

David Pilling's
Scanning
Software

David Pilling Software

© Copyright David Pilling 2005.

All Rights Reserved.

No part of this product may be reproduced in whole or part by any means without the written permission of the publisher. Unauthorised hiring, renting, loaning, public performance or broadcasting of this product or its constituent parts is prohibited.

This manual was produced with Ovation Pro DTP.

(Windows version 6th. April 2006)

If you have any comments, suggestions, bug reports or complaints please contact me. If you would like further copies of this program or details of the many other pieces of software available please visit my web site.

David Pilling
P.O. Box 22,
Thornton Cleveleys,
Blackpool.
FY5 1LR.
United Kingdom.

<http://www.davidpilling.net>

david@pilling.demon.co.uk

From time to time enhanced versions of programs will appear. You can upgrade your copy to the latest one from the above web site.

Contents

1. Introduction	1
2. Installation	3
3. Scanning something	4
4. Basic operation	5
5. Edit	13
6. Filter	16
7. Colour	20
8. Colour Management	31
9. Configuration	37
10. File formats	39
11. Snapper	55
12. Credits	57

1. Introduction

This is a program for manipulating bitmap images. Users of Windows will already be familiar many such programs.

The special features of DPScan are;

- It can work with 1, 2, 4, 8, 16, 24 and 32 bits per pixel images directly. Often other programs require images to be changed to another format before they can be edited.
- Various RISC OS file types are supported.
- It is also available for RISC OS and so will be familiar to RISC OS users.

Fundamentals

Images, or pictures come in two basic formats - bitmap or vector. A bitmap image consists of a huge number of coloured points. Usually these points will be referred to as pixels or dots. The dots are arranged in a rectangular array, the size of an image can then be specified by the width in pixels and the height in lines or rows of pixels. The principle is the same as TV pictures.

Vector images consist of lines and curves which specify filled areas. For Windows the typical bitmap image is a Windows BMP file, and the typical vector image a Windows Meta File (WMF). Mostly this program is concerned with bitmaps.

Given a bitmap image consisting of lots of pixels. Two fundamental numbers control how it will look. First the 'resolution' this is how close together the dots are. It should be obvious that the finer the dots, the better the image will look. To specify resolution, we talk about 'dots per inch' or 'dpi'. The resolution may be different across the image to down it. Therefore there is the X or horizontal (left to right) resolution and the Y or vertical (top to bottom) resolution.

Second how many colours each dot may be displayed in. This is controlled by the number of bits used to store each pixel. If one bit is used for each pixel, then the pixel can be in one of just two states, on or off. If 2 bits are used for each pixel, it can be one of four colours. Four bits per pixel give 16 colours and 8 bits 256 colours.

No mention has been made of how the states of these bits relate to the colour seen. For example in the one bit per pixel case, the on state might mean a black pixel or a white one with off meaning white or black. Conventionally a table is provided with the image that links the state of the bits with the real colours displayed on screen. This table is called the 'palette'.

We shorten bits per pixel to bpp or sometimes 'bits' in future. For more than

Introduction

8bpp, the palette would be too big, so for images with a larger number of bits, a different method is used, with the bits representing the colours directly. For a 16bpp image 5 bits are allocated to each of the colours red, green and blue (RGB), for 24 and 32 bits, 8 bits are allocated.

As expected the more bits per pixel, the better images will look. However they will use more memory.

Most Windows users have their screens configured to 96dpi horizontally, and vertically. In Windows terminology this corresponds to “small fonts”. The usual alternative is “large fonts” or 120dpi. More recent versions of Windows allow any screen resolution to be used.

Normally the Windows screen uses 24bpp, older computers may use 15 or 16bpp or even less. Dot matrix printer or Fax machine output is typically 200 dpi, laser printer or ink-jet output is 300 or 600 dpi. Books are typically printed at 2400 dpi.

It is possible for some pixels to be transparent. If pixels are completely transparent or completely opaque, it is common to say the bitmap image has a mask, or is masked. If the transparency is variable the image is said to have an ‘alpha’ channel.

Various methods can be used to support transparency, one way is to store another array of values which hold the transparency value for each pixel. Another technique is to specify that any pixel of a certain colour is transparent.

Images may be defined in various colour spaces. The common one is RGB (red, green, blue) used by all computer screens. Another is CMYK (Cyan Magenta Yellow black) used in printed material.

It is common to see scanners advertised as “1200 dpi, 24bit colour”. The inclination is therefore to scan everything at the maximum resolution and in as many colours as possible. However consider scanning a page of A4 at this resolution. The resulting data would occupy some $(1200 \times 1200 \times 10 \times 12 \times 24) / 8$ bytes, about 480Mbytes. This is larger than the memory on many computers, and would make the image difficult to handle. The lesson here is to use resolution sparingly. Often using lower resolution will not give visibly worse results.

There is usually nothing to be gained by scanning at a resolution higher than the “optical” resolution of the scanner, this is the actual resolution of the sensor and mechanism of the scanner. Most scanners offer to produce higher resolution images but do this by inventing extra data.

2. Installation

If you are using the program with a scanner, you will need to install TWAIN on your computer. You will also have to have some sort of physical connection between the scanner and the computer. Full details of this will be given in the documentation with your scanner.

Run the DPScan installer by double clicking on it. The installer is called "dpscani.exe".

The installer will prompt you for various pieces of information, the process should be familiar from other Window programs.

Once installation is complete you can use the Windows Start menu to run DPScan.

An empty DPScan window will appear.

To the left of the main window a small window will open showing tools. Below the main window there will be another window which shows information about the current image. These are referred to as the 'toolbox' and 'info palette'. Different tools can be chosen by clicking on them. When a tool is chosen the info palette display will change.

The main window is used to display images being manipulated, to set the area of pages to be scanned, and the position on the page of images being printed.

The caret can be positioned in either the main window, or the info palette by clicking on the main window or in one of the writable icons in the info palette with the left mouse button.

The Help menu (right hand side of the main window menu bar) has two entries, the first displays a help file which contains general information about the program. The second entry displays an information window. This will show you which version of the program you have. Please quote this in any correspondence.

You can quit the program by closing all open windows or by using Exit on the File menu.

3. Scanning something

This is the process of getting an image from paper into your computer. We can only deal with the computer side of the process here. You must have connected the scanner to your computer, usually via the SCSI bus, USB port, or possibly the printer port. The object you want to scan must have been inserted into the scanner, and the scanner switched on and made ready. Full details of this process should come in the documentation for the scanner.

The program controls scanners via a software interface called TWAIN. Before trying to scan anything, you must have installed TWAIN.

You must have a suitable TWAIN driver for the scanner you want to use.

Having completed the “obvious” steps above, this is a typical scanning session.

First load DPScan from the start menu.

Click on File on the menu bar. Choose 'Select Source...' This will open the TWAIN manager source selection dialogue box. Click on the name of the data source for your scanner, and then OK.

Now repeat the process but this time, choose 'Acquire...' from the file menu. The window for the source for your scanner will open.

Follow the instructions supplied with your scanner, eventually you will end up with an image in DPScan.

To save your image again use the File menu.

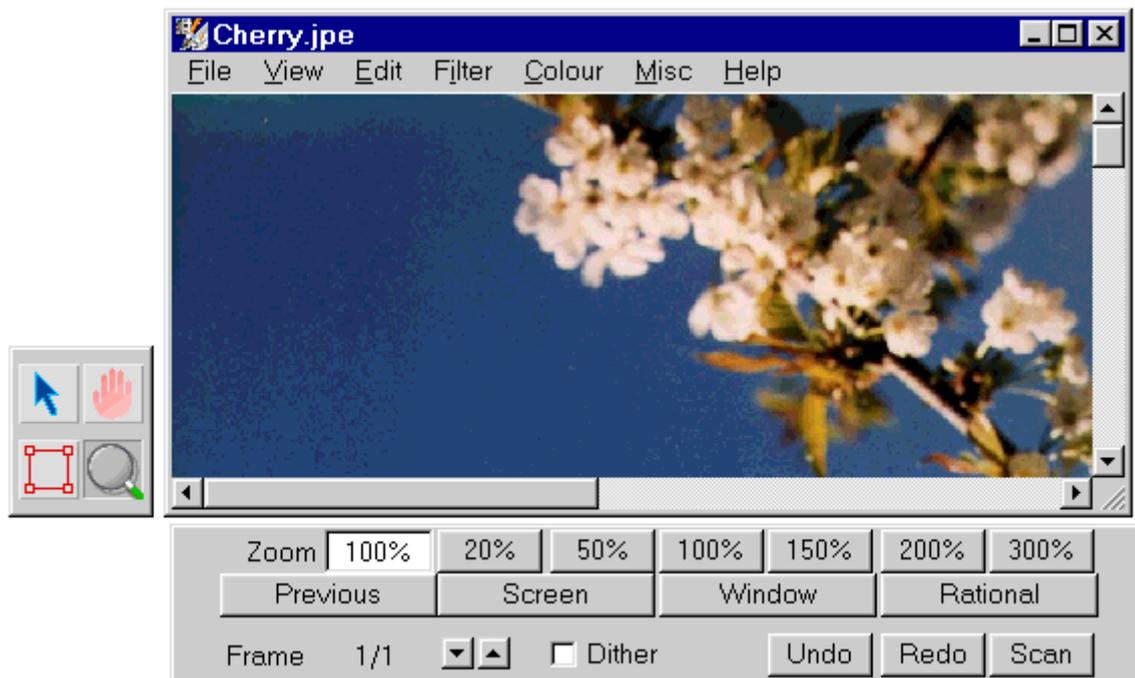
There are some keyboard short cuts to make the above process simpler. Shift F10 will open the TWAIN source selector, and if you only have one source installed, it is not necessary to go through this step. F10 will open the source window i.e. the same as Acquire...

Much of the rest of this manual is concerned with the things the program lets you do with images once you have scanned them, and various ways of making the above process easier or more accurate.

4. Basic operation

If you run DPScan from the Start menu you will get an empty window. In most bitmap programs there is no such thing as an empty window, however because the program is used for acquiring images, it is quite a sensible thing. Another difference to packages like Paint, is that the window is in general not exactly the same size as the image you have loaded. In this program the window represents a work area. This can be the area of your scanner, or the area of your printer. You can move images around the work area to control where they are printed on paper. Similarly by marking an area of the work space you can select which bits of a page your scanner will scan.

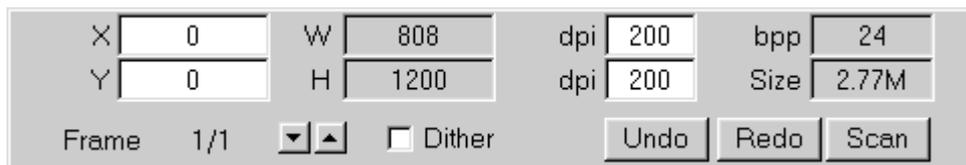
Usually when the main window opens, two smaller windows will open. The one to the left is the 'toolbox' the one underneath is the 'info palette'. By clicking on the toolbox icons, you can choose the current tool. The info palette displays useful information about the current image, and also allows you to control the image by either typing in values or clicking on the icons. The info palette display will change depending on the current tool. The tools are the pointer, the grabber hand, the rectangular selection and the magnifier. These will now be described.



The Pointer tool



This allows you to move the current image around the window. You do this by clicking within the image and dragging with the left button. Holding down the Shift key and clicking on an image will transfer the colour of the pixel clicked on to the palette window (see chapter 7). Holding down Control and dragging will produce a line. Attempt to make this lie along a line in the image which should be vertical or horizontal. When you release the mouse button the Edit menu rotate window will open with the correct angle to make the line in the image horizontal or vertical.



This is the info palette that accompanies the pointer tool. X and Y are the coordinates of the top left hand corner of the image, relative to the top left hand corner of the page displayed. Positive Y is downward - as if you measured the distance down the page with a ruler. W and H are the width and the height of the image. The measurements will be in the current units (see later) (pixels are shown above). The two dpi values are the resolution in the X and Y directions. bpp is the number of bits per pixel for the image. Size is the amount of memory the image uses - K signifies Kilo bytes (bytes times 1024 or just over 1000), M signifies Mega bytes (1024*1024).

Frame is the current frame number in the file, followed by the total number of frames in it. Some of the file formats the program supports, can hold more than one image in a file - you'll be familiar with this, if you've ever loaded a Sprite file into Paint. In general the large images that the program is designed for are best dealt with one image to a file. The program can read files with more than one image, but it can only save one image to a file. We term each image in a file a 'frame' or sometimes 'page'. By clicking on the arrows at the side of Frame you can step back and forwards through the images in the file.

Dither. This controls if the program will attempt to use dithering to give a better on screen display. For example if you are working with 24bpp images and only have a 8bpp screen, enabling dither will make it much easier to see what you are doing. Dither can be toggled by clicking on the small icon to the left.

The info palette can be moved by dragging inside it - like on a normal window title bar. The toolbox works in the same way.

Some of the icons in the info palette are writable. By left clicking you can position the caret in them and then type in values.

The writable icons include X, Y which allow more precise values to be entered. Finally the dpi values can be entered. This may seem odd. All it does is allow you to change the way the image is displayed. It can be useful if as often happens the image is supplied with incorrect resolution values. Values typed in will take effect when Return is typed.

The Grabber Hand

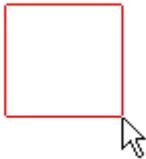


This is a simple tool. It just allows you to move around the window, without using the scroll bars, you do this by clicking inside the window and then dragging. The info palette display is the same as for the pointer tool.



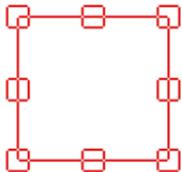
The Rectangular Selection tool

This is probably the most useful tool. You use it to mark off a rectangular section of the work area. This area can be an area of an image that you want to apply the image processing tools to, or it can be an area of the page that you want to be scanned.



To mark off an area, click inside the window and then drag with the left button.

As soon as you release the mouse button the area will display 8 small squares or 'ears'. These can be used to change the size of the marked area. You do this by clicking and dragging inside the small square with the left button.



The entire area can be moved by clicking and dragging inside it with the left button.

Finally you can clear the marked area by clicking outside of it with the left button.

The corresponding info palette, displays the coordinates of the top left hand corner of the marked area and its width and height.

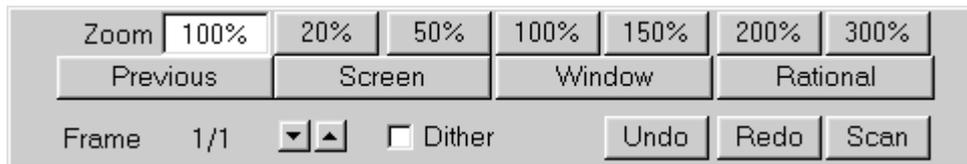
X	410	W	62
Y	238	H	70
Frame	1/1	<input type="checkbox"/> Dither	<input type="button" value="Undo"/> <input type="button" value="Redo"/> <input type="button" value="Scan"/>



The Magnifier

This tool allows you to change the scale of the display. Clicking and dragging with the left button will allow you to mark off an area, when you release the mouse button, the area will fill the window.

The top row of icons allow specific zooms to be set. The icon next to the word 'Zoom' shows the current zoom and new values can be typed in.



Previous will revert to the last zoom, Screen sets the scale so that the image would fill the screen, Window similarly sets the scale so that the image would fill the current window.

Finally, Rational sets the zoom so that one pixel in the original image, fills at least one screen pixel. You use it when you want to see all the pixels of the original image. Images will almost always look better when displayed in this way, rather than when scaled down so that some lines of pixels are missing.

In any icon where you type a value, you can type in an expression.

For example;

10+20

6*2.123

Operator precedence is

unary +/-

(and)

/ and *

units and %

+ and -

More advanced examples;

100mm+10%

5/6"

here " is short for inch, but the interesting bit is that the value of this is (5/6) inches and not 5/(6 inches).

Units supported and abbreviations.

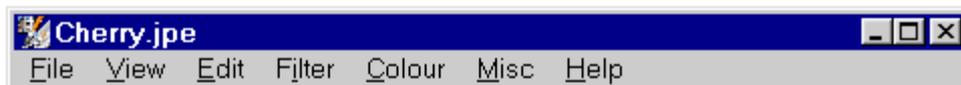
pixel	px (also default)
inch	in, ", i, I
millimetre	mm, m, M
didot	di, d, D
cicero	ci, CI
foot	ft, '
yard	yd
points	pt, p, P
picas	pi, PI
centimetres	cm, c, C
twips	tw

The didot is the equivalent of the point outside the UK and USA. There are 12 didot's in 1 cicero. A point is 1/72nd of an inch. A didot is 0.0148 inch or about 1/68th of an inch. There are 12 points in a pica. There are 1440 twips in an inch or 20 twips in a point.

Values may contain 0..9, the decimal point represented by '.' or ',', if prefixed by # they are treated as hexadecimal and may also contain a..f. Note that #100d is 4109 and not a value of 256 didots which could be written (#100)d.

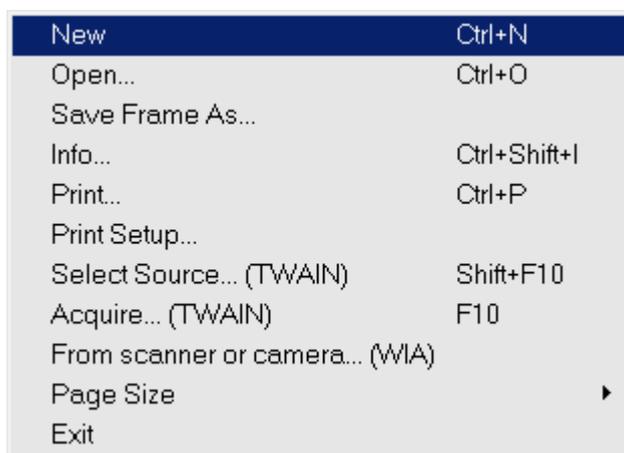
The Main Menu

Displayed at the top of the DPScan window is the main menu.



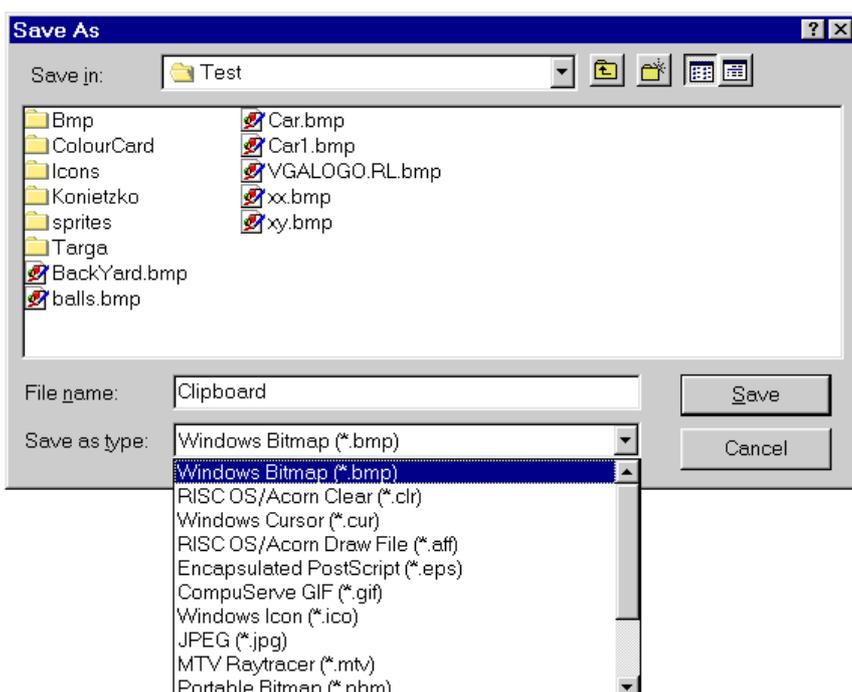
The general operation of the program is controlled by the File and View menus. As you'd guess the File menu is concerned with files, saving them, and settings for the entire file. Meanwhile the View menu controls how the image is displayed. The other menus allow the image to be processed, and are described in subsequent chapters. Here we will describe the File and View menus.

The File menu



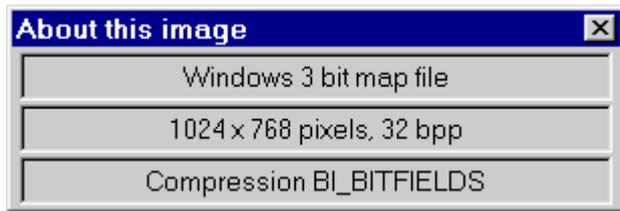
The first entry, **New** opens a new DPScan window.

Open launches the usual Windows dialogue for opening a file. The first image in the file you select will appear in a new DPScan window. Files can also be opened by double clicking them in Windows Explorer if they are associated with DPScan. You can also drag a file from Windows Explorer into a DPScan window.



Save Frame As... allows you to use the usual Windows dialogue to save the current image as a file. All the file types you can use will be shown.

Note that this option saves just the image on display. If the file you open contains more than one image they won't be saved.



The next entry on the File menu **Info**, displays a window of information about the current image. This can often be interesting.

Print

This opens the usual Windows Print dialogue.

Print Setup

This opens the standard Windows Setup window allowing you to configure printer settings.

TWAIN and Scanning

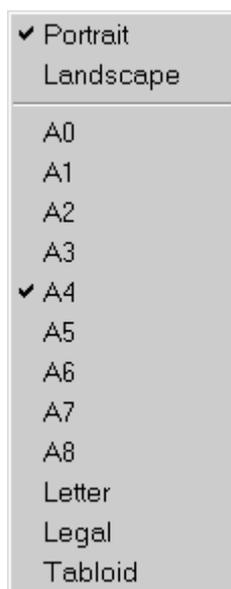
The next two items on the file menu are the standard TWAIN entries of “Select source...” and “Acquire...”. The first of these lets you choose the device you wish to acquire a new image from, the second lets you actually acquire it, via the TWAIN source dialogue box.

If you have used the rectangular selection tool, then the data source will be told that you only want it to supply an image covering that area.

More details of these options were given in the previous chapter.

From Scanner or Camera

This allows you to capture images from devices using the Windows WIA system (akin to TWAIN). Sometimes devices have both a WIA and a TWAIN driver, sometimes they only have one.



Page Size

The next entry on the file menu allows you to set the size of the main window work space. The top two items allow the orientation of the work space to be set. The others allow a range of standard paper sizes to be chosen.

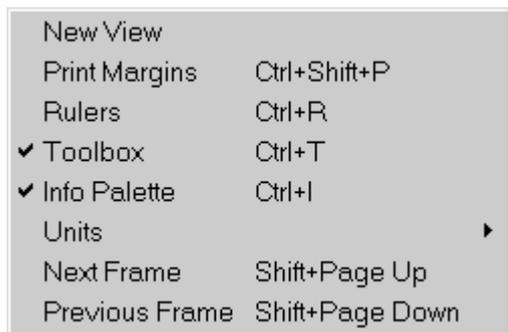
Exit

Closes all windows and removes the program from memory.

When you modify an image, the title bar of the window will display a *. If you try to close the window, a warning message will be displayed, allowing you to save the image.

More than one image window can be open at a time, you can achieve this by just repeatedly opening files.

The View Menu



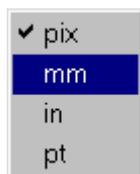
The first entry allows you to have another view of an image. This may allow you to see the image at a different scale, or to see a different area of it.

Print margins, controls if the area of the page available for printing will be displayed, as a light grey area on top of the normal white background and the image.

Rulers. This allows rulers to be displayed along the top and left hand sides of the window.

Toolbox. This toggles the toolbox window on and off.

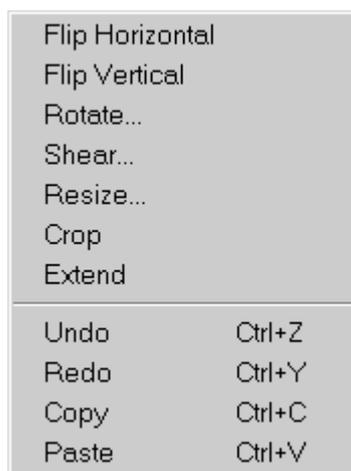
Similarly the next option allows the info palette to be opened or closed.



Units points to a small sub menu that allows you to set the units used throughout the program. The choice is of millimetres, inches, points (1/72nd of an inch), or image pixels.

The last two options allow you to move to the next and previous frame in the current file.

5. Edit



This menu is the third entry on the main menu. It controls operations altering the physical aspects (e.g. size) of the current bitmap.

Flip Horizontal

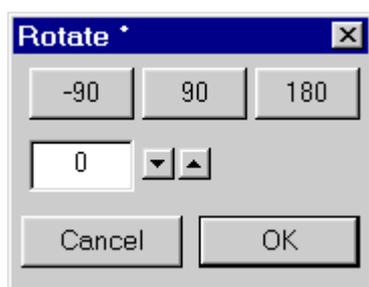
This has a similar effect to reflecting the image in a mirror running from top to bottom. The points on the left of the image, move to the right and so on.

Flip Vertical

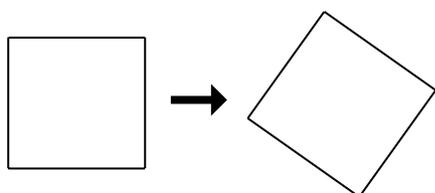
Similar to Flip Horizontal, but doing the reflection in a horizontal line.

Rotate

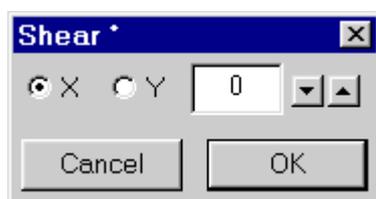
This allows the bitmap to be rotated. Rotations by 90 and 180 degrees can be done more efficiently, so extra buttons are provided for these. Other angles can be set in the writable icon, either by typing them in directly, or by using the arrows to step through the possible values. Clicking on OK will then carry out the rotation.



Angles are measured anticlockwise - a positive angle represents a rotation anticlockwise.



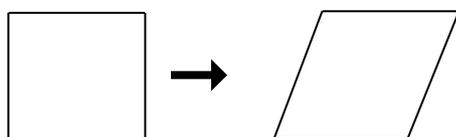
When a bitmap is rotated by an angle other than 90 or 180 degrees, it will cover a larger area - bitmaps must always be rectangular. The new area will be in the current background colour. This can be set from the palette window as described in chapter 7.



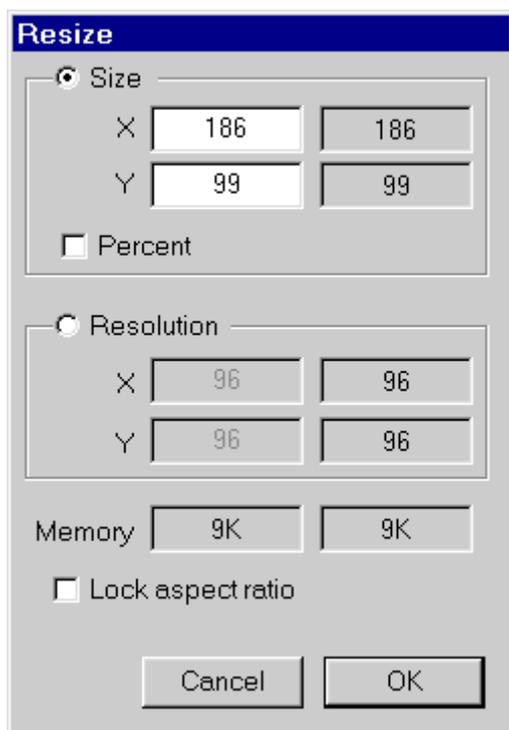
You can Control drag with the pointer tool to find angles if you're trying to make things horizontal or vertical, see chapter 4.

Shear

This pushes the bitmap over sideways. In other words it converts a rectangular shape into a parallelogram. Skew is another term for this effect.



It is possible to shear in either the X or Y direction, by selecting the appropriate button on the dialogue box.



Resize

This allows you to change the size of an image. In general there are various ways to do this. Changing the size is equivalent to changing the resolution of the image.

Simply because an image which is twice as wide will have twice as many points - as will one which has the same dimensions but twice the horizontal resolution.

The dialogue box lets you specify the scaling in one of three ways. Either a percentage change - select **Percent** and the **Size** option. For example 200 will make the image twice as big.

Secondly you can specify the new dimensions of the image - clear **Percent** and select **Size**.

Finally you can give the image a new resolution - select **Resolution**.

To use any of these options you first need to select the appropriate icons, and then type in the values you desire. Clicking on OK will start the scaling process. The text icons on the right hand side, show the old values for the parameters of the image. The two 'Memory' icons show how much memory the old image used, and how much the new one will. The Lock aspect button will when selected ensure that the image retains the same ratio between the X and Y dimensions.

Although scaling is a common option in bitmap programs, usually it is done in a simple way. For example, if you want to make an image twice as tall, the obvious scheme is to just duplicate each row in the image. However a moments thought will show that this is not the best approach. Imagine that two neighbouring pixels have the values 1 and 3, a good guess would be that any pixel inserted between them should have the value 2 i.e. $(1+3)/2$ or the average. This sort of technique is called 'anti-aliasing'. It can be done in various ways. You can specify the 'filter' to be used from the General configuration dialogue box. The choice of filter will also alter the results of rotation and shear operations. See chapter 8.

Antialiasing will give better results on images that support a wider range of intensity levels. It may well be better to convert your image to a larger number of bpp (see chapter 7) before you start processing it.

Crop

This entry on the Edit menu, will reduce the size of the image to that part of it which overlaps the rectangular selection.

Extend

The 'Extend' entry will make the image large enough to fill the rectangular selection.

Undo

The number of steps you can undo is set from the general configuration window.

Redo

Redoes anything undone.

Copy

Copy the current image to the Windows clipboard.

Paste

Paste the Windows clipboard contents into the current window. On Windows pressing the Print key copies the current screen to the clipboard, so Print followed by Paste is a quick way to do a screen snapshot. Similarly Alt+Print copies the active window to the clipboard.

6. Filter

Filters allow you to modify an image depending on its contents. For example using filters you can blur images or sharpen them. There are two main types of filters 'User filters' which can be set up and modified at will, and the filters which are built into the program. The filters are accessed from the fourth entry on the main menu. If there is a selected area, the filter will apply just to that, otherwise it will apply to the entire image.

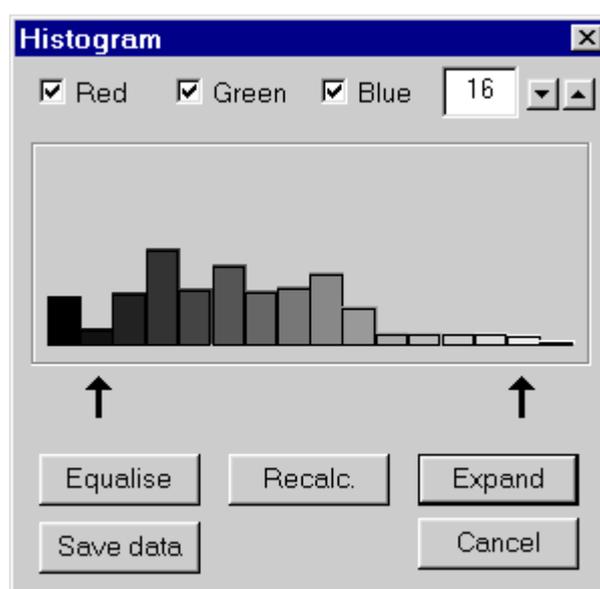
Built in filters

Invert

This replaces pixels in the image by their inverse. Black becomes white, colours are replaced by their complements.

Histogram

This is slightly more than just a filter. When you select it a histogram of the image will be calculated. This shows the number of pixels in each of a number intensity levels or bands.



The Red, Green and Blue buttons can be used to choose which components are displayed. You can change the number of levels with the writable icon at the top right and the associated up and down buttons.

The 'Equalise' button will attempt to spread out the pixels in the image so that there are equal intensities in each band.

Under the histogram, are two arrows which can be dragged from side to side. When Expand is clicked on, the range of intensities between the arrows will be expanded to fill the full range.

Expand and Equalise are crude tools. They may be of use with images where all

Filter

the pixel intensities have been squashed into a narrow band. Finally the 'Recalc' button can be used to recalculate the histogram if the image is changed.

Right clicking over the histogram window or clicking on the Save Data button will open a standard Windows save window that will let you save the histogram data as a text file. This might be useful to analyse data in another program like a spreadsheet.

Max

This filter replaces each pixel by the maximum of itself and the surrounding 8 pixels. The maximum of 1,2,3,4,5 is 5

Min

This filter replaces each pixel by the minimum of itself and the surrounding 8 pixels. The minimum of 1,2,3,4,5 is 1

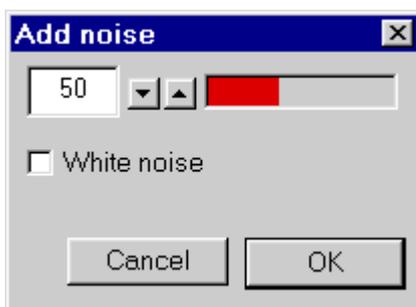
Median

This filter replaces each pixel by the median of itself and the surrounding 8 pixels. The median of a set of values is the value which splits the set into two equal parts. For example the median of 1,2,3,4,5 is 3. This filter is useful for removing noise from an image.

Average

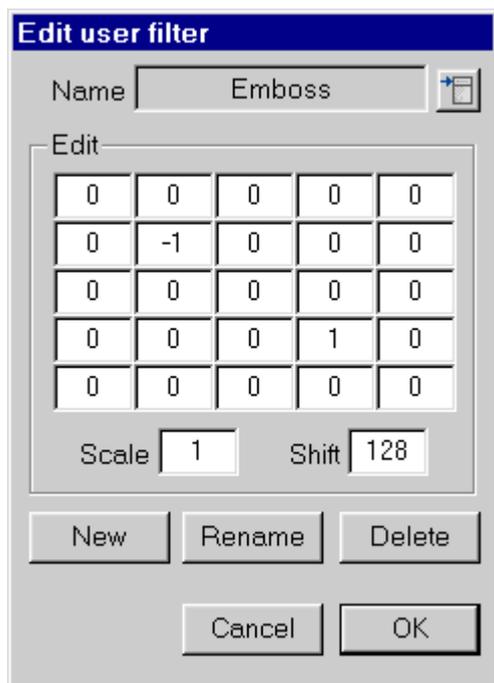
This filter replaces each pixel by the average of itself and the surrounding 8 pixels. For example the average of 1,2,3,4,5 is $(1+2+3+4+5)/5=3$.

Add Noise



This adds random noise to the image. The amount of noise is controlled by the slider and writable icon at the top of the dialogue box. The button 'White noise' controls if equal amounts of noise will be added to the colour components of the image. If 'White noise' is selected, the image will retain its colour. If it is not selected the image will tend to develop coloured dots. Noise may be useful in making

images look more life like. For example images which have been processed into a small number of brightness levels. By adding noise the obvious levels can be smoothed out slightly.



User Filters.

Technically speaking all the user filters are 'convolution filters' - each pixel is replaced by a combination of the values of its neighbours and itself. The combination is a fixed set of numbers - unlike for the Maximum and other filters above. The coefficient of each pixel can be specified, along with a global scale factor and a shift.

The user filter menu will always have as its first entry 'Edit...'. Selecting this will open the Edit user filter window. Underneath this the names of all the filters that have been defined will appear. These filters work like the built in ones, clicking on the name applies the filter. User filters will apply to just the selected area of the image, or to all of it.

Some of the user filters supplied are 'Sharpen' to make images look sharper, and 'Blur' (the opposite of sharpen) to make images more fuzzy. It can be useful to sharpen images before applying some of the colour dithering techniques described in Chapter 7.

Edit user filter

The centre position in the matrix of writable icons represents the current pixel, the other icons then represent its neighbours. Often only the central 3x3 area will be used. The bottom two icons show the overall scale and shift factors. Pixel intensities are always between 0 and 255.

Clicking on the right pointing arrow at the top of the window will open a menu showing all the filters that are defined and by selecting one, you can change the filter displayed.

The row of buttons along the bottom of the window allow you to carry out various operations on user filters.



New

This opens a window allowing you to enter the name of a new filter and the name of the filter to base it on.

Filter

Rename

This opens a window similar to New, but allowing you to change the name of the filter.

Delete

This deletes the current filter.

Cancel

Closes the Edit filter window, and leaves the current filter unchanged.

OK

Close the Edit filter window and saves the new filter definitions to disc. More technical users may wish to know that the filter definitions are kept in a text file called Resources.Filters.

Some examples may make how to use these filters more clear. The first example is a filter that does nothing, it just leaves the image unaltered.

0,0,0	Shift=0
0,1,0	Scale=1
0,0,0	

Only the central 9 entries are shown, the rest should be 0. So for each pixel there will be one contribution from the current pixel and none from its neighbours. The output intensity will be the same as the input.

Next a filter that inverts the image.

0,0,0,	Shift=255
0,1,0,	Scale=-1
0,0,0	

A filter that sharpens an image (Sharpen).

-1,-1,-1	Shift=0
-1,10,-1,	Scale=2
-1,-1,-1	

note that the sum of the central pixels and its neighbours is $(-8+10)/2=1$.

By changing the filter weights we can construct a filter that has the same effect but is not as strong (SlightSharpen).

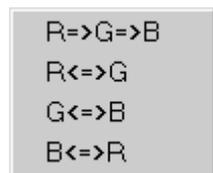
-1,-1,-1,	Shift=0
-1,20,-1,	Scale=12
-1,-1,-1	

again note that $(-8+20)/12$ equals 1.

7. Colour

Colour is the fourth entry on the main menu, as the name suggests it allows control of the colour of an image. This ranges from actually changing the colour to changing the number of bits used to represent it. Some of the colour options are very simple, others provide a lot of scope for experiment. As usual, the menu options will be dealt with in order.

Swop

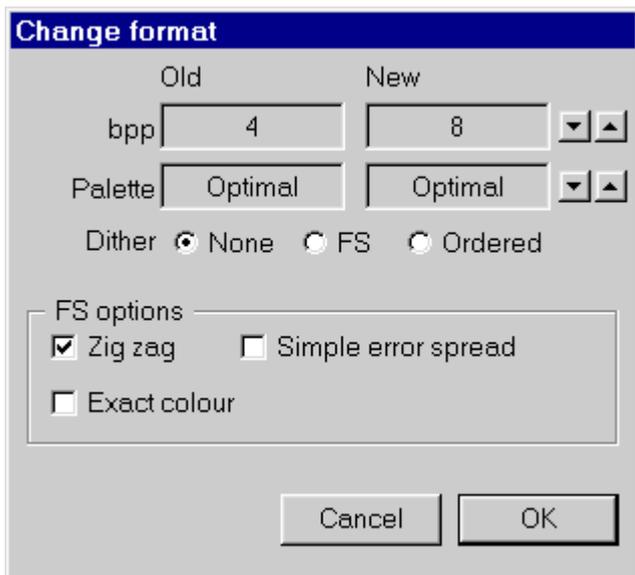


This allows you to interchange the red, green and blue components of the image. It may be of use if for some reason the components have been confused or for effects. This function can be applied to a selected area as well as the entire image.

The first entry rotates the colours red becomes green, green blue and finally blue red. The other options swop two components.

Grey Scale

This converts a colour image to black and white, or in general a series of shades of grey. The weights used to do this can be set from the General configuration dialogue box. See chapter 8. These weights define how bright a colour is perceived to be. Because the eye is more sensitive to green light, green is given a greater weight. If a non-palettred image (more than 8bpp) is grey scaled it will automatically be converted to 8bpp with a linear grey scale palette. This produces a useful saving in memory, and loses no quality. Throughout the program, there is special code to deal with grey scale images, although they are always stored as 8bpp images with a palette.



Change Format

This is probably one of the most complex functions in the program. The basic problem is taking an image with one bpp value and convert it into another image with a different bpp value. If the bpp value of the final image is greater than that of the starting image then usually there is a perfect mapping. The final image will look the same as the source image. Problems arise when the conversion is in the opposite direction from a larger bpp to a smaller one. There is also a

case of converting between the same bpp but with a different set of colours in the source and destination image palettes. Some examples may make this clearer.

Many printing processes can only reproduce black or white dots (a 1bpp image) yet often one sees illustrations in shades of grey (perhaps 4 or 8bpp). This has been done for centuries by a technique called 'half-toning' (also called ordered dithering) which allows shades of grey to be represented by combinations of black and white dots.

You may want a GIF file for a web site, such a file can only contain up to 8bpp yet images off your scanner may be 24bpp, or perhaps have colours in their palettes that do not correspond to those in the standard desktop palette. The familiar RISC OS program ChangeFSI manages to convert such images by a process called Floyd Steinberg integration. Colours are chosen so that on average they represent the colours in the original image. This is done by finding the closest colour to the current pixel, and then calculating an error. The error value is added to the value of the next pixel and the process repeated.

This is the Change format window. The icons on the left show the bpp and palette type of the original image. Those on the right the corresponding values for the new image. You can step through the values on the right with the arrow icons.

Three conversion methods are supported.

None

In this colours in the new image will simply be the ones closest to those in the old. This is sometimes called 'thresholding'. In the case of going from a grey

scale image of 8bpp to a black and white image of 1bpp, pixels would always be black or white depending on if the value of the original was greater or less than a threshold value say 50%. Usually for realistic images this option is of little use, it may be more useful for getting 'posterised' effects.

FS

This is Floyd Steinberg dithering. It can give good results. Typically you would use this option to convert a 24bpp image that you had scanned and manipulated to 8bpp before saving it as a sprite for use in other programs or where it will be displayed on the desktop.

Zig Zag controls how the images are scanned. When selected the scan will change from left to right and right to left on alternate lines. In general Zig Zag will give the best results.

Simple controls how the error in the colour of one pixel is shared with its neighbours. When Simple is selected all the error is passed to the next pixel to be converted on the same line. This simple dithering is the same as is used for the main program display when 'Dither' on the tools palette is selected. When Simple is not selected, true Floyd Steinberg dithering is used, with the error being spread over both the next pixel, and neighbouring pixels on the next line. In general Simple not selected will give the best results. Using Simple may be useful if you want to get an image that is exactly the same as the one on screen.

Exact. By default colours are found using a lookup table of nearest values - this has limited precision. If Exact is chosen, a search for the nearest value will be carried out for each point. Exact takes a lot longer to do the conversion, but can give the best results.

Ordered

This is traditional half-toning. Results can be almost as good as FS dithering, and are usually a lot faster.

New bpp

It is possible to choose any of 1,2, 4, 8, 16, 24 or 32 for the new images bpp. It is important to set this. There's little point changing to the same bpp, unless the palette is different.

New palette

This can take various values.

RISC OS

The RISC OS palette is defined as that used in the desktop by default. In 8bpp modes there are limitations on the possible palettes. This is the most common setting for the new palette, since usually images will be for display on the desktop.

Colour

Optimal

If this is selected, the image will be analysed and a palette created that represents it best. You use this if you want the best results, and are going to display the image on hardware that imposes no limitations on the palette.

RGB

This selects 'digital RGB' for the new palette. This means that the ranges of red, green and blue are split into equal parts. For example, in an 8bpp image, RGB can be given 5 levels each. This palette can be useful for producing images to be printed on printers that can reproduce red, green or blue dots with each dot being 100% on or off.

Linear GS

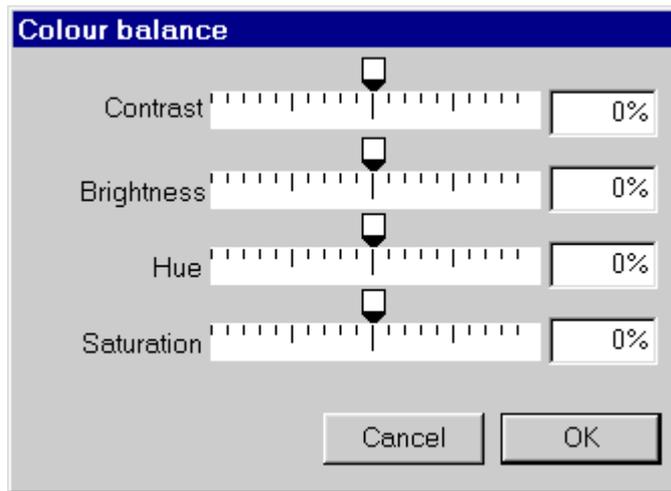
A linear grey scale palette. This can only be used on grey scale images.

None

This is used to indicate that the palette is the same as the original. However, when used in the old values icons, None indicates that the old image really does have no palette for example a 24bpp image.

There is a folder called "Palettes" in the folder where DPScan is installed. RISC OS Palette files in this will appear as destination palette options in the 'Change Format' window. One palette file suitable for 'Netscape' is supplied. See the section on the colour palette for details of how to load and save Palette files.

Colour Balance



This window, gives you a set of sliders that can be used to give global changes to the images colour. Changes can be applied to the entire image, or to a selected area.

The four values that can be changed, are related to the HSV colour model. The best way of understanding this, is via the colour wheel displayed by the palette window, and described in a later section of this chapter. The

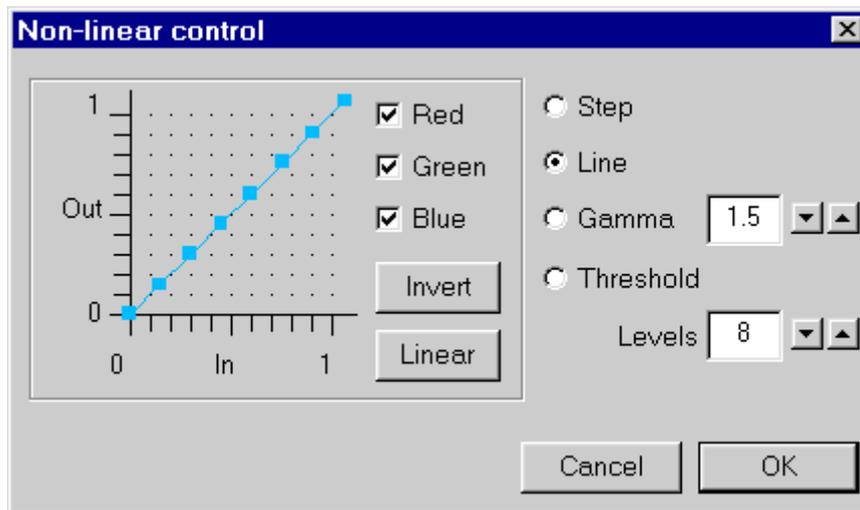
concept of brightness should be intuitively obvious. The brightness slider makes all colours brighter or darker. The contrast slider has a similar effect, except brighter colours have their brightness level increased more than dimmer colours. Note that the values shown are of changes to the image. If you enter a 5% contrast increase, and click on OK, then to get back to the original image, you have to enter -5% and click on OK again.

Saturation controls how close to grey colours are. By decreasing the saturation, you can gradually move an image towards being grey. Increasing the saturation makes the colours deeper and less like pastel shades.

Hue is related to the actual colour. Changing the hue, moves colours around the colour wheel in circles. It is almost like the colour swapping described above, instead of rotating RGB in a single large discrete step (green becomes red etc.), changing the hue allows this process to be done in small steps.

These operations are useful for correcting scans which have been done at the wrong values of brightness or are discoloured.

Non-linear



This allows the most sensitive control of the brightness of an image. The graph on the right of the window shows the output brightness for each value of the input brightness.

The resulting change can be applied to the entire image or a selected area of it.

The buttons on the far right allow the mapping to be applied to the red, green or blue components individually. The buttons on the left set the form of mapping.

Step

In this the input is mapped to a discrete set of output values. Commonly called 'posterisation'. The levels of the bands in the graph can be dragged with the mouse select button.

Smooth

A line is drawn through a number of values. You can drag the points forming the line around.

Gamma

A common form of correction is to relate the input level (i) to the output (o) by $o=i^G$ for G greater than 1 this will make the image darker, for G less than 1 brighter.

Threshold

In this the output is either 0 or 100%, and you can vary the point (threshold) at which the transition takes place.

Levels

This allows you to set the number of levels that there will be for the Smooth,

Step and Threshold options.

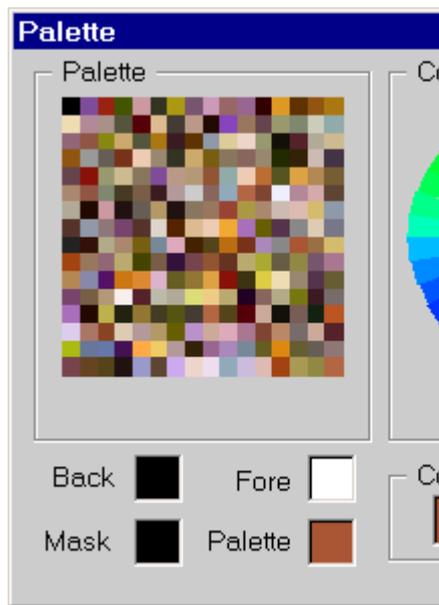
Invert

When clicked on this will invert the mapping.

Linear

This makes the levels linear. It may be useful for getting back to sensible values after a session of editing them.

Palette



Selecting this from the Colour menu will open the palette window. This can be useful just to see what the palette of the current image is. Using the palette you can choose the background colour to be used when an image is extended by the Edit tools (shear, rotate, extend), construct custom palettes, and finally set the foreground colour and tint to be used by the selections on the rest of the Colour menu - Tint, Filter, Add, Subtract and Reverse filter.

The top of the area, shows a small rectangle for each entry in the palette. The example on the left has 256 as it would for an 8bpp image. For images with less bpp, there will be less entries. For images with no palette this area will be blank.

The four squares at the bottom show the current foreground, background, mask and palette colours.

Mask, is the strength of the mask which is applied to new pixels created, for example when an image is rotated. The idea is similar to that of the background colour.

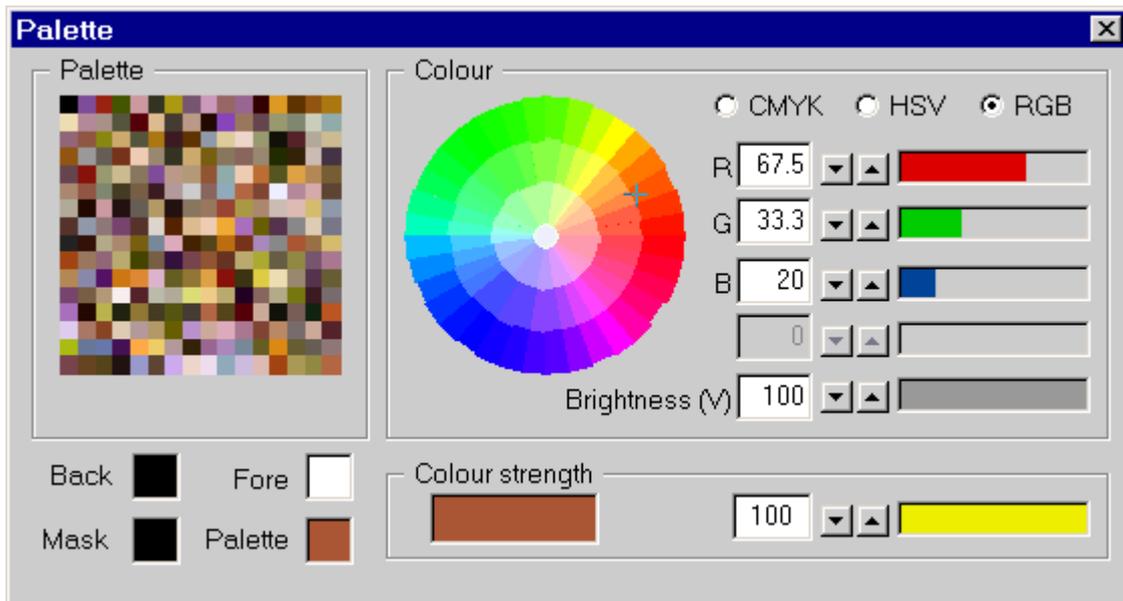
The mask value is taken from the colour strength slider. A value of 100% (white) means that all the brightness comes from the image (solid), a value of 0% (black) means that the all brightness comes from whatever the image is rendered on (transparent).

Left clicking with on these icons transfers the colour to the colour selector values (see later). Right clicking transfers the colour selector values to the colour.

If you left click on the palette area that colour will become the one shown on the palette icon and the colour selector. Right clicking in the palette area will transfer the current colour selector colour to the palette. A similar effect can be achieved by right clicking on the Palette icon.

As well as choosing a colour, you can also choose a strength to go with it - a tint. This controls how much effect the colour will have when mixed with other colours. The right bottom part of the palette window allows this strength or tint to be set. The icon under "Colour strength", shows the tinted and the original colour.

The right hand part of the window, displays the 'current colour'. It can be shown in one of three models, RGB (red green blue, as used by video screen displays), HSV (hue, saturation, brightness which has some intuitive value) and CMYK



(cyan, magenta, yellow, black, which is the same model as used in colour printing). The colour wheel displays a plane in HSV space - think of it as a slice through a cylinder.

The colour model can be changed by clicking on the RGB, HSV or CMYK buttons. Sliders will appear for the quantities that can be altered. In the RGB and CMYK models there is an additional slider to control the slice through the colour cylinder that is displayed.

You can select colours by typing in values (percentages of red, green blue etc.), by clicking or dragging on the colour wheel, or by stepping through values with the arrow icons.

The process for using the colour selector is as follows.

Left click on a colour icon i.e. the Fore, Back, Mask or Palette icons or the palette area icon.

Modify the colour using the colour selector.

Set a tint for the colour.

Finally store the colour by right clicking on one of the colour icons.



It is possible to get a colour from the current image, using the pointer tool, by holding down the SHIFT key and left clicking over the pixel of interest.

Right clicking over the palette area on the left will open a menu which allows you to export the palette as a RISC OS palette file or import a palette file. You can also import a palette file by dropping it into the main image window.

We now come to the final selections on the colour menu.

Tint

Tint takes each pixel in the image, and mixes its colour with that of the current foreground colour. The effect is like mixing light. If the current colour is Yellow then all pixels will become more yellow. How great this effect is, is set by the tint of the current foreground colour. Do not confuse this Tint option with the tint property of colours.

Filter

This works like viewing the original image through a colour filter (like looking at a picture through a coloured piece of paper). Colour filters have interesting properties. For example if you look at a pure green object through a pure red filter it will appear black - because no green light can get through a red filter. Only pure red objects will be seen at full strength through a red filter. As usual how much effect the filter has is set by the tint of the foreground colour.

Add

The foreground colour is added at equal strength to all the pixels in the image. Again the tint of the colour controls the effect.

Subtract

The opposite of Add, the colour is subtracted.

Reverse Filter

This reverses the effects of putting an image through a colour filter. Imagine you had taken a photograph through a coloured filter, then this option would undo the effects of the filtering, producing an image that looked as if no filter had been used.

One use for this is processing scans of colour negatives. This is what you should do with a colour negative scan.

1. Open the colour palette and left click on the film background in the image window, using the Select tool and with SHIFT pressed. Set the colour tint/strength to 100%. Finally right click on the foreground colour icon. You have now set the foreground colour to the same colour as the negative film.
2. Use Reverse filter from the colour menu.
3. Use Invert from the filter menu.

Your colour negative is now developed. Scanning at high brightness will help get better results, as will applying some gamma correction.

Profile

See the next chapter, Colour Management.

8. Colour Management

Typically the images which DPScan handles contain information on the colours of each pixel. For example the colour of a pixel might be 3 units of red, 6 of green and 10 of blue.

The question that arises is what is meant by blue? Is it sky blue or navy blue?

A related problem is that some colours appear differently in different situations - which is why people occasionally take clothing into daylight to see its true colour.

Colour management attempts to solve these problems.

In the early days of computers, an image pixel with 10 units of blue would be displayed with 10 units of blue for that particular system - i.e. combination of monitor and graphics card. That meant that images would look different on different computers.

To solve this problem an extra piece of information has to be added to the image, this is called a colour profile. The profile provides absolute information about the colours. So an image with a profile tells us not just 10 units of blue, but 10 units of an exact shade of blue.

Profiles can be contained in images (embedded) or they may exist on your computer separately from the image, and there may be some way of linking them to the image - for example by name.

Profiles apply not to just images but to devices like scanners and printers. In other words telling us what colour 10 units of blue from a scanner is, or what colour will appear on a printer if it is told to print 10 units of blue.

Many images don't have colour profiles. One reason is that many file formats don't support profile information. To cope with this a default profile has been invented by the computer industry - sRGB. It is common to find devices that produce images with this profile. By assuming this default profile applies in the absence of any other information, it is possible to get by without explicitly handling profiles.

For example if your scanner, monitor, camera and printer all have the sRGB profile it doesn't matter if the images you save contain no colour profile information.

In general this happy situation won't apply and it is necessary to have profiles for all your hardware and save profile information with images.

As far as Windows is concerned colour profiles are a bit like fonts. Colour profiles have the extensions .icc or .icm (the files are identical regardless of extension) and they all live in a common folder.

Colour Management

The folder is located at:

Windows 95/98/ME C:\WINDOWS\SYSTEM\COLOR\

Windows NT C:\WINDOWS\SYSTEM32\COLOR\

Windows XP/2000 C:\WINDOWS\SYSTEM32\SPOOL\DRIVERS\COLOR\

Usually devices will as part of their installation put their profiles into one of these folders. An installed profile can be used by any application.

You can also find profiles on the web. Adobe have a particularly useful set available for free from their web site.

Colour profiles have an associated colour space, for example RGB or CMYK. It is necessary to pick a profile with the right colour space for the situation. For example a CMYK colour space profile cannot be used for a monitor that uses RGB.

DPScan supports colour profile information for .bmp, JPEG and TIFF image file types. The File menu **Info** entry will tell you if an image contains a colour profile.

Colour Management is controlled by two windows in DPScan, these are available from **Profile** on the Colour menu and **Colour Management** on the Misc Menu. **Profile** controls profile information for the image in the current DPScan window. **Colour Management** controls settings for all images loaded into DPScan.

Colour Management

Monitor

Profile is the profile for your monitor. Only RGB colour space profiles can be selected from. The icon to the right sets the rendering intent. The **Colour Management** button controls if the profile will be used when displaying images on screen. The **Proofing for printer** button allows you to select a preview of what would be displayed on the printer. Note that this involves both the printer and the monitor profiles. **Use system profile**, reads the profile set in Windows for the monitor and sets it as the monitor profile in DPScan.

Printer

Profile is the profile for your printer. Only RGB colour space profiles can be selected. As with the Monitor a rendering intent can be selected. The **Colour Management** button controls if profiles will be used when displaying images. This can be useful because some printer drivers have colour management built in, if that is the case you may not want to use a colour profile for printing. **Use system profile**, reads the profile set in Windows for the printer and uses it when printing. If this option is selected the profile name displayed is that for the currently selected printer in DPScan.

Image Import

The remainder of the window is concerned with profiles applied to images that are loaded into DPScan.

CMYK Profile

This is the default profile for CMYK images i.e. if an image has no colour profile embedded this profile will be used. Only CMYK colour space profiles can be selected.

RGB Profile

This is the default profile for RGB images i.e. if an image has no colour profile embedded this profile will be used. Only RGB colour space profiles can be selected.

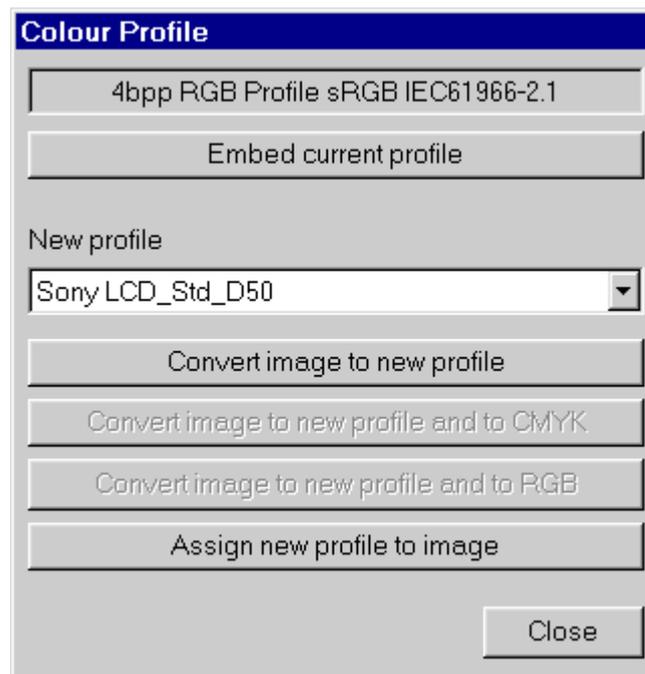
Scanner Profile

This profile is used for images acquired via TWAIN (typically from a scanner). Only RGB colour space profiles can be selected.

Use ICM

This controls which colour management system DPScan uses. When ticked it will use ICM which is part of Windows. To change this setting, you must change the value shown in the window, save choices, and then quit and re-run DPScan.

Profile



The top icon shows the colour depth (i.e. bits per pixel) and the colour model of the image (RGB or CMYK). Following this is the name of the colour profile prefixed by the word “Embedded” if the profile is embedded in the image.

Images with embedded profiles will be saved with embedded profiles (if the file format supports profiles). Images that are loaded and which don’t have an embedded profile will be associated with one of the default profiles (see the Colour Management window), however the profile will not be saved with the image unless it is embedded.

Embed current profile

This embeds the profile being used for the image.

New profile

This allows a profile to be chosen for the following options. All available profiles (CMYK and RGB colour models) are displayed.

Convert image to new profile

This converts the image from its existing profile to the New profile. This conversion results in a change of the data in the image. Say you had an image with an sRGB profile and wanted to convert it to Adobe RGB, you would select Adobe RGB for the New profile then use this option. This option is only available if the colour space of the new profile matches that of the image. Note that this operation will ideally result in no visible change to the image.

Convert image to new profile and to CMYK

This is similar to the previous option, however the image is also converted to CMYK colour model. You’d use this option to convert RGB images to CMYK.

This option is only available if the New profile colour model is CMYK. Note that this operation will ideally result in no visible change to the image.

Convert image to new profile and to RGB

This is similar to the previous option, however the image is also converted to RGB colour space. You'd use this option to convert CMYK images to RGB. This option is only available if the New profile colour model is RGB and the image colour model is CMYK. Note that this operation will ideally result in no visible change to the image.

The resulting RGB image will have 32bpp. It may be convenient to use the Format entry on the Colour menu to change to 24bpp. Some file formats (i.e. .bmp) can save RGB images with 24bpp or 32bpp, 32bpp involves some waste of space.

Assign new profile to image

This removes the existing profile from the image and assigns a new profile to it. The image data is not changed (compare with the previous options). You'd use this option if you decided that the wrong profile was associated with an image. Another use is to remove an embedded profile from an image. This option is only available if the image has the same colour space as the New profile. Note that this operation will usually result in a visible change to the image.

9. Configuration

The configuration dialogue box can be accessed from the choices entry on the Misc menu. Across the top is a strip of file type icons. Clicking on these icons, will open a sub configuration window for the file type.

The options provided for the file types are discussed in the next chapter. The arrows at either end of the strip allow the available file types to be scrolled through.

Across the bottom of the configuration window is another row of icons. These allow various aspects of the program to be controlled. Each icon representing a group of related parameters. Again configuration windows can be opened by clicking on the icons.

Finally there is a 'Save' button clicking on this will save the configuration of the program so that it is the same when the program is next used.

Memory

This controls the programs use of memory. The top slider lets you set the maximum amount of RAM that it will use. The next two bars cannot be adjusted, they show the amount of RAM in use, and the amount of disc space being used as "virtual memory".

Note that these bars can be confusing. The RAM bars show the proportion of the total RAM in the computer. The disc bar shows the proportion of free disc space in use. The free space on the disc may be much larger or possibly smaller than RAM. You should not make relative comparisons of the RAM and Disc bars.

The writable icon 'Temp file' allows the directory where the temporary file holding image data which will not fit in RAM is held. Clicking on the right pointing arrow to the side of this, will pop up a standard save box. The icon may be dragged to where you want the temporary file to be held. Note that if a temporary file already exists, you cannot change the location immediately this way.

It may be useful to leave the memory box open to give you some idea as to how much memory you are using. The RAM and Disc bars will be continuously updated.

The default temporary file path is;

<%TEMP%>

This means use the value of the Windows system variable TEMP. Other Windows system variables can be used when bracketed by %'s. It is also possible to insert the name of a Windows registry variable between < >'s. The temporary file should be kept on a large area of fast writable disc.

Configuration

Since Windows has a virtual memory system built in, the recommended setting for Max RAM is a large value so that DPScan does not use its own virtual memory.

Usually the temporary file will only become active when large images are being handled, in such cases you must have plenty of free space on the disc.

Given the well known fact that all hard discs are 25% too small, the situation of the hard disc with less free space than there is RAM is quite common. At any rate using virtual memory requires large areas of empty disc space.

Windows

The four buttons on this window control the default appearance of each new window. Whether the window will have rulers, a toolbox, an info palette or if dithering will be enabled for it.

Page

This allows the default page size, units and X and Y resolutions to be chosen, by clicking on the appropriate right pointing arrow and making a selection from the resulting menu. These values apply when a new, empty, window is opened.

General

There are several sections.

The first section allows the Aliasing filter used for the operations on the Edit menu to be set.

The second section consists of three writable icons which set the weights used when greyscaling images.

10. File Formats

In an ideal world there would be just one type of file used for storing bitmap images. In reality there are many thousands, some represent needless duplication, others have been designed to take advantage of different hardware, or to meet different requirements. In any case there is a need to be able to save and load files in different formats so that images can be exchanged with other programs and computers. Different file formats are also of use in different situations. It would be a mistake to assume that you should always use one file type, when often using a different file type can save you a lot of effort.

Given the potentially large size of bitmaps, many file formats provide some sort of data compression. By locating redundancies in the data and coding them more efficiently the size of the image files can be reduced. There are two types of compression in use. The first type is "lossless" in this the uncompressed image will be identical to the original image. This sort of compression is similar to that used in compressed filing systems and archive programs. The other type of compression is so called "lossy" compression. In this the decompressed image may be different to the original - the idea being that such changes will be small and not noticeable. Currently only one lossy type is supported - JPEG.

Many image file types were originally based on just saving the contents of the screen display. Such formats often do not save resolution data. The display prescription is to just display one dot of the image on one dot of the display. Commonly this will lead to high resolution images appearing too large. If this is a problem you should use file types like TIFF that do not have such shortcomings. It is also possible to change the display resolution - see the description of the info palette for the pointer tool in chapter 4.

We now go through the supported file types. For each one, the configuration options will be discussed. See the previous chapter for how to access these options.

All the file types have an "associate" icon. This lets you tell Windows to associate a given file extension with DPScan. Association makes DPScan load files when they are double clicked, it also makes Windows display an icon from inside DPScan for files with the relevant extension. Turning off association will not return the association to the previous setting. To do that you will have to use the facilities in Windows (Windows explorer View->Options) or in the program you wish the file extension to be associated with. As well as the normal extensions like .tif for TIFF files, hexadecimal extensions are supported (ff0 for TIFF files), these are what you may end up with when you transfer a file from RISC OS or may let you resolve any clash of extension names with other programs.

File formats

Many of the file types e.g. PCX, BMP, GIF, MTV, PBM, Clear, configuration windows allow a default resolution to be set. This is the resolution in dpi which will be given to images of these types when loaded. Setting it to your screen resolution (usually 96 or 120dpi) is a good idea.

Sprite

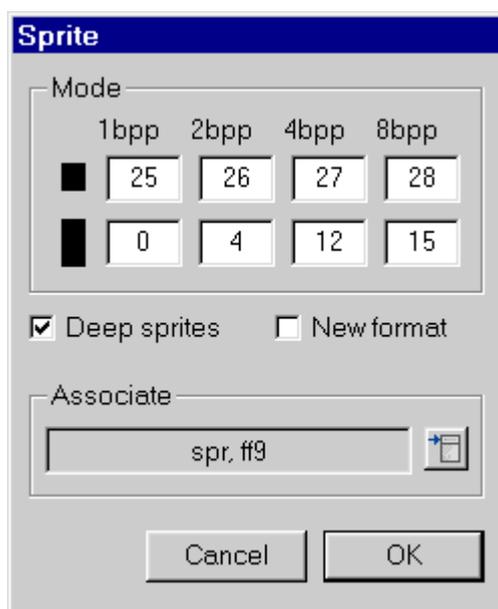
This is the standard RISC OS bitmap file type, as used by Acorn Paint and many other programs. Traditionally sprites come with a "mode" number which defines how they are to be displayed, bpp and resolution etc. The configuration options allow mode numbers to be defined for both rectangular (90x45 dpi) and square (90x90 dpi) pixels. Up until 1994, sprites only held 1, 2, 4, or 8 bpp data.

However after that time, a new variation on the sprite format appeared. This can hold 16 or 32 bpp data. For the 16 bpp data 5 bits are given to each of red green and blue, for the 32bpp data 8 bits. If you save a 24bpp image as a sprite it will be saved with 32bpp. Be careful 16 and 32bpp images cannot be used on older (pre 1994 RISC OS) software or hardware.

Sprite files can contain more than one image. Images may have a mask. Whilst the 1,2 and 4 bpp formats are straightforward. The 8bpp format has been obfuscated by hardware dependencies. Original 8bpp sprites had 64 palette entries of which only 16 were used - the full 256 colours being synthesised from the palette and the colour numbers. Later 8bpp images have a full set of 256 palette entries. When saving sprites the program will supply a 256 entry palette. This may upset some old software.

Deep sprites when selected allows 16 and 32 bpp images to be saved. If not selected these images will be converted to 8bpp. The conversion will be done with the current Change Format settings.

New format when selected makes images with 8bpp or less be saved in the new mode number less format. This format also only uses one bit per pixel for masks. Older sprites with a mode number use the same bpp for both the mask and the image. The old format does not support 16 or 32bpp, such images are thus always saved in the new format, regardless of this setting.



In effect no resolution data is saved with sprites. CMYK images are supported. Images may contain a mask and this is supported. Sprite files may contain many images.

There is a text file called "modes.txt" stored in the folder in which DPScan is installed. It is a common problem to find that sprites which you want to load have been created in some non-standard (soft) screen mode. The Modes file allows you to enter definitions for these mode numbers.

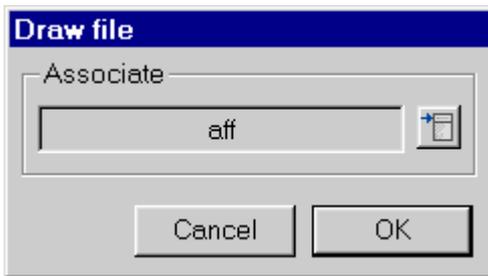
Extensions: spr, ff9.

Draw

Draw files are the standard RISC OS vector file format. They can contain RISC OS sprites. DPScan displays just the sprites in Draw files, letting you step through them.

Saving images as a Draw files is useful because Draw files effectively retain resolution data, unlike sprite files. So Draw is a format you can use on RISC OS which will preserve image sizes.

Extensions: drw, aff.



GIF

This is the Graphics Interchange Format. Originally invented by CompuServe (an online system or large bulletin board). The motivation was to allow pictures to be downloaded (transferred from one computer to another). It seems one of the design requirements was to allow images to be seen as they arrived. GIF (pronounced 'jiff') made use of the then new LZW data compression technique. GIF files give good results and are deservedly popular. GIF files can contain 1, 2, 3, 4, 5, 6, 7 and 8bpp images. They are therefore only of use for paletted images. GIF files can contain more than one image.

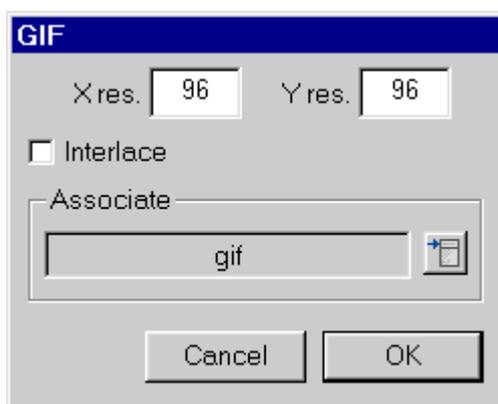
In GIF files transparent pixels are represented by one of the colours in the image palette. So for example if you have a 256 colour image with a mask, one of the colours must not be used by the image so that it can be used for transparent pixels. If all 256 colours are in use by the image, then the mask data will be ignored when saving the GIF file.

One configuration option is provided. Interlace controls if the image lines are stored contiguously (one after another) or with an interlace (every eighth line etc.).

In GIF files transparent pixels are represented by one of the colours in the image palette. So for example if you have a 256 colour image with a mask, one of the colours must not be used by the image so that it can be used for transparent pixels. If all 256 colours are in use by the image, then the mask data will be ignored when saving the GIF file.

No resolution data is saved.

Extensions: gif, 695.



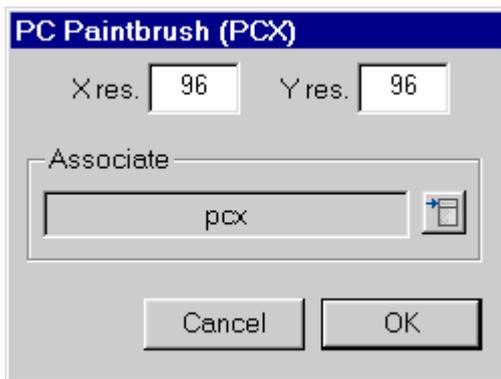
File formats

PCX

This is a file format popularised by ZSofts PC Paintbrush package on the PC. Files may contain 1, 2, 4 or 8 bpp. Simple run length data compression is implemented.

No resolution data is saved.

Extensions: pcx, 697.



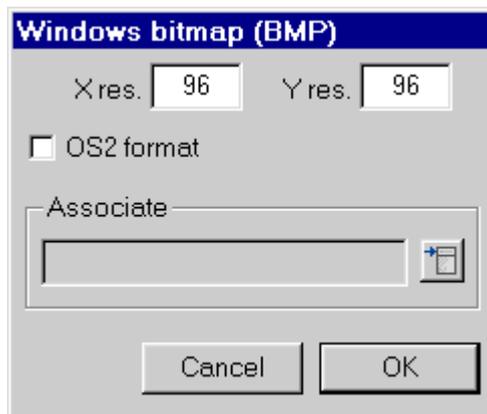
BMP

This is the Windows and OS/2 equivalent of Acorns Sprite file. Images may contain 1, 4, 8, 16, 24 or 32 bpp data. Simple data compression is implemented. The configuration option allows either Windows 3 or OS/2 headers to be written to saved files.

Resolution data is saved.

Colour profile information can be saved. The vast majority of .bmp files in existence have a Windows 3 header. A Windows 5 header is required for an embedded profile. What this means, is that if you save an image in .bmp format with an embedded profile some Windows programs may not be able to use the resulting file.

Extensions: bmp, 69C.

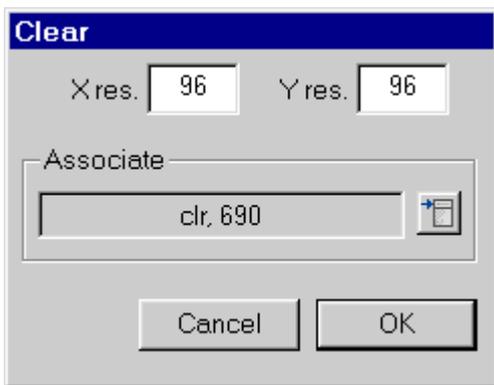


Clear

Given the confusion surrounding Acorn sprite files, and their original inability to support 24bpp, John Kortinck (of Translator fame) invented a file type that was simple and straightforward - hence the name 'Clear'. Clear files can contain 1,2,4, 8 or 24 bpp. There is no data compression. Quite a few RISC OS programs can handle Clear files.

No resolution data is saved.

Extensions: clr, 690.

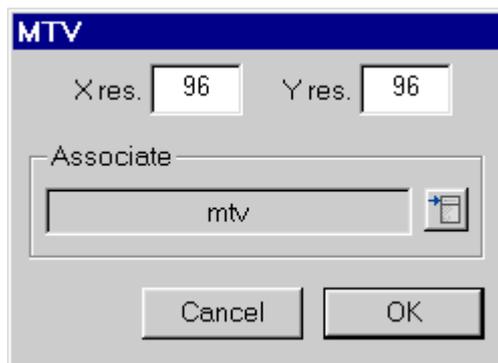


MTV

This file format was invented to save the output of the MTV ray tracer. It is a simple 24bpp format with no data compression.

No resolution data is saved.

Extensions: mtv, 699.

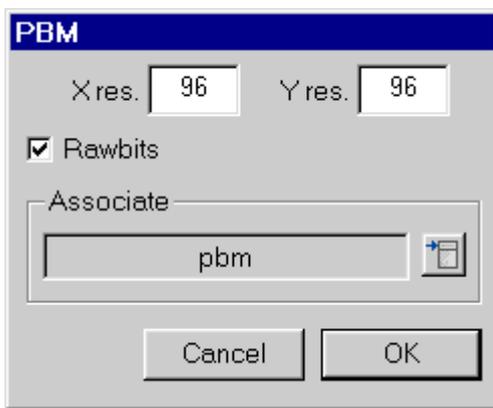


PBM

In the world of UNIX there are several large suites of general purpose image manipulation software. One such is Jeff Poskanzer 'PBM' toolkit. There are variations on the format to handle grey scale (PGM Portable Grey Map), 1bpp (PBM Portable bitmap) and colour images (PPM Portable Pixel Map). The configuration option Rawbits controls if PBM files will be stored as binary or text files. Binary (Rawbits selected) are much more compact than the text file option.

No resolution data is saved.

Extensions: pbm, 69e.



JPEG

This file format is from the Joint Photographic Experts Group who set out to produce a standard for lossy compression. Technically speaking the file format is JFIF (JPEG File Interchange Format). These files can cope with 24 bit images or 8 bit grey scale data. Because the compression is lossy you should only use JPEG files for saving the final result of image manipulation. However JPEG compression is often excellent, reducing huge files to manageable proportions.

The configuration option Quality (0 to 100%) controls how good a representation of the original the file will be. 100% gives no loss of data.

Resolution data is saved.

Colour profile information can be saved.

Extensions: jpg, jpeg, jpe, c85.



TIFF

This file format goes back to the early days of micro computers. It was an attempt to invent a file type for images that could be used by all programs. Developers of TIFF included Aldus, Lotus and Microsoft. TIFF stands for Tagged Image File Format. As the name hints TIFF files are extendible, more tags can be added to give more information about images. TIFF attempts to be everything to everyone. Unfortunately as a result it is possible to have two programs which both use TIFF files and are completely incompatible. However in many ways TIFF is the file format of choice.

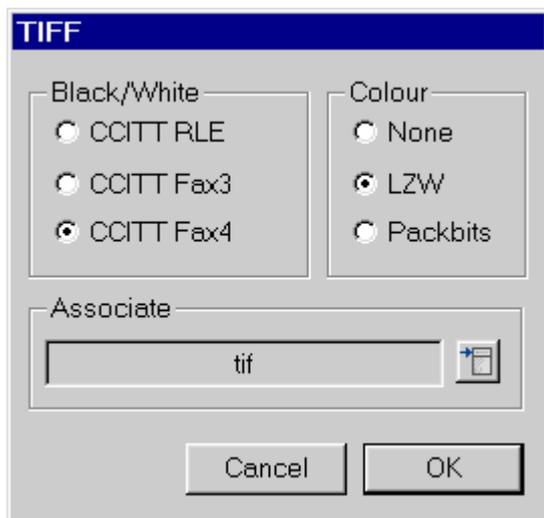
TIFF files can contain all bpp images, including paletted and non paletted images.

Resolution data is saved with the images, and a file can contain more than one image.

Configuration options allow you to choose the data compression method used. For black and white images, special coding techniques from the world of Fax are provided. For colour images there are other techniques. Note that all the black and white options can be toggled off (left click on a selected one) and the colour data compression will then be used for all images.

Colour profile information can be saved.

Extensions: tif, tiff, ff0.



ICO

This is a Windows icon. It is a small bitmap intended for icons in the Windows desktop. Files may contain more than one image, the standard sizes which you should stick to if you want your files to be useable everywhere are 16 by 16, 32 by 32 and 64 by 64 pixels. The idea is that all the images are of the same thing, and the one of the most appropriate size is used.

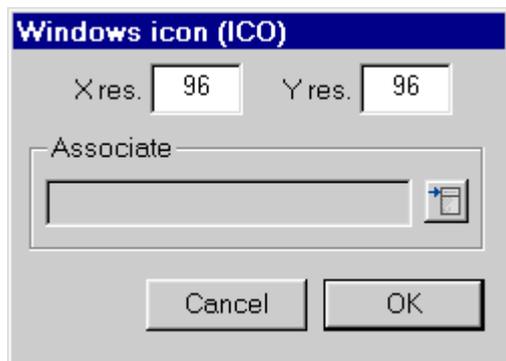
Ico files have a maximum width and height of 255 pixels. An error is raised if an attempt to exceed this is made.

Images support a form of masking, consisting of two sub images one of which is logically exclusive or'd and one which is logically and'd. The thing you need to know is that for masked images to be saved successfully they must have black in their palette (if they have one).

The sub images are related to Windows bitmaps (discussed above) so the same colour depths are supported 1, 4, 8, 16, 24 or 32.

There is no resolution data.

Extensions: ico, 132.

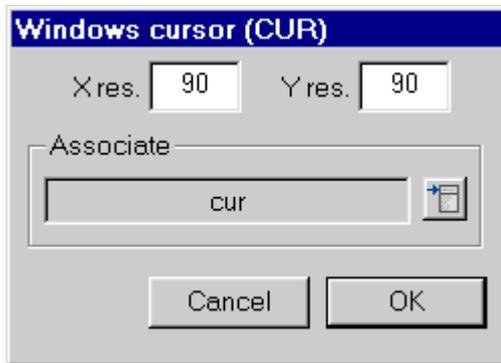


File formats

CUR

This is a Windows cursor. It is actually an ICO file with a different extension, so all the comments above apply.

Extensions: cur, 133.

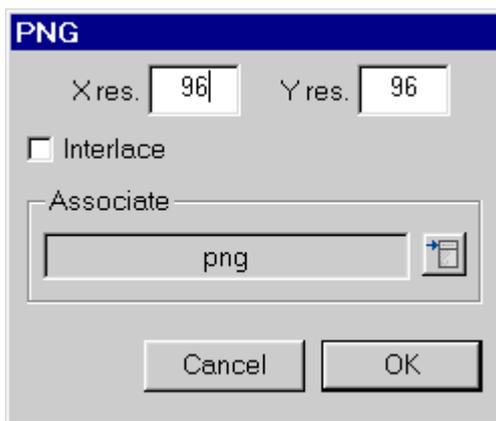


PNG

Portable Network Graphics. Patent problems with GIF files led to an attempt to define a new “open” file format that could be used in the same way as GIF files, but which used a compression method free of patents. In fact the compression technique used is ‘Deflation’ the same as in most Zip files.

It can sometimes offer much better compression than e.g. LZW TIFF’s on 24bpp images. Compression is lossless. The file format supports resolution data.

Extensions: png, b60.



EPS

Encapsulated Post Script. Consists of converting the bitmap into PostScript and adding a small amount of extra information. DPScan can only save this format, it cannot open EPS files.

The EPS configuration window provides options for a preview image and for saving the main image data as an embedded JPEG.

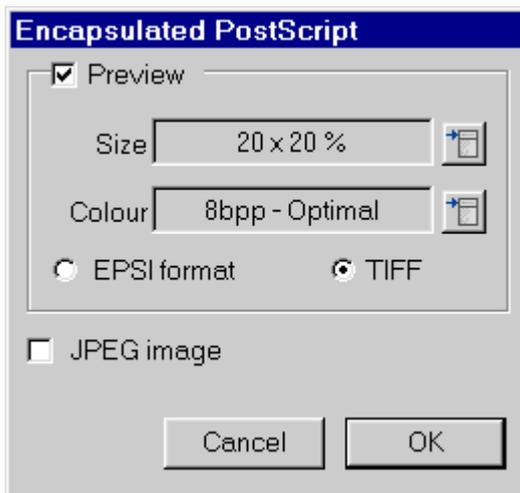
The preview section allows the main scale and change format windows to be used to generate the preview image. Good settings are probably a preview of size 20% at 8bpp with an optimised palette. The preview image can either be a TIFF file, glued on to the start of the EPS file - this is PC/DOS style, or in a truly platform independent form (EPSI) coded as hex in comments in the EPS file.

Remember that the current TIFF and JPEG configurations will be used for the EPS file - you may want to use TIFF's with no compression.

EPS saving supports CMYK, but only when the main image in the EPS is not JPEG'd.

EPS files contain resolution information. Masked images are not supported.

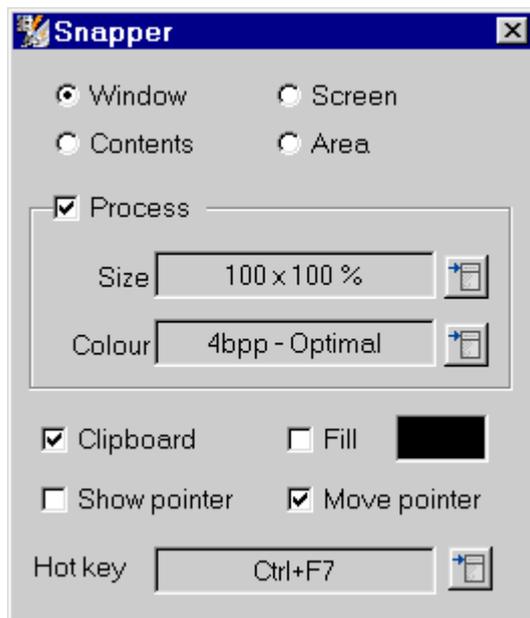
Extensions: eps.



11. Snapper

Snapper lets you capture images from the screen. To use it select Snapper from the Misc menu.

Snapper is only active whilst its window is open. When you press a given key (the hotkey) a snapshot of the relevant part of the screen is taken and passed to the DPScan window from which you opened Snapper. You might think that this means DPScan will be in your way when taking screen shots, however it can be minimised on the task bar.



The first set of options control what will be captured.

Screen

In this mode the entire screen area will be captured each time.

Window

In this mode the window under the mouse pointer will be captured. This includes the window title and scroll bars.

Contents

In this mode the contents of the window under the mouse will be saved.

Area

In this mode, a box will appear marking an area of the screen. The area of the box will be saved. The area can be adjusted by dragging it around the screen (with the sides), or dragging the corners to change its size.

Process

When this is checked the Resize and colour Format operations set up will be carried out. This is useful because often you will be snapping windows which only use a limited range of colours, Process can be set up to convert to a colour format with a limited number of colours which will use a lot less memory.

Clipboard

Makes the image be copied to the clipboard.

These two options make copying windows into another program quick. Just press the hotkey and then paste the clipboard (usually Control V).

Snapper

Show pointer

When checked the pointer will be shown on the image - if it was inside the area captured.

Move pointer

By contrast this moves the pointer before the image is captured. You might want to do this if you are capturing a menu. Moving the pointer will remove the inverted selection from the menu.

Fill

This allows you to specify the fill colour for the area window. If there is a fill, then only windows in front of the area will appear. The area window follows the usual rules for window order. Clicking on it or dragging with the left mouse button will bring it to the front. The fill colour is the background colour set on the DPScan colour palette (see chapter 7).

Hotkey

This is the key press used to trigger the screen capture. The icon to the right lets you open a small window, whilst that window is open it will display the key being pressed. Clicking on OK will close the small window and change the hotkey to the new setting.

You can use any combination of shift keys (Control, Shift, Alt and Windows) with another key.

12. Credits

DPScan was originally published by David Pilling in 1993 for the Acorn Archimedes/RISC OS ARM based desktop computer platform.

The sprites and windows were designed by John Ferguson.

This version was compiled using Microsoft Visual C++ version 6

The executables were squashed using UPX available from <http://upx.sourceforge.net>

TIFF handling uses libtiff to which the following notice applies;

Copyright (c) 1988-1997 Sam Leffler

Copyright (c) 1991-1997 Silicon Graphics, Inc.

Permission to use, copy, modify, distribute, and sell this software and its documentation for any purpose is hereby granted without fee, provided that (i) the above copyright notices and this permission notice appear in all copies of the software and related documentation, and (ii) the names of Sam Leffler and Silicon Graphics may not be used in any advertising or publicity relating to the software without the specific, prior written permission of Sam Leffler and Silicon Graphics.

THE SOFTWARE IS PROVIDED "AS-IS" AND WITHOUT WARRANTY OF ANY KIND, EXPRESS, IMPLIED OR OTHERWISE, INCLUDING WITHOUT LIMITATION, ANY WARRANTY OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. IN NO EVENT SHALL SAM LEFFLER OR SILICON GRAPHICS BE LIABLE FOR ANY SPECIAL, INCIDENTAL, INDIRECT OR CONSEQUENTIAL DAMAGES OF ANY KIND, OR ANY DAMAGES WHATSOEVER RESULTING FROM LOSS OF USE, DATA OR PROFITS, WHETHER OR NOT ADVISED OF THE POSSIBILITY OF DAMAGE, AND ON ANY THEORY OF LIABILITY, ARISING OUT OF OR IN CONNECTION WITH THE USE OR PERFORMANCE OF THIS SOFTWARE.

JPEG handling uses JPEGlib from the Independent JPEG group.

PNG handling use libpng from the PNG Development Group and the Zlib deflation algorithm library.

Roger Wilson for inspiring the exact FS code with CSA (1881/1895)