, Fractal image compression, *Byte* 19(11):195–202.

Avinash, G.B., 1993, Image compression and data integrity in confocal microscopy, *Proc. Microsc. Soc. Am.*, Jones and Begell, New York, 51: 206–207.

Barnsley, M.F., and Hurd, L.P., 1993, *Fractal Image Compression*, A.K. Peters, Wellsley, Massachusetts.

Compact Flash Association, 2003, *The Compact Flash Standard*, available from: http://www.compactflash.org/.

Cox, G.C., 1987, What went wrong with my micrograph? *Aust. EM Newslett*. 16:12–13.

Cox, G.C., 1993, Trends in confocal microscopy, *Micron* 24:237–247.

Cox, G.C., and Sheppard, C., 1993, Effects of image deconvolution on optical sectioning in conventional and confocal microscopes, *Bioimaging* 1:82–95.

Cox, C.G., and Sheppard, C., 1999, Appropriate image processing for confocal microscopy, In: *Focus on Multidimensional Microscopy* (P.C. Cheng, P.P. Hwang, J.L. Wu, G. Wang, and H, Kim, eds.), World Scientific Publishing, Singapore, pp. 42–54.

Deutsch, P.L., 1996, *DEFLATE Compressed Data Format Specification version 1.3*, available from: ftp://ftp.uu.net/graphics/png/documents/zlib/zdoc-index.html.

Digital Photography Review, 2003, *Digital Film Comparison*, available from: http://www.dpreview.com/articles/mediacompare/.

Mortimer, F.J., and Sowerby, A.L.M., 1941, *Wall's Dictionary of Photography, Sixteenth Edition*, Iliffe & Sons, Ltd., London.

Nathans, S.F., ed., 2003, *Writable DVD Drives*, available from: http://www.emedialive.com/r18/.

Nugent, W.R., 1989, *Estimating the Permanence of Optical Disks by Accelerated Aging and Arrhenius Extrapolation*, ANSI X3B11/89–101.

Pennebaker, W.B., and Mitchell, J., 1993, *JPEG Still Image Compression Standard*, Van Nostrand Rheinhold, New York.

Powell, E., 2004, *The Great Technology War: LCD vs. DLP*, available from: http://www.projectorcentral.com/lcd_dlp_update.htm.

Redfern, A., 1989, A discrete transformation, *Australian Personal Computer* 10(12):144–154.

Roelofs, G., 2003, *Portable Network Graphics*, available from: http://www.libpng.org/pub/png/.

Stinson, D, Ameli, F., and Zaino, N., 1995, *Lifetime of KODAK Writable CD and Photo CD Media*, available from: http://www.cd-info.com/CDIC/Technology/CD-R/Media/Kodak.html.