

IMAGE

A S T E R



David Pilling



Archimedes Software

© Copyright David Pilling 1993.

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.

ImageMaster, Trace, D2Font and Snapper were written using Acorn's ANSI C compiler release 4, and the Objasm assembler.

Credits and thanks to;

John Ferguson for sorting out the sprites and templates

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

The Independent JPEG Group - JPEG code

Sam Leffler - TIFF library

This manual was produced with Risc Developments Ovation DTP program.

(1st. printing 12th December 1993)

If you have any comments, suggestions, bug reports or complaints or would like further copies of this program or details of the many pieces of software available for the Archimedes, please write to;

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

From time to time enhanced versions of programs will appear. You can upgrade your copy to the latest one by sending your original disc with return postage to the above address.

Contents

1. Introduction	1
2. Installation	3
3. Scanning something	5
4. Basic operation	7
5. Edit	15
6. Filter	17
7. Colour	21
8. Configuration	29
9. File formats	31
10. Trace	35
11. D2Font	37
12. Snapper	39

1. Introduction

ImageMaster is a program for manipulating bit map images. Users of RISC OS will already be familiar with two such programs, Paint and ChangeFSI. ImageMaster provides more extensive facilities and tackles problems that these programs do not, in particular, acquiring images from scanners.

The special features of ImageMaster are;

- The ability to handle images bigger than memory, by using disc space. So called "virtual memory".
- ImageMaster can work with 1, 2, 4, 8, 16, 24 and 32 bits per pixel images. Many programs are limited to just one value.
- Images may be manipulated even if the display hardware does not support them. For example 24 bit pictures can be loaded, manipulated and resaved as 24 bit images. Real time dithering is used to display the image in screen modes with smaller numbers of bits per pixel.
- TWAIN. ImageMaster is the first RISC OS program to support this general interface to scanners. This means you are more likely to be able to use it with any given scanner.

Fundamentals

Images, or pictures come in two basic formats - bitmap or vector. A bit map image consists of a huge number of coloured points. Often 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 RISC OS, the typical bit map file is a Sprite, and the typical vector file a Draw file. Mostly ImageMaster is concerned with bitmaps, however one tool - Trace is provided which can convert Sprites into Draw files.

Given a bit map image 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 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 be larger. The typical displays on Acorn computers, have 90dpi horizontally, and 45 (or 90 dpi in 'multisynch' modes) vertically. Normally there are 4bpp (16 colours) or 8bpp (256 colours). 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 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 or 480 standard floppy discs. This is much larger than the hard discs on most computers, and even if your hard disc was big enough, the image would be difficult to handle. The lesson here is to use resolution sparingly. Often using lower resolution will not give visibly worse results. In fact most scanners physically scan at 400 dpi, the higher resolutions are synthesised by software - and can not be as good as a physical scan at the higher resolution.


2. Installation

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

Program Installation

ImageMaster is supplied on a single floppy disc. You should make a backup copy of this for day to day use. A text file called *!ReadMe* is supplied on the disc and notes new features and additions to the program. ImageMaster can be installed on your hard disc or another floppy by dragging the !ImgMaster application to the desired destination. ImageMaster is run by double clicking on the !ImgMaster icon, it will then install on the icon bar.

To use ImageMaster, you need to have version 3.75 (or greater) of the shared C library in the modules subdirectory of your !System directory (usually on Applications disc 1) or to have RISC OS 3.10 fitted. If this is not present, the program will fail to run and give an error message. A copy of CLib version 3.75 is included on the ImageMaster disc. To update your System resources (including the CLib module), you should insert the disc, open a window on it and double click the Sysmerge application and follow any instructions that are issued. This will ensure that both ImageMaster and other new programs will run correctly.



Once ImageMaster is installed on the icon bar, you can press the menu (middle) mouse button over the ImageMaster icon bar icon and get the ImageMaster menu.

The first entry leads to a program information box. This will show you which version of the program you have. Please quote this in any correspondence.

The next entry 'Choices' opens a window which allows you to set up many of the parameters of ImageMaster.

The final entry 'Quit' terminates ImageMaster and removes it from the icon bar. Clicking on the icon bar icon with the Select button (left) will open an ImageMaster window. This 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.

Clicking on the icon bar icon with the Adjust button (right) will open the configuration window (like 'Choices'). This allows you to set up the way ImageMaster works. Further details can be found in chapters 8 and 9.

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 'tools pane' and 'info palette'. Different tools can be chosen by clicking on them. When a tool is chosen the info palette display will change.

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 mouse select button.

Depending on the tool selected, the image may be manipulated by clicking and dragging on it in the window.

Clicking with the menu button over the main window, will open the main ImageMaster menu from which all the features of ImageMaster can be controlled.

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, 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.

ImageMaster controls scanners via a software interface called TWAIN. Before trying to scan anything, you must have installed TWAIN, and it must be visible to the ImageMaster program. Just like !Fonts or !System, that means the !TWAIN directory, must have had a desktop window opened on it before you attempt to use it.

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 ImageMaster by double clicking on its icon. It will install on the icon bar.

Next click on the icon bar icon with the left hand Select button on the mouse. A window will open.

Press the middle or menu button over the window. The main ImageMaster menu will appear. Follow the first entry 'File' to the right, and choose 'Select Source...'. This will open the TWAIN manager source selection dialogue box. Click on the name of the source for your scanner, and then OK.

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

The easiest thing to do, is to click on the 'Preview' button in this window. Preview will do a fast low resolution scan of what is in the scanner.

You should see an indication that the scanner is working, and hear it start to move.

Eventually an image will appear in the window.

Now choose the bottom left tool in the tools pane at the left side of the main window. This is the rectangular selection tool. You can use it to mark off an area of the image in the window. In general many of the tools in ImageMaster can be applied to such an area. However, here we just want to use the tool to mark the area to scan.

Look at the preview scan, and decide which bit of it is of interest to you. Then

drag on the area using the Select button. A rectangle will appear. Once you stop dragging, small 'handles' will appear on the rectangle. You can use the Adjust button to drag these handles and change the size of the rectangle. Dragging with Adjust in the middle of the rectangle will move it around.

By dragging the rectangle in various ways, position it over the area that you want to scan.

For some scanners, at this point you will have to re-insert the medium being scanned.

Now again, choose 'Acquire...' from the File menu. This time when the TWAIN source window appears, choose the colour and resolution that you want for your final image. If you are scanning a photograph, you may choose 200dpi and 24bpp colour. For a page of text you want to copy or fax, 200dpi and black and white or 1bpp. You can easily repeat the scan to experiment with different value. Finally click on the 'Scan' button, again the scanner will start to work.

Probably after a longer delay than for the previous scan, you will see the finished image appear in the main ImageMaster window.

To save your image for future use, again get the file menu, but this time, follow the first entry 'Save' on it to the right. A window will appear in this there is a row of icons that allow you to choose the type of file. For the moment, select Sprite. By dragging the Sprite icon at the top of this save box to a disc window you can save the image.

This saving method is the same as for most RISC OS programs. As usual you can also save directly to other applications. Not all file types are suitable for such direct saves.

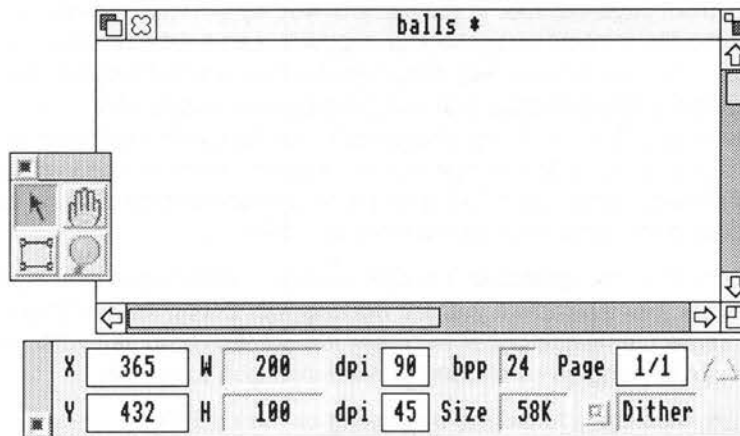
There are some keyboard short cuts to make the above process simpler. Shift F10 will open the TWAIN source selector, and if you only have one source installed, it is not necessary to go through this step. F10 will open the source window. Finally F3 will open the save box.

Much of the rest of this manual is concerned with the things ImageMaster 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

When ImageMaster is installed on the icon bar, there are several ways of making it do something. You can double click on a file type that ImageMaster is configured to load (see chapter 8), you can drop a file that ImageMaster understands on the icon bar icon. You can click on the icon bar icon. Clicking with the middle or menu button will display the icon bar menu. This gives you a chance to quit ImageMaster. You can also use it to access the ImageMaster configuration window (Choices...), this can also be obtained by clicking on the icon bar icon with the right or adjust button. Finally, clicking with select will open an empty ImageMaster window.

In most bit map programs, there is no such thing as an empty window, however because ImageMaster 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 ImageMaster 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 select button. There is one extra job for this tool. 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).

X	365	W	200	dpi	90	bpp	24	Page	1/1	/ \
Y	432	H	100	dpi	45	Size	58K	<input type="checkbox"/>	Dither	

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).

Page is the current page number in the file, followed by the total number of pages in it. Some of the file formats ImageMaster 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 ImageMaster is designed for are best dealt with one image to a file. ImageMaster 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 'page' or sometimes 'frame'. By clicking on the arrows at the side of Page you can step back and forwards through the images in the file.

Dither. This controls if ImageMaster 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.

Finally you can close the info palette by clicking on the small icon at the far left of the info palette. The info palette can be moved by dragging in the area around this icon - 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 clicking on them with the Select button, 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. The page number, so you can quickly jump to a frame. 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. The values typed in will take effect when Return is pressed.

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 select.



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 adjust.



The entire area can be moved by clicking and dragging inside it with adjust.

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


X	<input type="text"/>	W	<input type="text"/>
Y	<input type="text"/>	H	<input type="text"/>

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

The Magnifier



This tool allows you to change the scale of the display, like the usual Zoom facility in other applications. Clicking and dragging with select 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.


The Main Menu

Clicking over the work area window with the menu button, displays the main ImageMaster menu.




The general operation of the program is controlled by the File and View sub menus. As you'd guess the File menu is concerned with files, saving them, and settings for the entire file. Meanwhile the View menu control 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 leads to the usual type of Save box. You can gain quick access to this by pressing F3 (function key 3) when the caret is positioned in the main window or in the info palette.




All the file types that can be used for saving the image will be shown. By clicking on the appropriate button, you can change this. The button initially selected will be of the same type as the original file.

As usual saving can be accomplished by dragging the icon at the top to a disc window, or to another application.

Alternatively you can type in a file name and path and type Return or click on OK.


Function keys F1, F2 etc. allow the corresponding file type to be selected from the keyboard.



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

Print

This leads to a window allowing you to print images.



Copies is the number of copies of the image to be printed.

Portrait and Landscape control the orientation of the image on the paper.

Scale X and Scale Y can be used to set the scale the image will be printed at.

Corner X and Y are the offset of the bottom left hand corner of the ImageMaster work space relative to the bottom left hand corner of the printer page. Normally you will not need to use these values, because the image can be

more easily moved around the page by dragging with the pointer tool.

The Print icon will cause printing to start. To successfully print, you must have a printer driver installed on the icon bar. The name of the current printer will be displayed at the top of the window. ImageMaster will attempt to get the best results from your printer. Sometimes both the image and the printer will have much higher resolutions than the screen.

The Set icon will save the values in the window and close it without printing. It is possible to show the area that is available to print in the main window (see the View menu below). When you click on Set this area will be updated. So you can adjust these values, click on Set and then position the image relative to the printer area. Usually you cannot print the entire page area - printers cannot work right up to the page edge.

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.

Page Size

Page size	
<input checked="" type="checkbox"/> Portrait	
<input type="checkbox"/> Landscape	

A0	
A1	
A2	
A3	
<input checked="" type="checkbox"/> A4	
A5	
A6	
A7	
A8	

The final entry on the file menu allows you to set the size of the work space window. 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. This menu may be familiar from the Draw program.

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. This type of behaviour should be familiar from other RISC OS programs.

More than one image window can be open at a time, just by repeatedly clicking on the icon bar, or double clicking files.

The View Menu

View	
New view	
Print margins	^PRINT
Rulers	^R
✓ Toolbox	^F1
✓ Info palette	^P
Units	⇨
Next page	PgUp
Prev page	PgDown

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
mm
in
pt
✓ pix

Units points to a small sub menu that allows you to set the units used throughout ImageMaster. 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 ImageMaster menu. It controls operations altering the physical aspects (e.g. size) of the current bitmap.




Flip Horiz

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 Vert

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




Rotate

This allows the bit map to be rotated. Positive angles are anticlockwise. 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 at the top, 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.




When a bit map is rotated by an angle other than 90 or 180 degrees, it will cover a larger area - bit maps 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.



Shear

This pushes the bit map over sideways. In other words it converts a rectangular shape into a parallelogram.

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



Scale

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

Scale			
<input checked="" type="checkbox"/> Resize by	X	<input type="text" value="100"/>	%
	Y	<input type="text" value="100"/>	%
<input checked="" type="checkbox"/> Resize to	W	<input type="text" value="200"/>	<input type="text" value="200"/>
	H	<input type="text" value="100"/>	<input type="text" value="100"/>
<input checked="" type="checkbox"/> Resample	X	<input type="text" value="90"/>	<input type="text" value="90"/>
	Y	<input type="text" value="45"/>	<input type="text" value="45"/>
	Memory	<input type="text" value="58K"/>	<input type="text" value="58K"/>
<input type="checkbox"/> Lock aspect			
<input type="button" value="OK"/>			

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 increase - the 'Resize by' option. For example 200 will make the image twice as big.

Secondly you can specify the new dimensions of the image - 'Resize to'

Finally you can give the image a new resolution - 'Resample'.

To use any of these options you first need to select the appropriate icon, and then type in the values you desire. Clicking on OK will start the scaling process. The non-writable 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 bit map 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.

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 ImageMaster. The filters are accessed from the fourth entry on the main ImageMaster 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 16 intensity levels or bands.



The Red, Green and Blue buttons can be used to choose which components are displayed. These buttons operate in a slightly different way to usual. Clicking with Select will select just one of them, more than one can be selected by clicking with Adjust. The same principle usually applies to multiple selections e.g. files.

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 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.

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.

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 !ImageMaster.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 final entry on the main ImageMaster 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

Swop
R=>G=>B
R<=>G
G<=>B
B<=>R

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

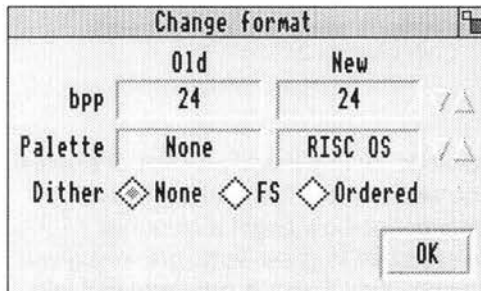
This converts a colour image to black and white, or in general a series of shades of grey. The weights used to do this can be set from the General configuration dialogue box. See chapter 8. These weights define how bright a colour is perceived to be. Because the eye is more sensitive to green light, green is given a greater weight. If a non-paletted 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 ImageMaster, 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 ImageMaster. 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.

The display on most Acorn computers can only display 8bpp yet images off other computers are often 24bpp, or perhaps have colours in their palettes that do not correspond to those in the standard desktop palette. The familiar program ChangeFSI manages to display 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.



By toggling the Change Format window to full size (click on the top right hand corner) you can access some extra options.

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 ImageMaster 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.

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.

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 ImageMaster palette window, and

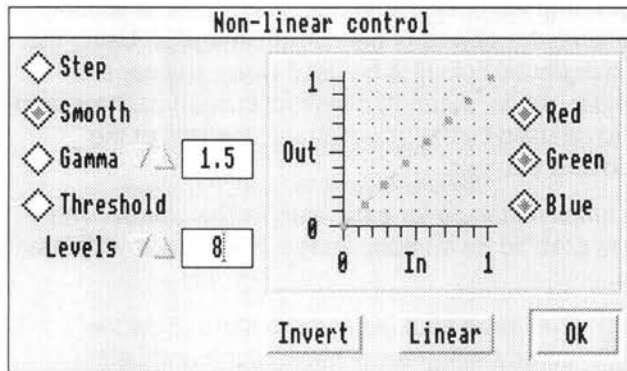
described in a later section of this chapter. The concept of brightness should be intuitively obvious. The brightness slider makes all colours brighter or darker. The contrast slider has a similar effect, except brighter colours have their brightness level increased more than dimmer colours. Note that 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 swopping 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 ImageMaster 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 colour and tint to be used by the selections on the rest of the Colour menu - Tint, Filter, Add and Subtract.

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.



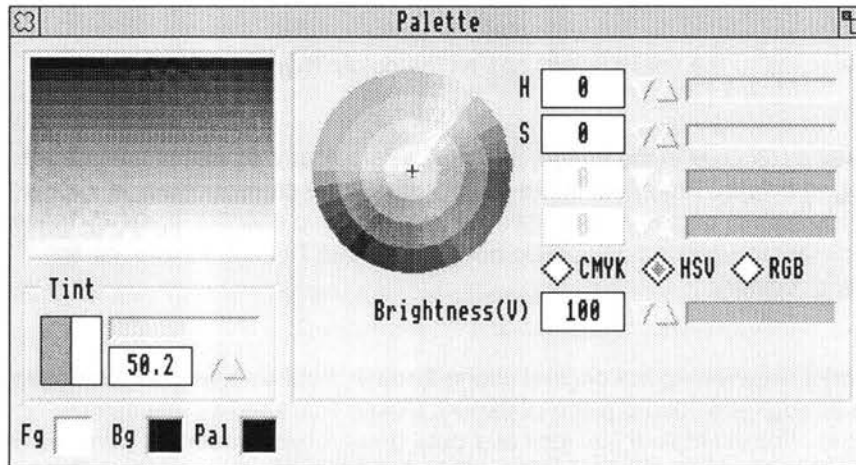
The three squares at the bottom show the current foreground, background and palette colours.

Clicking with select on these icons transfers the colour to the colour selector (see later). Clicking with adjust transfers the colour selector to the colour.

If you click on the palette area with select that colour will become the one shown on the palette icon and the colour selector. Clicking in the palette area with adjust will transfer the current colour selector colour to the palette. A similar effect can be achieved by clicking on the Pal icon with adjust.

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 central part of the palette window allows this strength or tint to be set. The icon under tint, shows the tinted and the original colour.

To display the colour selector, toggle the palette window to full size by clicking on the top right hand corner. The window will display this in addition.



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 point 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.

Click on a colour icon with select e.g. the Fg, Bg, Pal or the palette area icon.

Modify the colour using the colour selector.

Set a tint for the colour.

Finally store the colour by clicking on one of the colour icons with adjust.

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

We now come to the final 4 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.

8. Configuration

The configuration dialogue box can be accessed either by clicking on the icon bar icon with Adjust, or by selecting 'Choices' from the icon bar icon menu. Either approach will open a window. 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 ImageMaster 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 ImageMaster so that it is the same when the program is next used.

Memory

This controls ImageMasters 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, typically 4Mb. or 8Mb. 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 temporary file should be kept on a large area of fast writable disc. By default it is created inside the directory !Scrap. It is not a good idea to put the temporary file on a RAM disc. Doing so defeats the object, and ImageMaster can make use of any spare RAM more efficiently itself, rather than via a RAM disc.

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. There is no point having 4Mb of RAM, 2Mb of free disc space, and then expecting to be able to manipulate 10Mb images.

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.

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.

The final section allows you to specify what will happen whilst ImageMaster is busy working on a long process. The alternatives are to display the hourglass with a % sign, to multitask (which means you can carry on doing other things) or finally to blank the screen. Blanking the screen is useful because in screen modes with lots of colours and pixels the older Acorn computers run more slowly. Multitasking will be the slowest option, and it is not always possible to multitask for example when transferring files to other applications. It is possible to arrange for an audible 'beep' to be made at the start and end of each multi tasking process to give some sort of feedback.

Which of these options you choose will depend on how you work. For slow machines the multi tasking option may be useful. For large screen modes blanking will be most suitable. For a fast machine in a small mode, the hourglass is probably best.

9. File Formats

In an ideal world, one might expect there would be just one type of file used for storing bit map 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 bit maps, 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 option 'Double Click' this controls if the the file type will be loaded by ImageMaster when double clicked. You may well choose to have sprites not to be double click loaded. In general only after ImageMaster has been rerun after a machine reset will these double click values apply in every single case.

Sprite

This is the standard RISC OS bit map file type, as used by 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 ImageMaster will supply a 256 entry palette. This may upset some old software.

In effect no resolution data is saved with sprites.

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.

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.).

No resolution data is saved.

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.

BMP

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

No resolution data is saved.

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.

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.

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 Bit Map) 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.

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.

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 (click on a selected one with adjust) and the colour data compression will then be used for all images.

10. Trace

As previously described, pictures are stored in one of two formats - bitmap and vector files. Under RISC OS these are typically Sprites and Draw files. There are advantages and disadvantages to both types of file.

Sprite files use quite a lot of memory, and can't be easily scaled. If you attempt to make sprites larger they will develop jagged edges. Similarly, such 'jaggies' will show up when the sprite is printed on a high resolution printer. Trace, allows you to convert Sprites into Draw files and solve some of these problems.

Getting Started

Double click on the !Trace application icon, after a pause while the program loads and initialises, an icon will appear on the icon bar. Clicking with the mouse menu button over this, displays a small menu, this provides an entry to quit the program, and a dialogue box showing the version number of the program.

Clicking on the icon with the select button will open the main Trace windows.

Initially, these will be blank. Clicking the menu button over them, will display the main Trace menu. The two windows, show the Sprite file being traced and the resulting Draw file. A new sprite file can be loaded by dragging it into either of the windows, or onto the icon bar icon.

Main Menu Options

Save Draw

Moving the pointer to the right of this, will display a standard save box. Dragging the Draw file icon allows the Draw file to be saved to disc or direct into other programs. If no Draw file has been produced, this option will be shaded out.

Save Sprite

This allows the sprite image to be saved. It is useful if you have imported a sprite from Paint after making a change, and then decide you would like to keep the sprite.

Trace

This tells the program to trace the Sprite file currently loaded. Tracing can take a long time, whilst it is going on, the pointer will change to an hourglass. Once tracing is complete, the finished Draw file will appear in the Draw window. You can then use the Save entry to save the Draw file in the usual way

Zoom

Leads to a standard RISC OS magnifier that can be used to scale the images.

Choices

This leads to a sub menu with two entries;

Error

This allows you to enter the average error for each pixel traced. In other words how closely the Draw file corresponds to the Sprite file. Large values will produce Draw files that poorly correspond to the Sprite file. Small values will produce Draw files that are more accurate. However, if this number is too small, the Draw file will start to reproduce the individual pixels of the Sprite (just the opposite of what we are trying to do). Draw files will become larger as the error is decreased. In general, error has to be adjusted to get the best results for a given image.

Save

This allows the configurable options (Error) to be saved for future use. After selecting this entry, the next time you use Trace, it will start up with the current values.

Info

This displays information about the Draw and Sprite images.

Batch Mode

Because tracing can take some time, it is often useful to be able to go away and leave the program tracing a collection of sprites. To do this, assemble your sprites in a directory, then drop the directory on the icon bar icon. Trace will create another directory called 'Draw' in the same place as your original directory, and then commence to scan through your sprites, tracing each in turn and putting the resulting Draw files in the 'Draw' directory.

Escape

Tracing can be stopped by pressing the Escape key.

The Finished Product

After successfully tracing your sprite, it will often be best to load the Draw file into Draw. Many improvements can be made with Draw. Although a lot of Sprites will Trace well, some will give poor results, and Draw can often save these files.

The one type of Sprite which will give bad results, are those in which colours have been represented by dithering (lots of small dots of different colours). ImageMaster can be used to improve sprites considerably for tracing.

Trace will not work with non-palette sprites i.e. those with 16, 24 or 32 bpp. Such images do not have distinct enough edges to trace. So before importing sprites into Trace you should convert them to 8bpp. Naturally you should not do this conversion using FS or ordered dithering.

The more complex the sprite is, the larger the Draw file, and the longer Tracing will take. To start with try a simple Sprite and then as you get a feel for the program, build up to more complex images.

11. D2Font

This is a program that will convert Draw files into Outline fonts. One of the things you may trace, is text. Trace will convert these into Draw files, but you can't use Draw files as fonts in the common DTP programs. To use them, you need to turn the Draw files into fonts, this is what D2Font does.

To get started, double click on the D2Font icon. D2Font will appear on the icon bar. Clicking on the icon bar icon, will open a window consisting of a large number of squares, one for each character in the font. You now have to use Draw, to save each character of the font, dropping the save icon into the appropriate square.

Because filling the grid, may take some time, you can save the grid and all the Draw files in it. Pressing the menu button over the main window, will provide a menu with two save options. The first "save D2Font" is used to save the entire grid. Grids are saved in a special D2Font file. This can be reloaded by double clicking or dragging to the icon bar.




When you have finished filling the grid, the second save option saves the font. This presents a directory icon, which you must drag into your current !Fonts directory. The font will then appear if you double click !Fonts or next time you reset the computer. The directory, contains 'Outlines' and 'IntMetrics' files. It is important that the font is created with the correct name, unlike a normal file, you can't rename a font directory to change the font name.

Other options on the main menu, allow you to clear all the grid, or just one character.



If you double click on one of the character squares, a large version of the character will appear. There will be a red line under the character. You can use the Select and Adjust buttons respectively to drag the line up and down, and to make it bigger and smaller. The red line represents the baseline and width of the font. Normally, you need to make the width of the character slightly larger than the character outline. For letters like lower case 'g', the baseline needs to be half way up the character.

It is also possible to manipulate the base line for characters which don't have a Draw file loaded. The reason why this is useful, is because certain characters,



like Space, although they don't print anything, must have a width. So for such characters, you can use this feature to give the correct behaviour.

Because changing the base lines for every character, can be tedious. Two menu entries are provided.

'Mono space' makes the base line width for every character, the given fraction of the maximum character width. So entering 120% would make the base line for every character 120% of the width of the widest one.

'Prop space' make the base line width for each character, the given fraction of the characters width. For example, entering 120% (not a bad choice for most fonts) would mean that each character had a border of 10% of its width on each side.

Finally, a drag on a grid square using select, will allow you to save the contents as a Draw file to another application.

The lines which form font characters, have alternating black and white fills. This can lead to a small snag. Most of the things you trace, will have a white background. Trace will turn this into a large white rectangle at the back of the Draw file. If you import this into D2Font, you will produce a font in which black is replaced by white and vice a versa. The solution is to use Draw to select the bits that you want to import (click with adjust) and then the Save Selection entry to save them into D2Font.


Similarly since fonts are formed from filled curves you cannot use Draw files that only consist of lines, and not filled areas. Fortunately Trace always produces Draw files of the correct type.

There is a freely available font editor, written by Acorn, this is very useful for examining the fonts produced and for fine tuning them. Copies are available from Acorn support BBS's (e.g. Arcade 081-654-2212) and from many other sources in return for a blank disc and postage.

12. Snapper

Snapper is a screen capture program. It allows you to save an area of the screen as a sprite. Paint already has such a 'snap shot' facility. The reason for having another program to do this job, is that some of the most common tasks like saving the contents of a window or menu, are messy using Paint which requires the use of an imaginary rectangle and a time delay.

Snapper is run in the usual way by double clicking on its icon. It will install on the icon bar. The icon bar menu offers the usual Info and Quit options. In addition 'Control' will open the Snapper control window, and 'Save choices' will save the current Snapper configuration. The control window can more easily be opened by clicking on the icon bar icon with Select. The option to save the configuration is placed on this menu, because it will also save the position and size of the snapper control window.



The default control window, shows 4 buttons that can be used to select the mode the program will work in.

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.

Finally the 'Snap' button triggers the screen capture. When this happens, a standard save box will appear on screen. You can drag the icon from this in the normal way, or click outside the save box to get rid of it.

Obviously the Snap button is not always of use. For the Window and Contents modes, you would always end up with a picture of the Snapper control window. To get around this, it is also possible to trigger screen capture with a key press. By default this keypress is to hold down the left hand Ctrl and Alt keys.



toggling the snapper window to full size (click on the top right corner icon) reveals more options.

Key trigger

This allows you to set which key presses will start the screen capture.

Pointer

There are two options. Move will remove the pointer from the screen area before capturing the screen. This is useful if you want to capture a menu. With move selected, the menu will appear without the usual inverted selection where the pointer is. Show will display the pointer on any captured images.

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. Dragging it around with adjust will preserve its position in the window stack. Dragging with select will bring it to the front. The fill colour can be chosen by clicking on the right arrow and then making a selection from the menu.

Save Path

So far only screen captures from the desktop have been described. If you have a program that takes over the screen, none of the above facilities will be any use. By choosing Screen mode, and selecting the icon at the right of Save path, screen captures can also be achieved outside the desktop. The writable icon contains the path that will be used to save the screen. Clicking on the right pointing arrow will open a save box. This can be dragged to where you want the screens to be saved. The default setting is the same directory as that holding Snapper.

This full screen save from outside the desktop is implemented by the Snapper module. This is contained in the module Spell (it shares some code). The module supports one *command Snapper_Key which is followed by a number formed from one bit for each of the Left and right Alt, Ctrl and Shift keys. So Left Alt contributes 1, right Alt 2, left Ctrl 4, and finally right Shift 32. Files will be saved to the system variable <SnapperFile\$Dir>. This information should allow the use of the module outside of Snapper.