

Digital Media Tools

3rd edition

Nigel Chapman and Jenny Chapman © 2007

Published by John Wiley & Sons, Ltd.

The material in this sampler provides an indication of the content, style and teaching and learning features in the fully revised and re-illustrated 3rd edition of Digital Media Tools. We have included short excerpts from each chapter.

You can find a range of additional support materials, the full table of contents and preface from the book at the book's Web site, www.digitalmediatools.org. Instructors who are affiliated to a recognized educational institution can request an evaluation copy of the book from our publishers, John Wiley & Sons, or by using the evaluation request form under **Contact Us** on the support site.

This PDF document contains sample material taken from Digital Media Tools 3rd edition, by Nigel and Jenny Chapman.

By downloading this PDF you are agreeing to use this copyright electronic document for your own private use only. This includes use by instructors at educational institutions in connection with the preparation of their courses, but it does not include reproduction or distribution in any form, for which explicit permission from the Authors is required in every case.

Copyright © 2007 Nigel Chapman and Jenny Chapman

Digital Media Tools 3rd edition is published by:

**John Wiley & Sons Ltd.,
The Atrium,
Southern Gate,
Chichester
PO19 8SQ**

All Rights Reserved. No part of this publication may be reproduced, stored in a retrieval system or transmitted or distributed in any form or by any means, without permission in writing from the Authors, except under the terms of the Copyright, Designs and Patents Act 1988 or under the terms of a licence issued by the Copyright Licensing Agency Ltd,

The Publisher may be contacted at cs-books@wiley.co.uk or via their Web site www.wiley.com

The Authors may be contacted via their Web site www.macavon.com or via the book's support site www.digitalmediatools.org

2

Interface Basics

In This Chapter

- Working on Windows and Mac OS X platforms
- File, Edit and Help menus and commands
- Panel groups, docks and workspaces
- Panel controls
- Tools panel, Options bar and Control panel
- Layers, Navigator and History panels
- Automating repetitive tasks
- Rulers, grids and guides

Panel Groups and Docks

Panels are usually combined into *panel groups*. For instance, Figure 2.9 shows a panel group containing three panels. The group as a whole can be treated more or less like an ordinary window. You can drag the grey area surrounding the tabs to move it around the screen, minimize it (as shown at the bottom of Figure 2.9) or close it using the – and x controls in the top right-hand corner, or resize it by dragging out its bottom right corner. Note that these are not the standard system window controls, but are specific to Adobe’s user interface. The individual panels are revealed by clicking on the tabs along the top of the window to bring them to the front of the group.




Figure 2.9 A panel group

The arrangement of panels into groups is not fixed. You can customize your workspace in any of the applications in this book by dragging panels out of their groups by their tabs. If you drag a panel so that it is over the grey area at the top of a different group, a blue outline appears around the new group; if you then drop the panel it joins the group. If you drop the panel in an empty space away from any group, it appears in a group of its own. In this way you can organize your panels so that their grouping reflects the way you like to work with them.

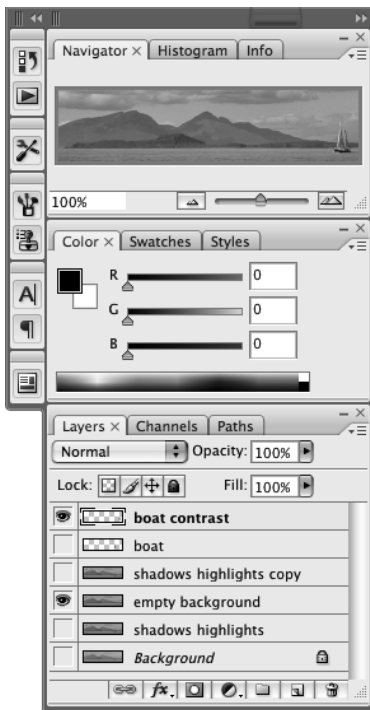


Figure 2.10 Panel groups in docks

You can close individual panels that you are not using by clicking the x control in the panel’s tab; you can also control their visibility using the **Window** menu, which is available in all the applications in this book. This menu has an entry for each panel and for any open document windows. If a panel is currently visible, there will be a tick next to it in the **Window** menu. Selecting a ticked entry closes the panel and removes the tick. Selecting a panel without a tick opens it. However, you should note that if you close or open a panel using a menu command, its entire group will close or open, contrary to what you might expect. Selecting any currently closed panel from the **Window** menu will therefore open the whole panel group that it was last a part of.

In Photoshop and Illustrator you can hide all the open panels at once – including the **Tools** panel – by pressing the **tab** key, and then bring them all back by pressing **tab** again. Pressing **shift-tab** hides (or reveals) all open panels except the **Tools** and **Options/Control** bar; this can be very useful if you need some extra screen space but also need to retain the tools and controls for working on a document.

Panel groups are further organized into *docks* and *stacks*.

A dock is a set of panel groups that are arranged vertically; docks themselves are arranged next to each other, usually with the rightmost fixed at the right-hand edge of your screen. Any other docks are placed up against their neighbour. If you refer back to Figure 2.2 you will see two docks to the right of the main image; Figure 2.10 shows them in greater detail. The dock on the right consists of three panel groups. The dock to its left only shows as a set of icons. This is because it has been *collapsed*, which simply means that it has been reduced to this minimal form to save screen space. In fact, there is an intermediate collapsed form, shown in Figure 2.11, where the palettes in the dock are reduced to icons with their names beside them. For beginners, this collapsed form may be the best option for docks, since it makes finding a particular panel easy, while not taking up much space. When you are familiar with the panels' icons, you may prefer to collapse right down to just the icons, freeing up a little more screen space. In either of the collapsed forms, broad horizontal dividers separate groups within the dock, as, for instance, above the Brushes panel in Figure 2.11.



Figure 2.11 Docks as icons and names

But how do you collapse a dock? If it is in its expanded form, like the right-hand dock in Figure 2.10, there will be a pair of arrowheads pointing right in the top right corner. Clicking on these will collapse the dock to icons and names. To collapse it fully, drag the pair of vertical lines in the top left corner. The names' area will shrink as you do so, the names becoming abbreviated, until there is no room for even the first letter, at which point the dock will snap down to a set of icons. (This may sound awkward, but is straightforward in practice.)




Figure 2.12 Temporarily expanding a collapsed panel

If a dock is collapsed, you reveal a panel by clicking on its icon. It will fly out from the dock, as shown in Figure 2.12. To put it back, click on the icon again, or click on any other panel's icon. In a collapsed dock, the arrows in the top right corner point to the left, and clicking on them expands the whole dock.

In all three programs that use the standard panels there is a checkbox somewhere in the preference settings labelled Auto-Collapse Icon Panels (Auto-Collapse Icon Palettes in Photoshop). If you tick this checkbox, panels will always revert to their collapsed form as soon as you click away from them, for example, when you select a tool or place a cursor on your document.




Figure 2.13 Creating a new dock




Figure 2.14 *Adding a panel group to a dock*




Figure 2.15 *Adding a panel group to a stack*




Figure 2.16 *A minimized stack*

Reorganizing entire docks is conceptually simple but may take practice. If you drag an undocked panel group or a single panel to one edge of the screen or to the edge of another dock, you will see a bright blue vertical line appear down that edge, as in Figure 2.13. If you now drop the panel group, it will become a new dock next to the existing one. Note, though, that the left edge of the screen is usually occupied by the Tools panel, described later. Placing a dock on the left of the screen will push the Tools panel out of the way, but it is possible to add a dock to its right, even though the Tools panel is not a dock itself and cannot be grouped with other panels.

If, instead of dragging a panel group to the edge of a dock, you drag it right over an existing dock, a blue outline will appear around the panel group your cursor is over, as in Figure 2.14, and dropping the panel group at that point will add all its panels to the outlined group. If you drop the group between groups in the dock, or at the top or bottom of the dock, it will be added as a group. You can remove groups or individual panels from a dock just by dragging them out.

The final possibility for organizing panels is as a stack, which, like a dock, consists of several panel groups arranged vertically, but it is not attached to the edge of the screen and, importantly, it cannot be collapsed. To add a panel group to an existing stack, drag it to the bottom of the stack. When a horizontal blue line appears at the bottom, as it does under the Histogram panel in Figure 2.15, drop the panel group. You can make a new stack from two panel groups. Stacks behave as if the panel groups in them are glued together, so you can drag the whole stack about the screen as a unit and close all the panels in it at once using the stack's x control. If you minimize the stack, all the panel groups are minimized, as shown in Figure 2.16.

TRY THIS

Practise using and organizing panels. In any application, bring each panel in a group to the front. Drag one panel out of a group and drop it into a different group. See what happens if you drag a panel out of a group and then back in.

Hide and reveal individual panels and note what happens to the groups. Hide all the panels, noting which interface elements remain.

Collapse the dock to icons and names, then just to icons. Practise revealing individual panels that are in the collapsed dock.

Remove several panels from the dock and combine them into a stack. Practise minimizing the stack.

Workspaces

The facilities just described give you considerable freedom in how you arrange panels. Some people will be happy to use *ad hoc* panel arrangements and change them from day to day, but others may prefer to be more systematic. A **workspace** is a named arrangement of panels. The names of all available workspaces appear on the **Window>Workspace** sub-menu. If you select a name from this sub-menu, the panels will move to the positions that constitute the corresponding workspace. (Even Dreamweaver has a similar facility in the **Window>Workspace Layout** menu.)

Each program provides one or more pre-defined workspaces. You can define your own and save it for future use, simply by arranging the panels in a way you like, and selecting **Save Workspace...** from the **Window>Workspace** sub-menu. (**Save Current...** in Flash or Dreamweaver.)

It would make sense to save custom workspaces if you share your computer with other people and want to make sure that programs are set up the way you like them. You may also find it useful to have custom workspaces for different types of job – this is the rationale behind the pre-defined workspaces in Photoshop, for example. For many of us, though, the default workspaces are the most useful, because selecting one of these immediately tidies up any mess made by moving panels around in an undisciplined manner. (A workspace in Photoshop includes keyboard shortcuts and menus, so changing workspaces may change these as well as the configuration of the panels. A warning message appears when this happens.)

TRY THIS

Tidy up after the previous exercise by choosing the default workspace.

Switch between the pre-defined workspaces in each application. See if you can work out the logic behind each arrangement.

Make your own arrangement of the panels in one program and save this workspace. Switch to a default and then back again to your own workspace.

Panel Controls

In the top right-hand corner of every expanded panel you will find a small symbol consisting of three horizontal line with a downward-pointing triangle to their left. Clicking on this causes a pop-up menu, known as the **panel menu** to appear, as in Figure 2.17. Here you can find commands that are related to the functions performed by the controls on the panel. Sometimes, these are duplicates of commands found on the main menu bar, or controls on the panel itself. In other cases, the panel menu provides the only way of carrying out some operations. The function of the panel menu is to group together all the relevant operations in one place. Because the panel menu is attached to the panel itself, functions and options relevant to the panel can be accessed with a small movement of the cursor, which is quicker, more convenient and less likely to contribute to repetitive strain injuries than having to use the main menu bar.



3

Bridge

In This Chapter

- **Browsing media files in different workspaces**
- **Viewing slide shows of images**
- **Organizing your files and folders**
- **Using other programs' tools from inside Bridge**
- **Adding your own ratings, labels and keywords to files**
- **Adding standard metadata**
- **Retrieving and sorting files**
- **Creating collections**

Rotating

When taking photographs, you will sometimes turn the camera on its side to take a picture in portrait orientation. When Bridge displays such a picture, it shows it on its side, in landscape orientation. You can rotate images in JPEG, TIFF and Photoshop formats clockwise or anti-clockwise, in 90° increments, using the buttons at the top right of the browser window, or the rotate commands on the **Edit** menu. Rotating an image only affects the way thumbnails and previews are displayed in Bridge, it doesn't actually change the data in the image file. It does, however, cause the image to be displayed in its rotated state if you open it in Photoshop.



Slide Shows

Instead of looking at all the images in a folder at once, you can choose **Slideshow** from the **View** menu. This causes the images in the current folder to be displayed as a slide show, which takes over the entire screen. Everything else is hidden, including the menu bar. Various keys can be used as slide show controls. For basic operation, press the space bar to start and pause the show, and press the escape key to stop the show and return to the browser window.

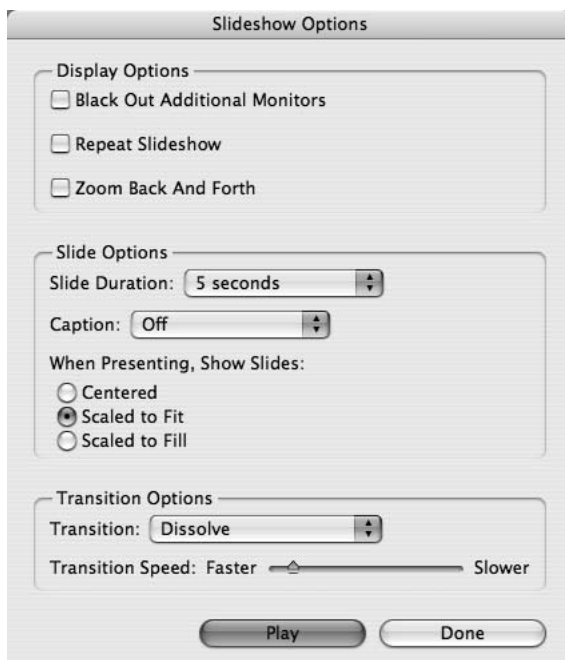


Figure 3.8 Slide show options

The slide show can be displayed in several different ways. Pressing **L** while it is playing, or choosing **Slideshow Options...** from the **View** menu, brings up the dialogue shown in Figure 3.8, which allows you to change some aspects of the display. There are three possibilities for the way in which the image is fitted to the screen: the image can be centred, with a border on all sides, it can be scaled to fill as much of the screen or window as possible while maintaining its aspect ratio, or it can be scaled to fill the whole screen area – this may cause some of the image to be cut off.

By default, each image is displayed for five seconds. You can choose a different duration or select **Manual** from the Slide Duration pop-up menu. In the latter case, you move on to the next slide by pressing the space bar or right-arrow key. By default, the slide shows stops after the last image has been displayed. If the **Repeat Slideshow**

checkbox is ticked, it will go back to the first and repeat indefinitely until you explicitly stop it. The transition used to move from one slide to the next can also be changed using the pop-up menu at the bottom of the Slideshow Options dialogue. The default Dissolve is probably the

easiest on the eye, while **None** is more austere. The other transitions are probably best avoided, as is the **Zoom Back And Forth** option, which seems designed to make you feel seasick. The speed of transitions is set using the **Transition Speed** slider in the **Slideshow Options** dialogue.

If the currently selected folder includes PDF files comprising more than one page, the slide show will display the pages of each file in order. Some additional navigation keys are available for such files. Pressing **h** during a slide show will show you a list of these, together with one or two additional options, which should be self-explanatory. Among these you will see some commands for simple editing and annotating from within a slide show. We will return to these later in the context of Bridge's annotation features.

TRY THIS

Run a slide show of all the images in one of your folders. Try altering the duration of each slide until you find a suitable speed for the show. Experiment with the different display options, noting any cropping that occurs when the slides fill the screen or window.

Compact Mode

A final option for displaying the browser window is Bridge's *compact mode*. In this mode, which can be selected from the **View** menu or by clicking on the icon in the top right corner of the browser window, everything except the thumbnails and basic navigation controls is hidden, as shown in Figure 3.9 (irrespective of which workspace is being used in the full browser window). The compact window can be resized by dragging out the bottom right-hand corner in the usual way. A panel menu with some useful commands is added at the top right corner.

By default, in compact mode, the Bridge window floats over any other open windows on your computer screen. This is convenient (if you have a big enough screen): it ensures that the browser is visible all the time, no matter which program you are actively working on. Photoshop and Illustrator users may find it useful to have a floating browser window to select files from, instead of relying on the normal file open dialogues. If you prefer the compact view to behave like an ordinary window, so that it is obscured when another program is at the front, deselect **Compact Window Always On Top** in the panel menu.



Figure 3.9 Compact mode

Some commands on the **View** menu still work in compact mode. For instance, the **View>Show Thumbnails Only** command has been used in Figure 3.9 to hide file information and display only the image thumbnails. A subset of the relevant **View** menu commands appears in the **View** sub-menu of the panel menu. You can also sort the displayed files, and use the slider at the bottom of the compact mode window to change the size of thumbnails. Filter settings are respected in compact mode, so only files that would be displayed in full mode are visible.

To return to the full browser view, select **Full Mode** from the **View** menu or sub-menu, or click on the icon at the top right of the compact mode window, next to the panel menu control. To toggle quickly between full and compact mode, use the keyboard shortcut **[cmd/ctl]-return**. If you are short of screen space, there is also an **ultra-compact mode**, in which the thumbnails are hidden, leaving just the controls, as shown in Figure 3.10. You can toggle between compact and ultra-compact modes using the other icon at the top right of the compact mode window, thus hiding or revealing the thumbnails with a single click. It can be convenient to leave the ultra-compact window floating at the top of your screen until you need an image, then switch to compact or full mode when you want to browse. (Note that if you return to full view and later select compact mode again, you will get the ultra-compact window if that is the view you last used in compact mode.)

Figure 3.10 *Ultra-compact mode*

Stacks

A large collection of files or folders will overflow the main content window. To help you manage large numbers of files more easily, you can group them into **stacks**. These are just virtual collections, although usually you would only group files in a stack if they had something in common.

To group a set of images or other files or folders together into a stack, select them all (they do not need to be contiguous thumbnails) and then choose **Group as Stack** from the **Stacks** menu.



Figure 3.11 *A collapsed stack*

The thumbnails will collapse, with only the first one displayed, as shown in Figure 3.11. The number in the top left corner of the stack tells you how many files it contains. If you click on this number, the stack expands so that you can see all the thumbnails in it. The files still remain stacked, though, which is indicated by an outline around them all, and the number, which now appears by the first file in the stack. Clicking the number again collapses the stack once more.

The slider and play button that you can see at the top of the stack in Figure 3.11 appear when you hover the cursor over a collapsed stack, providing there are ten or more items in it and the thumbnails are big enough to accommodate the controls. You use these controls to scrub or play through all the files in the stack – each one comes to the top in turn as you drag the slider or click the play button. The rate at which the files in the stack are played is set in the Thumbnails pane of the Preferences dialogue. However, the available values are all typical film and video playback speeds, suggesting that the play button is intended to be used for previewing sequences of video or animation frames. These speeds certainly don't give you enough time to look at individual images or other files, so if you just want to look through a stack quickly, use the slider instead.

You can add more files to a stack by dragging and dropping them onto it, or remove files by dragging them out; you can only drag individual files out if the stack is expanded, of course. To dismantle a stack, use the **Stacks>Ungroup from Stack** command. This conveniently leaves all the unstacked files selected. If you click on a collapsed stack, only the top file is selected, but you can select the entire stack by clicking on its border, allowing you to apply commands to all the files in the stack at once. In particular, you can attach keywords, labels and ratings, as we describe later.

TRY THIS

Familiarize yourself with the use of stacks by selecting several files within a folder and grouping them into a stack. Expand and collapse the stack, add some extra files and remove some others. Group a selection of image files into a stack (try an image sequence if you can) and experiment with the playback controls.

DON'T FORGET

Use Bridge to navigate and display collections of images and media files.

Choose between the default, filmstrip, metadata focus, file navigator and light table workspaces. Assign your favourite workspaces to the numbered buttons at the bottom right of the window, for quick switching between them.

Use a slide show to display selected images in sequence at a large size.

Use compact mode to keep a floating window of files available.

Group files into stacks to save space and organize them better.

4

Photoshop

In This Chapter

- Opening, importing and saving image files
- Resolution, resampling and colour modes
- Layers
- Making and using selections and masks
- Performing image adjustments and retouching images
- Applying effects and filters
- Painting and erasing pixels
- Creating animation

Moving and Transforming Layers



If you think about a layer as a sheet of acetate, then you can see that it makes sense to move layers about, each sliding across the one above or below it. By using the *move tool*, you can move any layer except the background layer. Make sure you have selected the layer you wish to move by clicking on it in the Layers panel, select the move tool from the toolbox, and then you can drag the selected layer around in the document window. You can move several layers together by using multiple selections or linking, as described in the previous section. Moving layers in this way is most likely to be useful if they contain text or separate objects with transparent backgrounds. For instance, in the example in Figure 4.10, the lighthouse layer could be moved down to reduce the amount of building showing at the bottom of the image and increase the amount of sky.

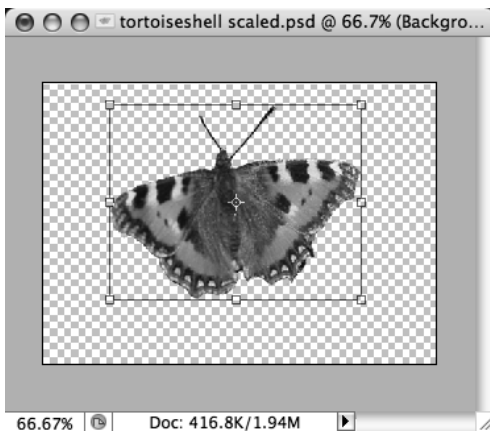


Figure 4.16 A Free Transform *bounding box*




Figure 4.17 *Layer rotation using Free Transform*

The move tool only allows you to move layers vertically and horizontally – you can't drag them round to change their orientation. The *Free Transform* command, which is on the Edit menu, is used to rotate layers and to apply other geometrical transformations. (As we will see later, you can apply transforms just to selected parts of a layer.) When you choose the *Free Transform* command, a bounding box is displayed, which surrounds the coloured pixels on the layer. Figure 4.16 shows a bounding box, which, as you can see, has small squares at each corner and in the middle of each side. These are called *handles* and are used to carry out transformations. In the middle of the bounding box you will see a symbol that looks like the cross-hairs in a gun sight. This marks the centre of the layer for the purposes of free transformation.

If you position the cursor outside the bounding box, it changes appearance to a curved two-headed arrow, which indicates that you can use it to rotate the image, which you do just by dragging round, as illustrated in Figure 4.17. This works like a lever: the further away the cursor is from the bounding box, the further you have to drag to get the same amount of rotation, so if you want precise control you should move the cursor well outside the bounding box. If you want to rotate the contents of the layer around a point other than its geometrical centre, you can drag the centre symbol to a different place.



Figure 4.18 *The Free Transform options bar*

Whenever you use the Free Transform command, two things happen. First, the Options bar changes as shown in Figure 4.18. You can apply precise transformations numerically by entering values in the boxes. Figure 4.19 shows the butterfly scaled up in size by entering a scaling factor (as a percentage) in the width (W) and height (H) boxes of the Free Transform Options bar. Second, the menu bar changes, with most commands being disabled. This is because a sequence of transformations must be finished and confirmed (or cancelled) before you can do anything else – working with free transformations is like using a modal dialogue box. When you have finished the transformation, you can press the return key or click on the tick mark (the *commit button*) at the right-hand end




Figure 4.19 *A scaled layer*

of the Options bar to confirm the operation if you have achieved the result you wanted, or press escape or click on the *cancel button* to abandon the transformation if you don't like it.

There are various ways to distort images using Free Transform, but many of these distortions are more likely to be useful in vector graphics work (and so should be carried out in Illustrator) than when working with typical bitmaps. For example, you can scale the image disproportionately by hand: pulling one of the handles in a vertical side will scale in a horizontal direction only; pulling a handle in the top or the bottom will scale vertically, and pulling one of the corner handles scales in both directions. (You can also scale proportionally by hand if you hold down the shift key while pulling a corner handle.) You can reflect the image by pulling a handle all the way through an image and out of the other side of the bounding box, but unless you get the distance exactly right the image will also be distorted. For a precise reflection of the original image it is better to use Edit>Transform>Flip Horizontal (or Flip Vertical). More esoteric transformations are performed by holding down modifier keys while dragging the handles of the bounding box. If you press [cmd/ctl] while dragging a corner handle the corner will move independently of the other points, causing distortion. If the shift key is held down at the same time as [cmd/ctl], and a handle in the middle of one side is dragged, the image will be skewed.

When you use free transformation on a bitmapped image you should also bear in mind that, in general, changing the geometry of a bitmap will necessarily result in it being resampled and this may cause some loss of quality in addition to the distortion.

TRY THIS

Paste any small image onto a new layer of a much larger image. If your image has several layers, hide them all except the new one. Select the new layer with the small image and then practise scaling, rotating, distorting and skewing it with the Free Transform command.

Layer Alignment and Distribution

A collection of layers can be arranged in various ways using the **Layer > Align** and **Layer > Distribute** sub-menus, or by using the alignment and distribution icons which appear in the **Options** bar when the move tool is selected. The icons for alignment only have one line running beside or



through the pairs of tiny rectangles, whereas the icons for distribution have two. (If you are in doubt, hovering over an icon with the cursor will bring up a tool tip which identifies that particular icon.) Linked layers can be lined up horizontally along their top edges, centre lines or bottom edges, or vertically along their left edges, horizontal centre lines or right edges, as indicated graphically by the icons on the buttons (reading left to right). These operations align the areas of the layer which are not transparent, that is, those areas that you have painted on or pasted an image into. In effect, the objects on the layer are being

aligned. Distribution equalizes the space between layers' edges or centres in the same sense, and the same options – top, vertical centre, bottom, left, horizontal centre and right (reading from left to right) – are available as for alignment.



Figure 4.20 *Aligned and distributed layers*

Layer alignment and distribution can be used to produce compositions with a grid-like layout. In the simple example shown in Figure 4.20, five small images were roughly pasted onto separate layers in a document, as shown in the upper

illustration. The result shown at the bottom was achieved by selecting those five layers and then clicking on two icons: **Align bottom edges** and **Distribute horizontal centers**. While this could have been done by eye, with the aid of some guides or a grid, it is much quicker to do, and more likely to be accurate, if you use the **Align** and **Distribute** commands.



TRY THIS

Using **Align and **Distribute**, arrange the small images that you pasted onto separate layers in earlier exercises in all the possible grid-like layouts that their number permits. For example, if you have six images, arrange them as two rows of three columns, and three rows of two columns, as well as one aligned row.**

Flattening and Merging Layers

Photoshop experts working on complex graphics often use a large number of layers – a hundred is not uncommon. Many of these may be hidden, having been used as experiments or interim stages. Because each layer can occupy as much space as a single image of the same size, using this many layers requires a machine with a lot of memory and a large scratch disk. Before saving a final version of an image for display, it is common to *flatten* the layers, that is, combine them all into a single layer. (Be sure to save a safe copy of the image file in its layered form before you do this. You'd be surprised how often you need to go back to an image after it is supposedly finished.) If you save the image in certain formats, including the common Web image formats, layers will be flattened automatically, because those formats do not support layers. However, there will be other times when you need to flatten explicitly. This is done by selecting **Flatten Image** from either the **Layers** panel drop-down menu or the main **Layer** menu. You can also selectively merge layers: the **Merge Layers** command (which also appears on both the **Layers** panel menu and the **Layer** menu) merges all the currently selected layers. (Note that it does not merge all linked layers automatically if only one of them is actually selected.) If no more than one layer is currently selected, the available merge options change to **Merge Down**, which merges the selected layer with the one below it, and **Merge Visible**, which merges all currently visible layers into one layer.

Layer Groups and Layer Comps

Although the examples in this chapter only feature a few layers, Photoshop images often comprise tens or even hundreds of layers. To help you organize large numbers of layers and reduce the **Layers** panel to a manageable size, you can collect layers together into *layer groups*. A layer group resembles a folder on your hard disk, in that it provides a container for a collection of related items. Indeed, the icon at the bottom of the **Layers** panel for creating new layer groups looks like a folder. It works in the usual way: clicking creates a new group with a default name; **alt/opt**-clicking brings up a dialogue in which you can name the group and set some options. You can group layers by selecting them in the **Layers** panel and dragging and dropping them on to the group's entry. To create a group and populate it in one go, you can drag selected layers on to the new group icon. Transparency and blending mode can be set for a layer group. They are applied to the group as a whole, on top of any settings that have been applied to the individual layers.



Layer comps help you work with layers in another way. Comp is an abbreviation of composition; a layer comp is a composition formed by a combination of layers. Professional designers often create different versions of an image to show to a client by combining different sets of layers. For instance, Figure 4.21 shows two different versions of our composite lighthouse image. Each is made by combining two layers, as we explained earlier. All four layers are kept in the same Photoshop file (together with some invisible layers used as the basis for the final ones). Each of the two images is made by making two of the four layers visible, as you can see from the **Layers** panels shown beside each one.

A layer comp is a record of the state of the Layers panel. It records which layers are visible, their position and any layer styles that have been applied to them. Here, we changed the visibility of some layers, so we can create a layer comp for each state. This is done using the **Layer Comps** panel. Clicking on the **New** button opens the simple dialogue shown in Figure 4.22, where you can give the comp a name and specify which properties of the layers it should record. If you then change the state of some layers, in this case by altering their visibility, you can make a new layer comp corresponding to the changed state. You can then show the image in one of its states by applying the corresponding layer comp. This is done by clicking in the box next to it in the panel. A small icon identifies the comp that is currently applied. At the right of the composite images in Figure 4.21 we have shown how a layer comp was applied to obtain each of the two images. You can also cycle through a collection of layer comps using the arrows at the bottom of the panel. There is always a layer comp called *Last Document State* at the top of the panel. Applying this restores the image to the state it was in before you applied any of the named comps.

Layer comps are only useful for images, such as collages, which are made by combining several layers. You will also come across a very similar concept later in this chapter when we describe how to create Web animations in Photoshop.

TRY THIS

Create a large document and paste half a dozen small images into it, leaving them on separate layers. Make several different arrangements of the small images, using only some each time, and create layer comps to record the position and visibility of the layers in each arrangement. (Hide the layers containing the images you are not using in each case.) Show each layer comp in turn as if you were presenting the alternative compositions to a client.

DON'T FORGET

Hide or show layers by clicking the eye icon in the Layers panel.

Pasting creates a new layer.

Set opacity and blending mode to control how layers interact.

Select multiple layers by shift-clicking or cmd/ctl-clicking in the Layers panel.

Use the move tool and Free Transform command to move and transform layers.

Enter numerical values in the options bar for precise transformations.

Use the Align and Distribute menus or option bar controls to produce regular arrangements of linked or selected layers.

Flatten or merge layers to save space, but always keep a layered version.

Use layer comps to create different versions of a composition in a single file.

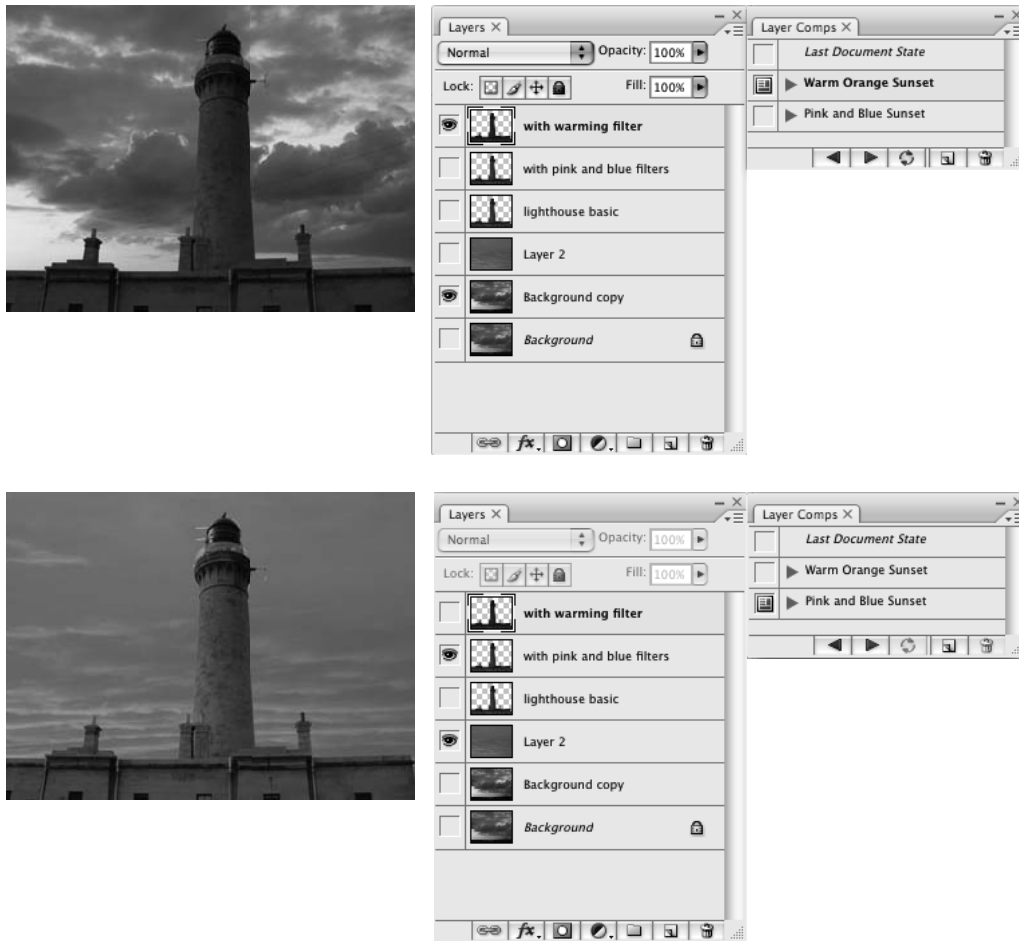


Figure 4.21 Layer comps

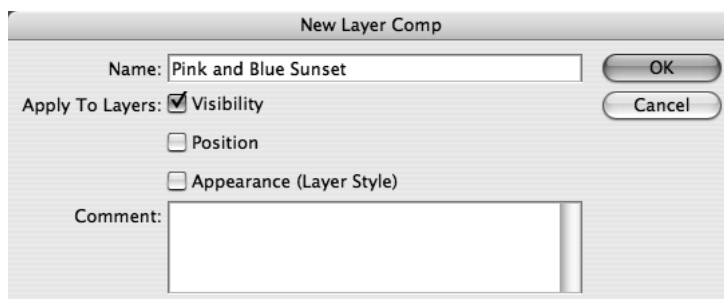


Figure 4.22 Creating a new layer comp

5

Illustrator

In This Chapter

- **Technical drawing: strokes, fills and shape tools**
- **Freehand drawing: pencil tool and brushes**
- **Selection, transformation and reshaping**
- **Complex fills and strokes, Live Paint, styles, effects and 3D**
- **Symbols and instances, symbol sets and symbolism tools**
- **Combining bitmaps and vector artwork**
- **Live Color, guides and alignment, layers and type**
- **Exporting Web graphics**

Freehand Drawing

Some sorts of drawing, particularly technical diagrams, require the precision of the shape tools described in the previous section. You could think of this style of drawing as being essentially *algorithmic*: the objects are created by entering numerical values or from simple gestures that have the effect of setting values. The pen tool, which is described in Chapter 8, uses a similar algorithmic approach to creating smooth curves. When you draw with the pen, you drag lines that relate to the mathematical parameters of the curve; you don't actually draw the curve itself.

People who can't draw very well may feel more comfortable with this algorithmic approach to drawing, particularly if they are used to thinking about mathematical descriptions of things, but artists and designers who are used to drawing with conventional physical tools and media usually find this a completely alien way of thinking about drawing, and prefer to work in a way that more closely resembles their usual experience of making marks. Photoshop's brushes, which are described in Chapter 4, attempt to simulate natural media and processes, which seems quite natural when you are applying colour to the actual pixels of a bitmapped image, but in Illustrator, too, drawing tools can be used in a freer way than its vector basis may lead you to suppose.



We will start with Illustrator's *pencil*. Superficially this is used like a real pencil, and works best if you have a graphics tablet. You select the pencil from the **Tools** panel, and then draw with the graphics tablet pen to make a line. If you only have a mouse, you must drag the pencil about with the mouse button depressed, releasing the button when you have finished. (This is harder to control than a graphics pen, as well as being bad for your hand.) As you draw, a dotted line follows the cursor, as shown on the left of Figure 5.10; when you stop drawing, this is converted to a finished line using the current stroke, as you can see on the right of the figure. Hence you can draw pencil lines of any colour and weight (thickness).




Figure 5.10 *Drawing a pencil line*

If you play with the pencil, just scribbling some lines, you will probably see that it really behaves differently from a real pencil: the finished line does not exactly follow the track of the graphics pen or mouse. Small irregularities are levelled out and the final line is a continuous smooth curve of uniform thickness.

There is also a difference between the pencil and the shape tools we previously described. Whenever you finish drawing a line with the pencil, you will find that, whatever you may have set it to before, the fill has been set to none. As we will see later, when you draw any path, it is treated as if its ends were joined with a single straight line, of width zero. If the fill was not set to none, the path would be filled in, which is rarely what you want to do when making marks with this tool. If you do want paths filled as you draw, you can set an option to make this happen.

Beware, though, that if you do this and inadvertently set the fill to white instead of none, you may find parts of your lines mysteriously disappearing.

If you want to draw a closed shape without having to return the pencil to its exact starting point, you can hold down the [opt/alt] key while drawing (not before you start) and keep it held down until after you have let go of the mouse button or lifted your graphics pen. The start and end of the line you have drawn will be joined up with a straight line in the current stroke.

TRY THIS

Try scribbling some lines with the pencil tool, using different stroke weights and colours. Experiment with holding down [opt/alt] while you draw, to make closed shapes with the pencil.

Make a simple drawing composed entirely from pencil lines.

Since Illustrator is a vector-based program, the marks you make with the pencil tool must be converted to some mathematically representable form. Bézier curves, which are described in Chapter 8, are used for this purpose. The precision with which the curve approximates the movement of the pencil is determined by the tolerances specified in the pencil options dialogue, shown in Figure 5.11, which is invoked by double-clicking the pencil in the Tools panel. Two values control the tolerance. The Fidelity might be better called the infidelity: it is the number of pixels by which the approximation may deviate from the path followed by the pencil. The Smoothness is a measure of the extent to which the bumps in the path are smoothed out. If the smoothness is set to zero or a low value, a curve that reflects any jaggedness in the drawing motion is produced. High values smooth this jaggedness out, so if you want smooth curves (especially if you are trying to draw with a mouse) you should set the Smoothness value quite high. Ticking the checkbox labelled Fill new pencil strokes causes the filling behaviour we described earlier. We will explain the significance of the other checkboxes in the Options section later.

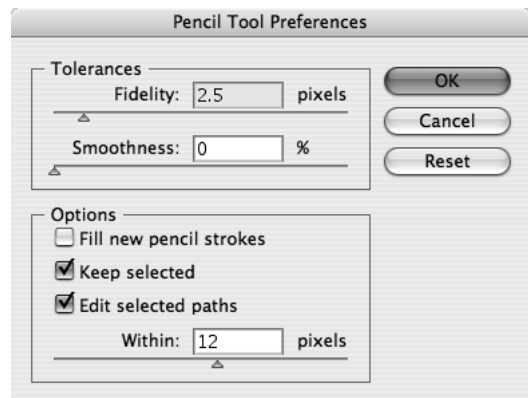


Figure 5.11 Options for the pencil

TRY THIS

Scribble a line with the pencil tool, making a zig-zag movement with your graphics pen or mouse. Change the fidelity and smoothness of the pencil tool a few times, making a similar scribble each time and noting the different appearance of the resulting curves.



The closest Illustrator comes to natural media drawing is in the use of its *paintbrush tool*. Again, this is best used with a graphics tablet and a pressure-sensitive pen if possible. In use it is similar to the pencil: you drag the tool around and a path that approximates your movements is built out of Bézier curves. Double-clicking the paintbrush in the **Tools** panel brings up an options dialogue identical to that for the pencil. The difference is that instead of a uniform stroke being applied to produce lines that look as if they have been drawn with a technical draughtsman's




Figure 5.12 Marks made by calligraphic brushes

round-nibbed pen, a variable stroke is used. The exact appearance of the stroke depends on the particular type of brush you have chosen to use. Illustrator provides four broad categories of brush. To begin with, we will consider its *calligraphic brushes*, which produce marks that look as if they have been applied with some sort of calligraphic pen or an italic nib. Figure 5.12 shows some examples of the different strokes that are available from among Illustrator's default calligraphic brushes.




Figure 5.13 The Brushes panel

You need to set the stroke colour before using the paintbrush (usually you won't want to use a fill), and then select a brush from the **Brushes** panel. In the default configuration of the panel, shown in Figure 5.13, only the top row contains calligraphic brushes; we will come to the other types shortly. You can access several more designs by choosing *Artistic_Calligraphic* from the *Artistic* sub-menu of either the **Window>Brush Libraries** sub-menu, or the **Open Brush Library** sub-menu on the **Brushes** panel menu, to display an extra panel of calligraphic brushes. With a brush chosen, you can set tolerance values and other options by double-clicking the brush in the **Tools** panel.

The way that stroke weight interacts with brushes is quite confusing. If you hover the cursor over a brush in the **Brushes** panel, you will see that each brush has a stroke weight as part of its name. When you first select a brush in the panel, the stroke weight, as shown in the **Control** panel, is set to 1 pt, but the mark that is made by the paintbrush will have the weight indicated by its name. For instance, if you select the 5 pt oval brush, you will get a brush stroke roughly 5 pt wide. (The width of the stroke won't be uniform.) If you set the stroke weight after you choose the brush but before making any marks, the width of the stroke will be multiplied by the value you set. For example, if you set the stroke weight to 2 pt, you will get a stroke twice as wide as if you had left it at 1 pt. So, in the example just given, using a 5 pt oval brush with the stroke weight set to 2 pt, the marks actually made by the brush would be roughly 10 pt wide.

When the paintbrush is selected, its style and stroke weight can be set from the **Control** panel.

Rather than setting the stroke every time you use a brush, you may find it more convenient to create a new brush that incorporates your preferred weight. For calligraphic brushes this is fairly easy. Since the other parameters that control the marks made by such a brush are easy to understand, you can create your own brushes if you find that you use calligraphic brushes frequently.

In the **Brushes** panel menu, select **New Brush...** Then, in the simple dialogue that is displayed, select the **New Calligraphic Brush** radio button and click **OK** to get the dialogue shown in Figure 5.14, in which you can set the characteristics of your brush, starting with its name. The other characteristics are the diameter, angle and roundness. All of these can either be

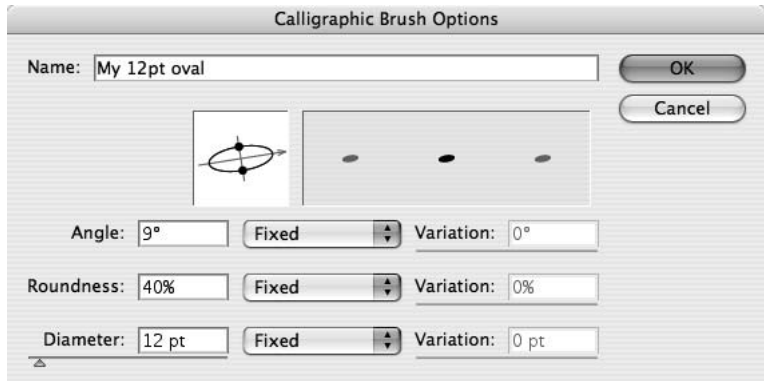


Figure 5.14 *Defining a calligraphic brush*

set numerically or by direct manipulation: the diameter with a slider, the other two values using the iconic representation in the middle of the dialogue. To adjust the angle, pull the arrowhead round; to change the roundness of the brush (the eccentricity of the ellipse, technically speaking) pull one of the black dots in or out. The preview on the right shows the resulting shape of the brush tip. Brush strokes are made to look as if that shape was being pulled across the paper at the specified angle, hence, if the shape is a pronounced oval, vertical and horizontal strokes will be different widths, as though you were using a pen with an italic nib. Using the pop-up menus next to the value of each parameter, you can choose to make that parameter respond to pressure when a graphics tablet is used, or you can add random variation, to avoid the mechanical look that tends to result from applying exactly the same brush stroke to different paths.

Creating a good brush that makes the marks you want requires plenty of practice. If you wish to create custom brushes it is worth taking the time to experiment extensively.

TRY THIS

Make some marks with each of the calligraphic brushes, using different colours. When you feel that you understand what sort of marks you can make, create a picture using the brushes.

Make a new calligraphic brush that is wide and nearly flat. Practise altering each of its characteristics. Draw some lines and curves with your modified brush each time you change a value, until you can see how changing the parameters affects the marks you make.

6

Flash

In This Chapter

- Flash documents and movies: the stage and the timeline
- Creating and importing artwork
- Graphic symbols and instances
- Motion tweening, easing and motion guides
- Editing animations and applying effects
- Movie clip symbols and animated graphic symbols
- Tweening filters and shape tweening
- Sound and video in Flash

Selection

Even if you avoid drawing in Flash by importing artwork from another program, you will sometimes need to be able to select parts of a drawing so that you can move or transform them. When we come to animating in Flash, for example, you will see that one way to make animation is by moving or transforming selected shapes or parts of shapes. Flash provides a *selection tool*, whose icon is the conventional black arrow. When you have selected something with this tool, you can move or transform it, or change its stroke and fill using the controls in the **Properties** panel. Flash's way of selecting and manipulating objects is not conventional, however. It has some unexpected features, which can be confusing for people familiar with other media tools programs.

Crucially, the way in which objects behave when they are selected depends on which *drawing mode* was in force when they were created. The default drawing mode – and therefore the one you are likely to experience first – is *merge drawing*. In this mode, whenever you draw something it interacts with anything that is already on the stage. Objects that are drawn in this way exhibit three distinctive characteristics. First, strokes and fills can be selected independently. Second, by default only individual segments of a compound path are selected. Third, where shapes overlap, they are split into separate pieces. This can lead to much confusion and frustration at times.

Consider the filled oval shown at the left of Figure 6.14. Clicking inside the shape with the arrow tool does not, as you might expect, select the entire filled ellipse. Instead, it just selects the fill. Once it has been selected, the fill can be manipulated independently of the enclosing stroke that is apparently its boundary; for example, it can be dragged out of it, as shown. If you want to select an entire shape that was made in merge drawing mode – that is, the fill and stroke together – you must double-click in the fill.

Where an object is drawn from separate strokes, as, for example, you might draw a triangle with the line or polystar tool, clicking on a line only selects the part drawn with an individual stroke. With freehand lines, you may find that when you click on a line a much smaller part of it is selected than you expect. Double-clicking a line selects all the lines connected to it, so in the example of a triangle, double-clicking any one side of the triangle would select all three sides (but double-clicking inside the triangle would select its fill as well).

Things become more complicated when shapes overlap, because in Flash, when a line crosses a shape, it divides that shape into segments which will be selected separately by the arrow tool. The line is also divided into segments at the points where it crosses the shapes. After moving the fill of our ellipse we end up in just this situation: the elliptical line that was the stroke crosses the oval fill. If we now click on the line and drag it, only the visible segment moves. The original oval line has been split. The same thing happens again in the new position. The small segment at the top can therefore be selected and dragged independently, as shown at the right of Figure 6.14.




Figure 6.14 Selection of line and fill for objects drawn in merge drawing mode

In order to select the areas or lines you need, you can hold down the **shift** key while clicking or double-clicking. If you do so and click on an unselected object it will be added to the selection; if you **shift-click** on an already selected object, it will be removed from the selection.

The arrow tool can also be used to select several objects or parts of objects at a time. If you hold down the mouse button and drag with the arrow tool, a selection rectangle is pulled out. Any object inside the rectangle becomes selected when you release the mouse button. If the objects were drawn in merge drawing mode, where an edge of the selection rectangle passes through an object, both strokes and fills are split along the edge, and only the parts that lie within the rectangle will be selected. This is another potential surprise if you are used to the behaviour of selection tools in other programs, but it is consistent with the way Flash treats overlapping shapes made in merge drawing mode.

Selections can also be made with the **lasso tool**. If this tool is used without modifiers, it will produce a freehand outline and anything inside the outline will be selected. Again, if the objects were drawn in merge drawing mode, both strokes and fills will be split wherever the selection outline intersects the objects. If you select the lasso tool and then click the polygon mode button at the bottom of the Tools panel, the lasso behaves like the line tool, enabling you to draw polygonal selection areas. In this case, you must double-click to close the selection polygon; for standard lasso selections it is sufficient to let go of the mouse button.

TRY THIS

Draw some shapes or a simple scene or object of your choice, using both strokes and fills (in the default merge drawing mode). Experiment with altering parts of the drawing by making selections with the arrow and lasso tools and then changing line characteristics, fill colours and the arrangement of the shapes.

Merge drawing mode may not seem odd if you have never used any other vector drawing program, although the way that selection splits objects and strokes can be very disconcerting. If you have used another vector drawing program, however, you will probably think that the objects you draw – ellipses, rectangles, and so on – should retain their identities. For instance, in Figure 6.14, you would expect that you could select the entire filled ellipse and drag it as a single object. You can indeed do this in Flash, but only if it was drawn in **object drawing** mode.



If you look back at Figure 6.6, you will see that, as well as the smoothing options, there is a small round button in the options area of the Tools panel. This appears when any drawing tool is selected, and it is used to switch between the default merge drawing mode and object drawing mode. To use object drawing mode you must explicitly turn object drawing on by clicking this button – but don't forget to turn it off again if you wish to return to merge drawing mode. Any shape which was drawn in object drawing mode can be selected with a single click using the arrow tool, and it does not split any object it intersects. Furthermore, if a selection rectangle intersects an object, it selects all of it, without splitting it into pieces. This is often much more useful. (If you import artwork from Illustrator, as we shall describe shortly, the objects behave more or less as if they had been created in object drawing mode. That is, they continue to behave roughly in the same way that they did in Illustrator.)

The brush tool is, by its nature, a special case. You can use it in either object or merge drawing mode, which affects how the brush marks interact with other objects. However, even in object drawing mode, if a single brush stroke (i.e. one that you make in a single drag) overlaps itself, the partial strokes merge, so you can, for example, paint an area by scribbling with the paintbrush.

It is possible to convert objects from merge drawing mode to object drawing mode. Select one or more objects, and choose the Union command from the Modify>Combine Objects sub-menu. They will then be treated as a single object and won't affect objects that they intersect. Conversely, if you select an object, and then choose Modify>Break Apart, its fill and stroke will become separately selectable, and will interact with any objects they overlap, as if they had originally been created in merge drawing mode.

TRY THIS

Practise using all of Flash's drawing tools in object drawing mode. What is the effect of drawing shapes in merge drawing mode on top of shapes made in object drawing mode and vice versa? Experiment further with selection.

The Modify>Combine Objects sub-menu has several commands besides Union. Draw several overlapping shapes in object drawing mode and experiment with using the different available Combine Objects commands.

You can use the Edit>Cut and Edit>Copy commands to move or copy any selected objects to the clipboard, in the usual way. Flash provides two different commands for pasting: Edit>Paste in Place and Edit>Paste in Center, depending on where you want the pasted object to be placed: in its original position or in the middle of the stage.

If you select several objects, you can group them together, using the Modify>Group command. As long as they remain grouped, the objects can be manipulated as a single entity. In particular,

the group can be moved as a whole, and transformed in the various ways we will describe in the next section. If you need to work with the individual members of the group again, you can use the **Modify>Ungroup** command to separate them out.

Transformations and Reshaping

Once you have selected an object on the stage you can transform it in various ways. Broadly speaking, there are two approaches you can take to transformation: either you can do it by hand and eye, manipulating objects directly on the stage, or you can do it numerically, by entering values in the **Info** and **Transform** panels. The first approach may feel more natural for artists and designers; the second provides greater accuracy.

The simplest form of transformation is movement. This is done by selecting the object or objects to be moved with the arrow tool and then dragging; holding down [opt/alt] while dragging causes the objects to be copied, as you might expect. Note that, although a bounding box appears round an object drawn in object mode when you select it, you still have to drag inside the object itself to move it. If it has no fill, you have to drag the stroked outline.




Figure 6.15 *The free transform bounding box*

Other transformations are performed using the *free transform* tool. This can be used like the arrow tool to select and move an object, but it has extra capabilities. When you select anything with the free transform tool, a bounding box with handles at the corners and half way along each side appears, as shown in Figure 6.15. If you place the cursor directly over one of the handles, it changes to a double arrow to indicate that you can scale the object by dragging. If you are over one of the handles in the middle of a side, you can scale in the direction perpendicular to the side, as illustrated in the middle of the top row of Figure 6.16; if you are over one of the corner handles, you can scale in both directions at once by dragging diagonally. If you hold down the shift key while dragging a corner, the object's proportions are maintained as its size changes.




Figure 6.16 *Free transformations*

If you hold the free transform tool near, but not actually over, one of the corner handles, the cursor turns into a circular arrow; this indicates that you can rotate the object by dragging it, as illustrated at the top right of Figure 6.16. (The corner handles can actually be dragged in more or less any direction to make the object rotate,

but moving them as if you were actually pulling the object round to its new position produces the most intuitive feedback.) By default, the object is rotated around its centre point, but you can move the centre of rotation by dragging the circular icon – visible near the middle of the shape in Figure 6.15 – to a different position.

The free transform tool can also be used for performing skewing transformations. When you move the tool near to a side of a bounding box, but not over any handle, the cursor turns into a pair of parallel arrows, and you can skew the object by dragging parallel to the side. Finally, by holding down [cmd/ctl] while dragging a corner handle, you can distort the bounding box and its contents. Some effects of skewing and distortion are illustrated at the bottom of Figure 6.16.

The **Modify>Transform** sub-menu offers alternative ways of invoking the scaling and rotation tools. It also provides some further useful transforms, including reflections in the horizontal and vertical axes (**Flip Vertical** and **Flip Horizontal** respectively).




Figure 6.17 *The Info panel*

If you know the exact coordinates to which you want to move an object, you can select the object and enter the values in the X and Y fields at the top right of the **Info** panel, shown in Figure 6.17. (Note that the X and Y values in the bottom right of the **Info** panel show the current cursor coordinates.) Click in the appropriate square in the little grid to the left of the X and Y fields to choose whether to set the coordinates of the top left corner or of the centre of the object. Similarly, if you know the exact size to which you need to scale an object, you can enter the width and height in the W and H fields.




Figure 6.18 *The Transform panel*

The **Transform** panel, shown in Figure 6.18, offers a numerical way of applying scaling, rotation and skew transforms to selected objects. The two fields at the top of the panel are for entering percentage factors by which to scale the object horizontally and vertically (indicated by the tiny double-headed arrows). By selecting the checkbox marked **Constrain** before you enter any values, you can ensure that the scaling is uniform – that is, that the height and width are scaled by the same proportion. In this case, you only need to enter one of the values as the other is automatically set to match it.

The remaining three fields are for providing values, in degrees, for a rotation or skewing. Use the radio buttons to select either rotation and skew – they cannot be applied at the same time. (But separate amounts of horizontal and vertical skewing can be applied at the same time.) Rotations are measured clockwise, so use negative values if you want to rotate objects anti-clockwise. In the bottom right corner of the panel there are two easily overlooked buttons. The one on the right



resets all transforms, so long as the object is still selected – clicking it returns the object to its original state. The left-hand button is used to copy the object while transforming it.

TRY THIS

Draw a simple asymmetrical object and practise using the free transform tool to alter its scale, position on the stage and angles of rotation and skew. Try moving the circle that marks the centre of rotation, and see what effect it has on the different transformations. What happens if you drag the handle in the middle of a side all the way through the object and out the other side?

Try performing the same transformations using the Info and Transform panels and the commands on the Modify>Transform sub-menu.

Now draw several objects, group them together, and experiment with applying transformations to the whole group.

Alignment

A different way of moving objects is by aligning them with each other. The Align panel, shown in Figure 6.19, provides a means of laying out objects in various ways. This has many uses – for example, when laying out controls in a movie that provides interactivity via buttons. When you click on one of the icons in the Align panel, all currently selected objects are moved in the way indicated. Most of the icons are self-explanatory. For example, clicking on the icon at the left of the top row causes all the selected objects to be aligned along their left edges – the selected object that is originally furthest to the left is used to set the position to which the others are aligned. Similarly, objects can be aligned along their right, top and bottom edges, or centred vertically or horizontally. They can also be distributed (using the second row of buttons) so that the distance between corresponding edges is the same. Unless the objects are all the same size, this often produces unexpected results. More often, what you want to do is distribute the objects so that there is an equal amount of space between each of them. This is what the two icons labelled **Space** in the bottom right of the panel do. Selecting the button at the far right of the panel, labelled **To Stage**, causes the alignment or distribution to be performed relative to the stage. For example, if you select this option and then click the left align icon, all the objects will be lined up on the left-hand edge of the stage.

You can also use the Align panel to scale objects so that they are all the same size. This is done with the **Match size** buttons in the bottom left of the panel, which match the selected objects' horizontal or vertical dimensions, or both. In each case, the size is set to match the largest selected object. Again, you can select **To Stage**, in order to make objects as tall or as wide as the stage.

TRY THIS

Draw nine different objects and try out all the different ways of aligning and distributing them on the stage. Finally, arrange them in an evenly spaced three by three grid, centred on the stage.




Figure 6.19 The Align panel

7

Dreamweaver

In This Chapter

- **Web pages and sites, accessibility and validation**
- **Dreamweaver's user interface and preferences**
- **Page layout and simple formatting**
- **Creating and using CSS rules for typography and layout**
- **Adding images, Flash, tables and forms**
- **Using behaviours and the Spry framework for interactivity**
- **Setting up and updating sites**
- **Using templates**

JavaScript

Web pages that consist only of styled text and images are *static* – they don't do anything once they have been loaded into the browser. In contrast, pages with *dynamic* elements respond to events such as mouse clicks and rollovers, by changing in some way. Apart from CSS rollovers, implementing dynamic page elements requires *scripting*. That is, a small program must be written which is executed in the browser when some event occurs, and causes some elements of the page to be changed. Scripts are written in a relatively simple language called *JavaScript*.

As we explained in connection with image-based navbars, relying on JavaScript may cause problems, because a significant number of people disable it or use devices that do not support it. Hence, it is important that you make sure that your page can be used – links can be followed, forms can be submitted, and so on – even if the scripts don't work. Furthermore, in the past there have been considerable incompatibilities between different browsers' support for JavaScript. With the arrival of a standard *Document Object Model*, which specifies how scripts can access and change elements of the page, these incompatibilities are becoming less of a problem. Nevertheless, it is important to test pages that use JavaScript in as many browsers as possible.

The code for responding to events and updating the page is usually fairly trivial, but any kind of programming requires practice and skills which many designers lack. Dreamweaver provides ready-made code for certain commonly used types of interaction, concealing the scripts almost entirely. If you have no programming experience, you can still add dynamic behaviour to your pages by using these scripts. They fall into two classes, which were added to Dreamweaver at different times. We will describe them in the order they were developed.

Behaviours

The older class consists of a set of pre-built parameterized scripts, called *behaviours*, which you can attach to elements of a page to perform some *action*, such as replacing one image with another, or hiding a part of the page. Each behaviour has some *parameters*, such as the URLs of an image and its replacement, for which you specify values when adding the behaviour. You also select an *event*, such as a mouse click or rollover, that should trigger the behaviour. Whenever that event occurs, the action will be performed in response.

Behaviours are fairly crude, and many of them either cause actions that are not recommended nowadays, such as opening a pop-up window, or are ineffective, such as trying to detect the type of browser being used. Others, such as page elements that move across the screen, have become clichés that hark back to the days when “Dynamic HTML” was in vogue.

We have already seen an example of the use of behaviours, when we looked at image-based navbars. The underlying image swap effect is available on its own, as one of the image objects on

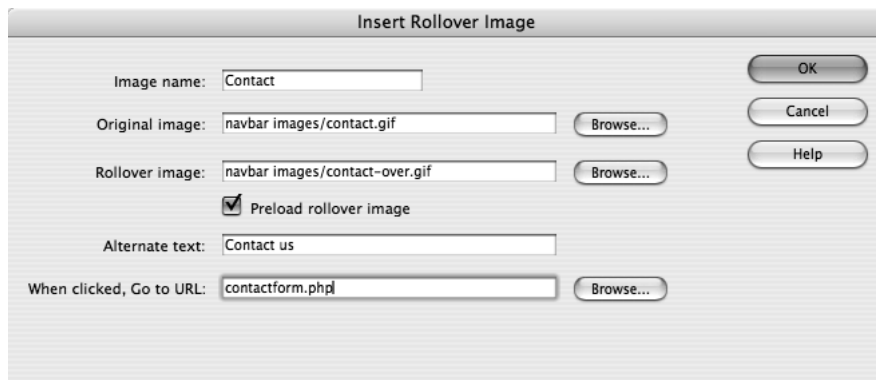


Figure 7.71 Making an image rollover

the Insert bar. (Look back at Figure 7.4.) When you click on this icon, the dialogue box shown in Figure 7.71 will appear. You choose two images, one for the normal state and one for the rollover state. (There is no provision for a down state.) You can add a URL, so that the image functions as a link, and you must provide a textual alternative.

The Behaviors Panel

Rollovers and image-based navbars are special cases of interactivity, in which the element involved in the behaviour and the event that triggers the action are pre-determined. This makes it possible to use the Insert bar to add these behaviours. In the more general case, you need to select the element and event, as well as setting parameters for the behaviour; this cannot be done just by clicking an icon on the Insert bar. Most behaviours are therefore added using a longer sequence of operations. First, an element is selected; second, a behaviour is added to the selected element, and its parameters are set using the Behaviors panel; finally, if the default is not suitable, the event that should trigger the action is specified. You can add any number of behaviours to an element, although you are rarely likely to need more than one.

The Behaviors panel is docked with the Attributes panel in the Tag Inspector, so when you select an element its name is incorporated into the title bar to show you which element the behaviour is being attached to. The panel resembles the top pane of the navigation bar dialogues: +, - and up and down triangle buttons are used to add, delete and rearrange elements of a scrolling list of behaviours. Clicking on the + button causes a menu to drop down, displaying all the available behaviours (as illustrated in Figure 7.72). Any behaviours which cannot legally be attached to the currently selected document element are greyed out.




Figure 7.72 Choosing a behaviour




Figure 7.73 *Setting the parameters of a behaviour*

Note the presence of the `~Deprecated` sub-menu. This contains behaviours that you should not use; they are only there for compatibility with pages made in earlier versions of Dreamweaver.

Selecting a behaviour from the list causes a dialogue to be displayed, in which you enter the appropriate parameters. Figure 7.73 shows an example. As with all the behaviours provided with Dreamweaver, the meaning of the parameters is fairly self-evident. However, in case of doubt, every parameter dialogue for a behaviour has a **Help** button, which causes a detailed explanation of its usage to be displayed in Adobe Help Viewer.

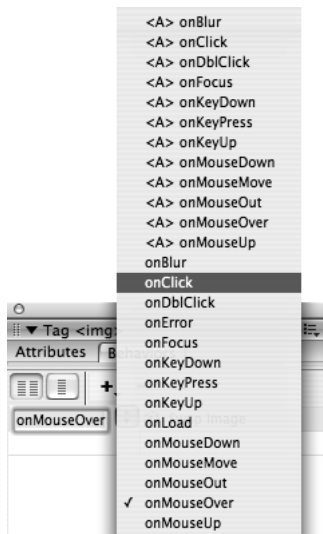


Figure 7.74 *Choosing the event*

Once you have provided the behaviour's parameters, it is added to the list in the **Behaviors** panel. This list has two columns: the right-hand column lists all the behaviours attached to the selected document element; the left-hand column shows the event which will trigger the behaviour listed next to it. A default event, which is usually the most sensible, is initially chosen. If this is not, in fact, the event you wish to use to trigger the behaviour, you can select the behaviour and, as you can see in Figure 7.74, the entry in the events column acquires a pop-up menu from which you can select an alternative event. Most events can be associated with most types of element, although a few, such as `onFocus` and `onBlur`, can only be associated with form controls and links. The entries in the menu of events that begin with `<A>` will cause Dreamweaver to put a dummy link around the selected element to receive an event which otherwise it could not. This is not likely to be something you need to do unless you expect visitors to your site to be using very old browsers.

Space does not allow us to describe all the available behaviours; in any case, since new behaviours can be added, no such description could be complete, so we will restrict ourselves to describing



Figure 7.75 An image gallery

a single example showing how behaviours can be used to create a page that reacts to user input in a useful way. The Dreamweaver Help facility should provide adequate assistance in using the others, but bear in mind our earlier remarks about their usefulness, or lack of it.

The Swap Image behaviour is the basis of rollovers, but you can use it explicitly to achieve different related effects. In particular, you can make the image that is swapped different from the one that receives the event, to create a so-called “remote rollover”, where moving the cursor over one image causes a different one (somewhere else on the page) to change. Alternatively, you can associate the behaviour with the onClick event, so that clicking an image causes another one to change. We will do this to create the image gallery illustrated in Figure 7.75. (We did tell you not to take Dereq Ffisq’s artwork seriously....)

In order to be swapped, an image must have a name, which you give it by typing in the text field at the left of the Properties panel while it is selected. Once this is done, creating the effect is simple.

The image which is to serve as the control is selected and the **Behaviors** panel is used in the way just described to add a behaviour to it, in this case **Swap Image**. The parameters dialogue is the one shown in Figure 7.73. The list at the top shows the names of all the images on the page; you must select one of these to be replaced by clicking on its name in the list. The field labelled **Set source to** is filled in with the URL of a replacement image. Normally, this will be within the site, so you can use the **Browse...** button to navigate to it.

When the behaviour is triggered, the image selected is replaced by one contained in the chosen URL. If the **Preload Images** checkbox is ticked, the images will be downloaded with the page to ensure a smooth swap. If **Restore Image onMouseOut** is selected, the original image will be restored when the mouse moves away from the control. This is desirable if you are using the default **mouseover** event to trigger the swap; otherwise, you may prefer to deselect this option, as we did in the present example.

The intention is that when a user clicks on one of the four small images down the left side of the page, the large image to the right will be replaced by a full-sized version of the thumbnail that was clicked. We prepared the four large images, all at the same size so they could be swapped without resizing or distortion, and, for each one, we also saved a scaled-down version to act as the thumbnail. The thumbnails were added to the page in the usual way, and the large version of the first image was added in the right column, and given the name **main**. All we then had to do was select each thumbnail in turn, and attach the **Swap Image** behaviour to it, with the event set to **onClick**. In each case, **main** was selected in the parameters dialogue and the URL was set to point to the full-sized image corresponding to the selected thumbnail.

— TRY THIS —

Make your own image gallery for your site, using Swap Image behaviours.

By attaching a behaviour to the body element triggered by the `onKeyPress` event, make the gallery revert to displaying the first image whenever the user presses a key. (The easiest way to select the body is by clicking its name in the tag structure display at the bottom of the document window.)

The Spry Framework

Dreamweaver's newer mechanism for adding interactivity is based on a JavaScript library known as the *Spry framework*. There are three facets to this framework. *Spry effects* can be used to add simple visual effects to elements of a page, such as making them fade in or out. *Spry widgets* are user interface elements that go beyond the simple form elements provided by HTML, and more closely resemble some of the interface elements found in desktop applications. *XML data sets* are used to display data in the form of XML on the page. The data is retrieved from the server, where it may be generated dynamically by a script.

The technique of using JavaScript to retrieve and display XML data dynamically is referred to as **AJAX** (which stands for Asynchronous JavaScript and XML). The term's meaning is often extended to include any use of JavaScript to provide sophisticated interactivity, so Spry widgets and effects might also be included.

It should be apparent that over-use of effects can easily become tiresome, and that using non-standard interface elements is likely to confuse some visitors to your Web site. Retrieving data dynamically and using it to rewrite the page can also be confusing and lead to usability problems, such as an inability to bookmark a particular state of a page. Furthermore, AJAX has potential security problems, so for all of these reasons it should be used judiciously.

Unlike behaviours, which are copied into the HTML document when you add them, the JavaScript code for the Spry framework is kept in external files, which are referenced from the document. The first time you use a Spry effect, widget or XML data set, Dreamweaver will copy the relevant part of the Spry framework into your site.

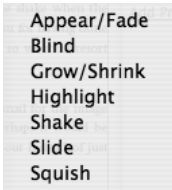
XML data sets can only be used effectively as part of a Web application which includes server-side computation. Otherwise, they are largely redundant, and add unnecessary complications. An account of Web applications is beyond the scope of this book, so we will confine our description of the Spry framework to its effects and widgets.

Spry Effects

Spry effects are easy to use, but also easy to abuse. Before you make part of your page shake when the cursor rolls over it, you should ask yourself whether users are going to thank you for having done so. In fact, we would be hard pressed to find a compelling use for Spry effects, so we must resort to a somewhat contrived example.

Looking back at the image gallery in Figure 7.75, you can see that the thumbnail for the image currently being displayed is redundant – clicking it won't have any effect. Perhaps it would be a good idea to make it disappear, and perhaps it would be nice for it to fade out instead of just vanishing. One of the Spry effects, **Appear/Fade** can be used to achieve this.

Spry effects are applied as behaviours, using the procedure described earlier. They can be found on the **Effects** sub-menu of the pop-up menu used for adding behaviours. They are mostly time-based transitions, such as you might use in video editing, so they include a parameter that sets their duration, that is, the time it takes for the effect to occur. Figure 7.76 shows the parameter dialogue for the **Appear/Fade** effect. The other effects use similar dialogues.



- Appear/Fade
- Blind
- Grow/Shrink
- Highlight
- Shake
- Slide
- Squish

8

Shared Concepts

In This Chapter

- **Colour models and modes**
- **Swatches, colour panels and colour pickers**
- **Gradients**
- **Text frames and text boxes**
- **Formatting characters and paragraphs**
- **Bézier curves**
- **Web image file formats**
- **Optimizing images for the Web**

Colour Panels

If you do not wish to work with a restricted set of colours using the **Swatches** panel, you can specify colours explicitly in the **Color** panel, shown from Photoshop and Flash in Figure 8.5. Illustrator's **Color** panel is almost identical to Photoshop's, except for the two swatches in the top left corner. In Photoshop, these represent the foreground and background colours; in Illustrator, they show the stroke and fill colours. The swatch pairs are identical in appearance and function to the two swatches at the bottom of the **Tools** panel in these programs. You click on one or the other to set the corresponding colour. To set a colour value precisely, you can use the sliders in the middle of the panel or the numeric entry fields. You will only be able to type in numerical values if you know them, which means having some numerical colour specification to hand, but by using the sliders it is possible to produce any desired colour within the colour model you are using: the chosen swatch in the top left corner of the panel updates to show the colour corresponding to the sliders' positions as you move them.



Although you can get used to it with practice, mixing colours from red, green and blue components in this way is not entirely intuitive. Moving the sliders does not always have the effect you expect, and some trial and error is usually required to get the colour you are aiming for. As we noted above, many people find the HSB colour model more comfortable. You can select HSB Sliders from the **Color** panel menu to change the R, G and B sliders shown in Figure 8.5 to H, S and B and use them to adjust hue, saturation and brightness.

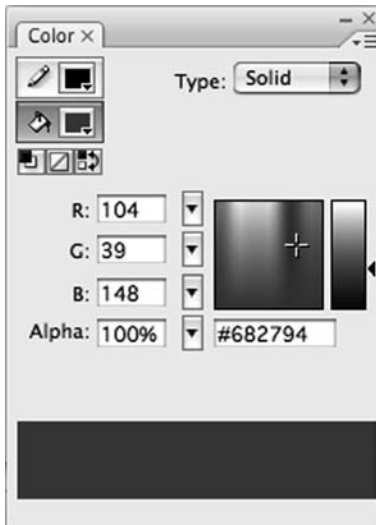


Figure 8.5 *The Color panel in Photoshop (top) and Flash (bottom)*

Flash's **Color** panel, also shown in Figure 8.5, provides a similar interface to Photoshop's and Illustrator's. The sliders are concealed under the triangles to the right of the colour component numbers, but they function in just the same way as in the other programs. There are significant small differences, though: in Flash, clicking on either of the stroke or fill colour icons causes a palette of colour swatches to pop up, from which you can choose a colour, as you do from the **Swatches** panel. (In fact, this pop-up palette effectively is the **Swatches** panel, and any changes you make, such as deleting a colour from the **Swatches** panel, are reflected in the pop-up palette.) The **Color** panel has an additional field and slider, labelled **Alpha**, for adjusting the opacity of the colour. (In Photoshop, opacity is a property of layers, while in Illustrator opacity adjustment is made independently in the **Transparency** panel.) Similarly, in Flash the **Color** panel is where you define gradient fills, but in Illustrator this is done in a separate panel, as we will describe later.

In Dreamweaver, there is no Color panel, but when you click on a colour swatch in the Properties panel, to set the colour of some text for example, a palette of swatches pops up, as in Flash.

—TRY THIS—

Practise setting foreground, background, fill and stroke colours by adjusting the sliders in the Color panel. Choose a colour in an image and try to match it using the Color panel sliders. Try RGB and HSB sliders. Do you find one of the colour models easier than the other to use for this purpose?

Sampling Colours

If you want to match a colour that is already present on screen, use an *eyedropper tool* to sample colours. The basic operation of the eyedropper is the same in all programs: clicking with this tool samples the colour of the pixel where you click. However, eyedroppers can be used in different contexts and behave slightly differently in some programs.



The eyedropper that you will find in the Tools panel in Photoshop, Illustrator and Flash is selected like any other tool. In Photoshop, clicking anywhere in the document sets the foreground colour; [opt/alt]-clicking sets the background colour. There is a pop-up menu in the Options bar that allows you to specify how the sample is taken. The default, **Point Sample**, uses the colour of the pixel where you click. The other options specify the size of a square in pixels; the sampled colour is then the average colour within a square of that size surrounding the pixel you click on.

In Illustrator, the eyedropper sets the fill and stroke colours to match the object you click on, and it applies these attributes to any object that is currently selected. In fact, in Illustrator, the eyedropper samples other appearance attributes, too, including transparency, stroke width and style, and various text attributes. Double-click the eyedropper in the Tools panel to specify which attributes are sampled and applied. Illustrator's eyedropper can sample colours from anywhere on the screen. If you click with the tool and keep the mouse button held down, you can move the cursor anywhere on the screen, and when you let it go the colour will be sampled and used as the fill colour.

Flash's eyedropper behaves in a broadly similar way: if you click the fill of an object, it samples the fill colour and then changes into the paint bucket tool, so you can apply it to another object. If you click the stroke, it samples the width, style and colour of the stroke, and changes to the ink bottle tool. In other words, in each program, the eyedropper samples attributes so that they can be used for other objects as if you had set them explicitly.

The Tools panel is not the only place where you can find an eyedropper, though. In Photoshop and Illustrator there is an area at the bottom of the Color panel, called a *colour ramp*, filled with

a spectrum of colours, as you can see in Figure 8.5. When you hover the cursor over the ramp it turns into an eyedropper, to indicate that you can select a colour by clicking. The eyedropper also appears when the cursor is over the Swatches panel, again indicating that clicking will sample the colour under the cursor. As we described in Chapter 4, an eyedropper is also used in the Color Range and Replace Colors dialogues to sample a colour to use for making a selection.

Finally, when you use the Swatches panel in Flash, or a pop-up palette of swatches in Flash or Dreamweaver, the cursor changes to an eyedropper, as it does in Photoshop and Illustrator, but this eyedropper can be moved anywhere on the screen to sample the colour of any pixel.

TRY THIS

Practise using the eyedropper tools in all programs to sample colours. In Photoshop, set the foreground and background colours to values sampled from a photograph. In Illustrator and Flash, draw some objects, set the stroke and fill of one of them and copy it to the others. In Dreamweaver, set the colour of some text to match a colour from your desktop image or pattern. Make sure that you understand what is being sampled by the eyedropper in each program, and how the sampled values are used.

Colour Pickers

The most elaborate way of choosing colours is by using a *colour picker*, which presents the colours that can be represented in the chosen colour model in a graphical form. Both Windows and Mac OS X have system colour pickers, which can be accessed from most programs, and these are used by Flash and Dreamweaver for picking colours. Adobe have created their own, platform-independent colour picker, and this is used in Photoshop and Illustrator.

The Apple system colour picker can be presented in several different ways. Three of them are shown in Figure 8.6. The leftmost display is an HSB picker. Because monitor screens are two-dimensional, the HSB model must be displayed as a disk, representing hue and saturation for a particular value of brightness, together with a separate slider representing the brightness cylinder. As the brightness changes when the slider is moved, the hue and saturation in the corresponding slice through the cylinder are displayed in the circle. A user chooses a colour by clicking; the hue and saturation can be deduced from the coordinates of the click. The position of the slider provides the third component of the HSB value for the selected colour.

The other views of the colour picker are more easily understood. The middle one illustrated just consists of sliders, like the Color panel. The pop-up menu can be used to select which set of sliders (i.e. which colour model) is used. The right-hand view in Figure 8.6 is the crayon picker, which is more entertaining than useful. Its mode of operation should be evident. As well as these three,

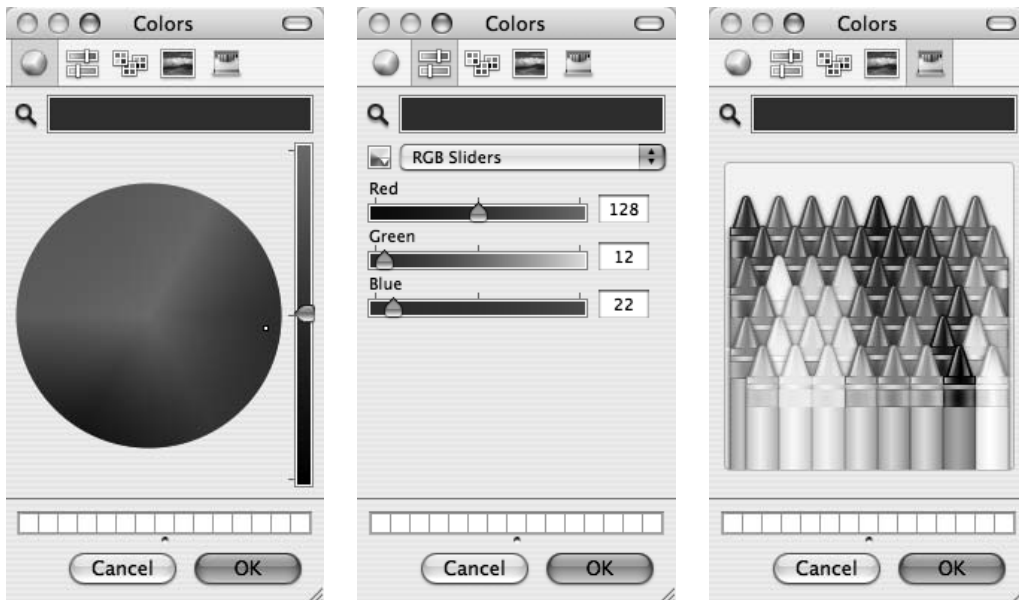


Figure 8.6 The Mac OS X system colour picker

you can also display swatches of named colours from certain standard palettes, or you can choose colours from an image of a spectrum. Thus, the colour picker duplicates the function of all the other available methods of choosing colours, but in a single dialogue. You switch between the different displays using the icons at the top of the dialogue.

In contrast, the standard Windows colour picker, shown in Figure 8.7, presents different methods of selecting colours in a single dialogue. On the left is a set of swatches for a very limited range of colours. On the right is an HSB picker, except that, unlike the Apple picker, hue and saturation are arranged along two perpendicular axes, giving a square instead of a circle, which distorts the colours' distribution somewhat. Colours are picked in the usual way, by moving the brightness slider and clicking to set the hue and saturation components. Alternatively, numerical values can be entered, either in HSB or RGB. As with the Apple colour picker, it is possible to save colours that you use frequently, and have their swatches shown in the Windows colour picker whenever you invoke it.




Figure 8.7 The Windows XP system colour picker