

1

What is DTP?

Long ago (it seems), when people wanted to put together a brochure, leaflet, poster or newsletter, they were quite content to bash away at a typewriter, use felt tip pens and sometimes, for that really special look, rub down instant lettering. How happy we all were in our innocence. These days everyone expects so much more, in finished quality, in speed and in convenience. The factor that has raised our sights above our former naïveté is the relatively recent blossoming of desktop publishing, commonly called DTP.

DTP is the latest in a series of technological advances in printing: as Gutenberg's invention (and Caxton's bringing to England) of movable type changed the way that books were produced and made them accessible to a wider public for the first time, so DTP has brought a substantial part of the process within the reach of everyone. With DTP you have almost complete control over every element of the page: the size and style of lettering, and the size and position of all sorts of pictures from all kinds of sources, can be changed in a moment.

Such power is not without its responsibilities, however. DTP relies on the careful dovetailing of a number of separate skills and is not something that you can expect to get the hang of straight away. Even the manuals can only tell you what to do to achieve certain effects; they won't explain what you ought to be trying to do. Don't be discouraged, though: nothing worth accomplishing comes without effort.

The other side of the coin is the enjoyment of spending time on a project and achieving a printed product that not only pleases you but impresses your family, your friends – and your clients. Make no mistake about it, DTP is more than just an absorbing way to pass the time (though it is all of that): it is creative and fun, and can be profitable.

From Manuscript to Type

The Computing Perspective

To put DTP into perspective we need to look at what it is and what it replaces. On the one hand we are all familiar with typewriters and what they can do (in my case, not a great deal), but DTP is much more than a replacement for the mere mechanics of typing. It involves a whole new way of looking at committing the written word to print. The printed word is a combination of many factors, from the purely functional to the aesthetic, and by the end of this book you should have a fair idea of what they are and how to put them to work for you.

Let us start with the simplest of these computer aids: the word processor. No doubt you will already be aware of the basic idea even if you haven't had practical experience with one. The original concept, which came from the office environment, was to liberate the typist from the chore of using correcting fluid on errors and the drudgery of completely retyping the whole of a document every time part of it was rewritten.

Word processors reduced the work involved in making minor and major changes to texts, and did so to such an extent that documents can now go through many revisions, each version being printed out for checking, for often quite trivial reasons. A cautionary point I shall return to every so often is that just because the technology makes something possible, it doesn't necessarily make it the best approach. That is not to say that we shouldn't use the facilities on offer, merely that we should think about what we are trying to achieve before doing so.

Word processors developed from simple tools for typing on screen, correcting mistakes and moving bits of text around, and became quite sophisticated. For example, they allowed 'mail merging', the production of personalised letters from a master copy and a list of names and addresses. The text could be automatically split up into pages with footnotes and page numbers inserted at the correct places, and different styles and sizes of type could be sprinkled about (as long as your printer could cope with them). Spelling could be checked during or after typing, and some software will even attempt to improve your writing style (as this is not yet available for the Archimedes, I have had to manage without).

Until recently, though, if you wanted to include diagrams and pictures in what would otherwise be a solid mass of text, word processors couldn't help you. You were forced to print the text first and then add the pictures with scissors and gluepot. If you also fancied having the text in two or three columns on the page then you might have had to cut the text into strips and assemble the whole thing by hand. The term

'paste-up' referred as much to what happened to your fingers as what went on the page: it was a potentially messy, frustrating and time-consuming job that was hard to get looking right. It was best left to experts, as the risk of transposing or losing entire passages was very great.

It was therefore natural for software designers to look for ways round this and try to achieve printed output that didn't just mimic a typewriter. They recognised that businesses would want to smarten up their correspondence and enliven their reports by trying to emulate traditional printing. The advantages were predicted to be twofold: all the processes would be under the direct control of the company and would therefore be faster, and there would be cost and time savings through not needing to farm out typesetting to outside firms.

What happened in practice was that the separate jobs that would normally have been done by skilled designers and typesetters tended at first to be pushed down the hierarchy to the office secretary, because 'the software would take care of the technical stuff'. This short-sighted attitude resulted in a proliferation of ill-designed and poorly presented documents, which clearly demonstrated that DTP software was not the whole answer. In the world of fair-sized businesses there is now a well-defined niche for DTP operators who have had training in a variety of skills.

In the home, school and small business environment the luxury of such job demarcation is either irrelevant or impracticable, so let us move away from the historical perspective on the computing side and examine the traditional procedures that DTP replaces, and also the skills that it doesn't.

The Printing Perspective

In traditional publishing terms, all that DTP does is to provide a different way of preparing a finished single copy of a document ready for duplication. This was once the exclusive domain of printing houses (strictly, the preparation of camera copy is done by *typesetters*), and some mystique surrounded the printer's skills. In the book world this was (and still is) interwoven with the complex set of processes performed by the publisher, starting with the author's written word and culminating in printed copies being distributed.

Briefly, the publisher's first job was to make sure that the author's spelling and grammar were correct and that any illustrations were of high enough quality. Crucially, though, before passing the typescript to the printer the publisher would draw up a design specification for the publication that would be aesthetically pleasing, appropriate for the material and also suitable for the expected type of reader. For instance, student text books can be very densely packed and have narrowish margins (because that makes them cheaper, and students can't afford to

pay for blank paper or too many pages), whereas ‘coffee table’ books often have lots of white areas to complement the overall appearance of the text and pictures. All the different type faces and sizes were carefully chosen to give a balanced appearance.

The printer’s job was effectively to retype the whole text on a sophisticated system which resulted either in a back-to-front image of the words cast from molten lead (there were several such hot metal processes) or in a photographic image of the words on film (phototypesetting). The printer would provide rough prints (proofs) for the author and publisher to check, even though any corrections that had to be made in either process were expensive.

At some stage the typeset text needed to be integrated with any illustrations as specified in the design, and an index might have to be prepared (the author’s responsibility!). Once the publisher was satisfied that all was well, the printer would print off the copies required.

Thus we can identify six main areas in producing a publication: writing, design, typesetting, page layout, preparing a final perfect copy and duplicating. All of these require skill and experience. What is true of books is true of everything else that DTP can be used for, from posters to pamphlets, from messages to magazines. DTP can provide only the tools for typesetting, page layout and preparation of the printed image; the rest is up to you.

But there is one more phase, which is implicit in every stage: checking. Taking the traditional route meant that a good few people other than the author saw the job as it progressed and could correct those simple but embarrassing errors that we all make. With DTP you have no one to blame but yourself!

Screen versus Paper

Word processors were originally designed to handle only *characters*, i.e. letters and numbers, whereas DTP deals with everything, both text and pictures, as images, or *graphics*. This lets you make as many changes (and mistakes) as you like to the balance of all the elements on the page before committing the result to paper. The term WYSIWYG (What You See Is What You Get) was coined to express this philosophy; as you might expect, it is only partly accurate. Quite clearly a laser printer that can distinguish 300 dots to the inch will give a different result from a monitor screen that can resolve only perhaps 80 dots to the inch horizontally. Some sort of approximation is called for on screen.

Figure 1.1 shows how this is done for text on the Archimedes: the same letter is portrayed in four guises. In (a) we see an interpretation of how it is stored on disk, as an outline form made up of points joined by straight lines and curves. The outline is

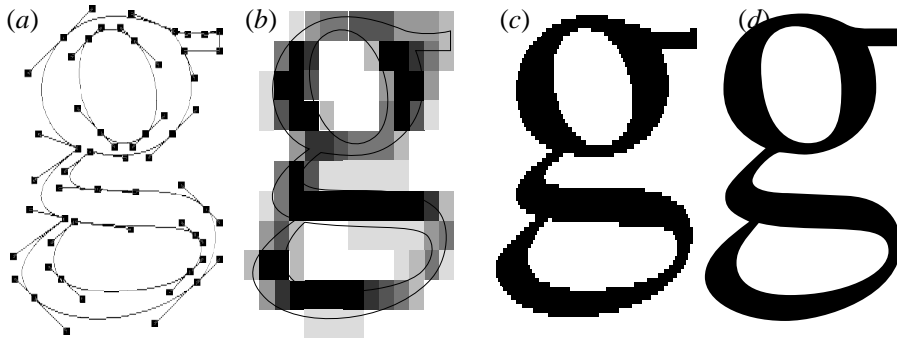


Figure 1.1 Different representations of a letter: (a) outline; (b) small on screen, much enlarged; (c) larger on screen; (d) as printed out

a mathematically smooth image that can be converted to give an appropriate set of rectangular dots, or *pixels*, at any letter size.

The problem with small letter sizes on screen is that they tend to become unreadable because the number of pixels is small, distorting the letter shape. The Archimedes avoids this by a technique called *anti-aliasing*, which makes text surprisingly readable even at quite small sizes. Each pixel is given a grey level depending on how much of the outline passes through it; (b) shows a letter that was small on screen, much enlarged and overlaid with the outline. Half-close your eyes or hold the page a long way away and you will see a fairly smooth shape, if a little blurred. Try constructing the same letter out of black rectangles in a grid only 14 wide by 11 high and you will see just how effective anti-aliasing is.

At larger letter sizes on screen anti-aliasing is unnecessary for legibility, although some jaggedness remains, as shown in (c). Lastly, (d) shows a laser printout, illustrating the difference in resolution between screen and printer. Anti-aliasing can't be used in printing, because ink tends to come only in black, not shades of grey.

Similarly, these differences in resolution will lead to differences in appearance between any pictures on screen and in the final printout. Despite these discrepancies, using DTP gives you a very good picture of the finished product because you can zoom in on the page for close-up work and then look at the whole page on screen to get an overall impression, even if details cannot be seen clearly.

Some Printing Terminology

Many of the terms used in DTP have their roots in the history of typesetting and printing technology, so I make no apology for including a section on them here.

They are fundamental to any description of the processes involved, although purists might object to subtle mutations in meaning that have happened during the evolution from traditional to computer-based technology.

Measurements

Points and picas. The fundamental unit of measurement in printing is the *point* (usually abbreviated to pt), which is roughly equal to 1/72 inch, and individual letters are often described as having a certain *point size*. One point sounds a very small amount, but the eye is surprisingly sensitive to fine differences in letter size: even a 1pt variation is very noticeable up to about 16pt, and the trained eye can pick out deviations of less than a point.

Figure 1.2 gives you an idea of how point sizes translate into appearance, ranging from tiny (now rare, but it was sometimes used in bibles) to larger than is sensible in a book. Tabloid newspapers use still larger sizes to grab attention from a distance.

This is 4pt type
This is 6pt type
This is 8pt type
This is 10pt type
This is 12pt type
This is 14 pt type
This is 16pt type
This is 20pt type
This is 24pt type
This is 28pt type
This is 36pt type
This is 48pt type
This is 60pt

Figure 1.2 Text at a range of sizes

Having 72 units to the inch doesn't make for convenient figures for measuring larger things such as text widths. An alternative, still fairly common, is to use the *pica* (it rhymes with 'biker'), which is equal to 12 points; there are 6 picas to the inch. A measurement of, say, 5 picas and 6 points could be abbreviated to 5pi 6pt or 5,6 (note the comma) or 5½pi. It is worth noting that the standard line spacing on a dot-matrix or daisywheel printer is one-sixth of an inch, or one pica.

The terminology of type sizes goes back to the days when typesetting, or *setting*, was still done on movable pieces of lead type. Figure 1.3 shows names for various parts of a letter and also certain defined heights, which to some extent determine the characteristic 'look' of a type face. The size of the piece of type was the same as, or very slightly more than, the distance from the top of the highest ascender to the lowest descender. On printing, this could give a tiny white space between one line and the next. The amount of this white gap could be increased by putting thin strips of lead between the lines of text, which gave rise to the term *leading* (pronounced 'ledding') to describe this situation. So the text of this book, which is 10pt with 2pt leading, would be described, somewhat colloquially, as '10 plus 2'. The reason for including leading is to make the text more readable: the longer the line, the more likely the eye is to get confused, so the more leading is needed.

Another way of looking at leading is to measure the distance between successive *baselines* (see Figure 1.3) and call it *line spacing*, in which case this text comes out as 10pt type with 12pt line spacing, or '10 on 12'. But beware: the term 'leading' is sometimes used indiscriminately as though it meant line spacing; this is wrong, confusing and to be avoided.

Text that was typeset without any leading was described as *solid*, and if leading had been added it was naturally called *leaded*. It was normally impossible to set lines in lead type closer together than their point size, but with DTP there is no such restriction because the lines of text are generated purely electronically rather than




Figure 1.3 Terms used in measuring and describing text

being made of metal. Thus *negative leading* can now be used for special effects. Rather than say '12 minus 2', though, it makes more sense to speak of '12 on 10'.

Ems and ens. Points and picas are *absolute* measures, as values expressed in them are constant. On the other hand it is sometimes useful to describe a distance by a *relative* measure, i.e. one that is related to the type size, in the way that a word takes up more space in a larger point size even though the letters are the same. For this we use the *em*, which is easier to visualise than it is to put into words.

An em is simply a horizontal distance equal to the current point size of the text. It is named after (but is no longer necessarily anything to do with) the space occupied by a lower-case m. A 14pt em space is therefore 14pt wide, and similarly a 10pt em space is 10pt wide. One of the invariable conventions of printing used to be to start a new paragraph by inseting the text by one or two ems; other, more modern, conventions also exist, as this book demonstrates. But the beauty of using ems for such a specification is that the exact amount of space used is independent of the point size that the text is eventually set in: it stays in proportion visually.

A simple subdivision of the em is the *en*, being half an em in width and formerly related to the width of the letter n.

Fonts and Faces

A book, newspaper, newsletter, brochure or poster (or, for that matter, just about any printed object) set in only one sort of type is likely to look pretty boring, even if different sizes of type are used. That is why the chapter headings, subheadings and figure captions in this book are set in a variety of different-looking types, to catch the eye. Each different form is called a *font* (or *fount*; also *typeface* or *face*). A full font consists of all the letters and numbers, punctuation marks and special symbols. (I realise that I am perpetuating an incorrect use of strict terminology by using 'font' as a synonym for 'typeface', but this is a case where current usage has overtaken definition.)

Despite their differences, all the fonts used in this book have a kind of resemblance to one another. This is because they all belong to the same type *family*, in this case designed to match the typeface Times (but Acorn have called it Trinity for reasons of copyright). A family usually consists of at least four faces: roman (the normal upright version), italic or oblique (a sloping version), bold, and bold italic (or bold oblique). The difference between italic and oblique is that italic is a completely separate design, with many of the letters showing different characteristics, whereas oblique is more like a sloped variant of roman.

Trinity Medium (roman)
Trinity Medium Italic
Trinity Bold
Trinity Bold Italic
Trinity Medium, sloped
 Trinity Medium Italic, back-sloped

Figure 1.4 The Trinity family of fonts, and two extra faces created by electronic sleight of hand

Corpus (Courier)	↑*□●⊙⊗*⊕⇒▼ (Dingbats)
NewCorn (New Baskerville)	Fancy Dress
Homerton (Helvetica)	Γρεεκ (Greek)
Pembroke (Palatino)	Katie (Friz Quadrata)
Trinity (Times)	<i>Sophie (Park Avenue)</i>
Vogue (Avant Garde)	<i>Tamsin (Brush Script)</i>
<i>Chance (ZapfChancery)</i>	Warsaw (Tiffany)

Figure 1.5 Some standard DTP fonts (left) and a few more exotic and specialised ones (right), with conventional names in parentheses; all at 18pt

In Figure 1.4 we see all four Trinity faces, as well as a roman sloped by 15° forwards to show the difference between italic and oblique, and also an italic back-slanted to make it more or less upright for comparison with roman.

Times (or Trinity) is a fairly plain, straightforward face with a neutral appearance. Many other fonts are now available for the Archimedes if you want to give your message a specific flavour (warm, exciting, frivolous, formal, informal or whatever) or if you have special needs for symbols. A selection of fonts is shown in Figure 1.5, together with their standard names in parentheses.

All the DTP packages discussed in this book come with several fonts. Be warned that some fonts for the Archimedes have the same appearance, but different names,

because they were designed for different software suppliers (Pembroke/Paladin and AvantG/Vogue are notable look-alike pairs).

Fonts can be divided into two fundamental camps: *serifed* and *sans serif*, besides other categories like script, ornamental and possibly even outrageous. The main distinction is based on whether or not the letters have decorative ‘ticks’ on the ends of the strokes. This is easier to demonstrate than to put into words, so here are some letters in 18pt Trinity (a serifed font) and in 18pt Homerton (sans serif):

HETbiv HETbiv

Serifed fonts have been around in various designs since about 1470, which tends to give text set in them a ‘traditional’ air. The first sans serif fonts, which were designed in the 19th century, looked so odd to people used to seeing serifs that they were described as Grotesque. Text set in sans serif type generally has a more ‘modern’ appearance.

You might imagine that an 18pt letter in Trinity would be the same height as the same letter in Homerton or Vogue. A glance at Figure 1.5 and the example above will show you that they aren’t: both capital and lower case letters can differ in height from their equivalents in other fonts (compare Trinity and Vogue in particular). This isn’t a mistake, it is a design feature; however, it does make mixing fonts of different families more complicated. Fortunately, fonts within a family are designed to work together and are therefore compatible in size.

The scale differences, as well as the presence or absence of serifs, is partly what determines the overall look of a face: if the *x-height* (see Figure 1.3) is larger the ascenders and descenders may be proportionately shorter, although in some older fonts the ascenders rise above the capital height to give a distinct elegance as well as more open space between lines. A book set wholly in sans serif like Vogue or Homerton will be less easily readable than one set in Trinity because there is less shape distinction between the letters.

Paper and Pages

The commonest size of paper that you will find in offices these days is undoubtedly A4, measuring 210mm × 297mm. This is the standard size for photocopiers and office stationery. In fact, the A in A4 represents a whole series of paper sizes, each exactly half the size of the one above, starting at A0 (the largest, 841mm × 1189mm, with an area of one square metre). You can picture the series as resulting from successive foldings along the longest side each time, as Figure 1.6 indicates.




Figure 1.6 A scale diagram of the A series of paper sizes

This is fine when it comes to separate sheets of paper, but in computing we often use continuous stationery. The commonest page size here is wider and shorter than A4 (8½ inches × 11 inches), although ‘true’ A4 is available. The important point is that whatever paper size you use, you must remember to set the page size in your DTP program or word processor accordingly.

The normal orientation of a page is like this one, with the longer side vertical. This is known as *portrait* orientation. The other way round, where the width is greater than the height, is called *landscape*. We can distinguish between the orientation of the printing on a page (which is the same as the shape of the page on the screen: text is always set horizontally) and the rotation needed to print the page out on a printer that may have a maximum print width of about 8 inches. It is therefore correct to describe a page on screen as landscape but to think in terms of printing it sideways.

If you have done any typing at all you will be familiar with *margins*, the blank areas

either side of the typed area. With DTP, because each page is treated as an image, there are four margins to consider: top, bottom, left and right. That no doubt seems obvious, but as with the portrait/landscape distinction, the limitations of your printer also come into the picture because there may be a region around the edges of the page where it is physically impossible for your printer to print. This is the *print border*, and most DTP software gives you the chance to see it on screen so that you can keep your page inside it.

Take a look at almost any book and you will see that the left and right pages are arranged differently. There are many ways in which this can manifest itself, but for now it will be enough to note that the margins on any page are often asymmetric, those at the outer edges of the pages being wider or narrower than the margins towards the spine, and the top (or 'head') margin being greater than that at the foot of the page. The overall effect is to make the left page a mirror image of the right, resulting in an attractive balance.

Types of Illustration

In any kind of printing work there are two fundamental types of illustration, by which I mean any non-textual material. The first, and simplest, is *line art*, which is anything that can be printed as solid lines or areas of the ink colour (normally black). Figure 1.6 is an example of line art.

The second sort is the *half-tone*, which is rather more complicated. Consider a normal black-and-white photograph: some areas may be completely black and others completely white, but mostly there will be areas that are lighter or darker shades of grey. This situation is called *full-tone*, because the graduation from white to black is smooth and effectively continuous.

It is impossible to print an accurate representation of a full-tone picture as an example in this book because, as mentioned earlier, printing ink comes (for our purposes) in black only: the ink is either there or it isn't. Figure 1.7(a) shows the unedifying mess that results from trying to print a photograph without artificial aid.

The way round the problem is to divide the original into very small areas and measure the average greyness in each, then convert every measurement into a tiny dot that covers only a fraction of the white paper around it. The outcome is that each sampled area becomes represented by a partial covering of ink that averages to approximately the amount of grey in the photograph in that area. This used to be done photographically with expensive specialised equipment; Figure 1.7(b) demonstrates what a half-tone looks like when the conversion is made by computer software, using only 16 different grey levels (more would have been better). You can see the different-sized dots if you look closely.






Figure 1.7 (a) An attempt to print a full-tone picture; (b) after conversion to half-tone

Related to, but rather simpler than, the half-tone is the *tint*, which is an area filled with uniform fine dots giving a grey appearance. This is often used for clarifying or defining an area on a piece of line art (see Figure 1.3 for an example) or, in DTP especially, as a panel behind text to make it stand out on the page. On screen, the tint areas show up as solid grey or as a colour; the dots are produced during the printing process.

