Imaging Services in a Multimedia Environment

Image manipulation tools, compression and decompression functions, picture quality adjustment techniques, and support for industry standards are some of the features included in the HP Image Library.

by Andrew Munro and Ahmad H. Shekarabi

On UNIX*-system-based computers words are the traditional means of communication whether the user is creating a business report or presentation, writing a functional specification, or sending electronic mail. While the topic might benefit from some visual content, the user usually finds it easier to just use words.

Unfortunately, words may not be enough. The user could be describing CAD graphics that remote colleagues need to see. Perhaps, the user is an insurance adjuster who is sending a report to the home office that describes photographs of damaged property. Another user may be describing the appearance of a computer screen that a customer support engineer needs to see.

Without imaging capabilities on the computer, the user may resort to noncomputer means, such as the postal service, to send images.

This article describes the major parts of the HP imaging solution, how it meets the characteristics required of an imaging system, and its application to HP MPower.

Image Files

Image files contain computer graphics and digital records of physical objects, such as photographs, pages from books, and faxes. The images of these objects are represented as collections of bits called pixels.

Image data comes from several sources including screen captures, video frames, and external devices such as scanners, video recorders, or fax machines. Once the image data is collected in a file it can be displayed, modified and saved, or printed.

The data in image files is stored in different formats. These image formats are more commonly called image types. The most popular image file format is called tagged image file format, or TIFF. The following are some of types of images stored in TIFF files. These image types are listed in order of image quality—from a simple monochrome to the highest-quality color:

- Bitonal. Bitonal images contain pixels with two levels—black and white—but they can be displayed using any two colors. Each pixel takes one bit. Bitonal images are frequently referred to as monochrome images. They often contain text, such as a page of a fax.
- Grayscale. Grayscale images contain pixels that identify levels of gray, from black to white. Grayscale images are some-times also referred to as monochrome images. They are often used for images of scanned photographs.
- Palette. Palette images contain pixels that index into a color map containing the red, green, and blue values for the pixel. The color map has three contiguous sections: red, green, and blue, each with 256 16-bit entries. The color of a pixel is determined by using the pixel value as an index into each of the red, green, and blue sections to obtain the desired color. This is the lowest-quality color image and is also called a pseudo color image, a type commonly used for computer-generated images.
- RGB. An RGB image contains pixels with three samples: red, green, and blue. Each sample identifies a level of that primary color. This is a good-quality color image type that is typically used for photographs. RGB format provides better picture quality than the palette format because RGB provides 2^{24} color variations for images, whereas the palette provides 2^{8} color variations for images.
- YCbCr. A YCbCr image contains pixels with three samples in the order Y, Cb, Cr. (YCbCr images are often incorrectly termed YUV images.) The Y sample is called the luminance sample; it represents the gray level of each pixel. Cb and Cr are called the chrominance samples. Together with Y, they represent the color of each pixel. An RGB image can have the same high quality as a YCbCr image, but a YCbCr image often requires less disk space.

These image types are supported by the image libraries described in this article.

Using Image Files

In trying to provide a software alternative that makes using images as easy as using text, an application developer is usually faced with the following issues from potential customers.

- I don't have enough disk space to store images.
- If I compress my images, won't they become slow to display?
- How can I clearly display photographs when the screen's resolution is only 5% of the photograph's resolution?
- Aren't there too many computer formats of image files to deal with?

Because of these issues, Hewlett-Packard took a comprehensive look at the imaging requirements of end users and application programmers. From this investigation, the image library project team determined that an imaging solution with the following characteristics would be needed to fulfill the needs of both application developers and end users:

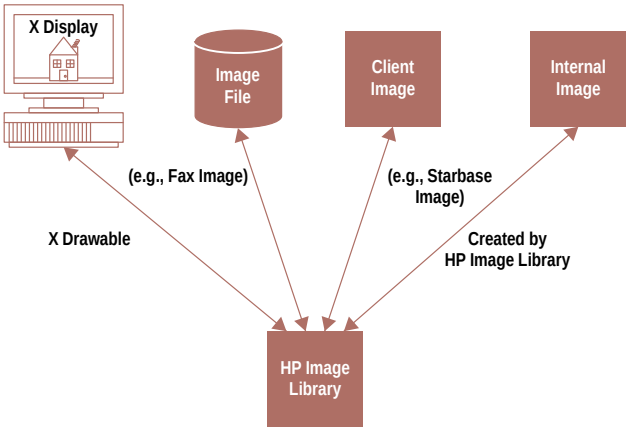

Fig. 1. The four types of images the HP Image Library can read, write, and display.

- Support for industry-standard image file formats
- Minimum disk space requirements for storing images
- Fast, high-quality display of compressed and uncompressed images
- Basic image manipulation, such as rotating and cropping
- Extensibility to allow programmers to add custom image functions and custom image file formats.

The solution we created to help fulfill these needs has three parts: libraries of image functions built on the X Window System, end-user tools built on these libraries, and an image developer's kit with an extensible application interface.

HP Image Libraries

To make imaging functionality a standard capability on HP 9000 Series 700 workstations and Series 800 multiuser systems, the image team decided to create two libraries and build them into the standard run-time environment. These two libraries, the *image library* and the *extensible file support library*, contain high-level functions that C programs can use to access and manipulate images. These two libraries are collectively called the HP Image Library. With the HP Image Library functions, applications can read and write images that exist in four forms:

- A file image. This is an image in a single or multipage file (e.g., a file containing the image of a fax cover page plus one or more other fax pages).
- A client image. This is an image in memory that is created and managed by a client application separately from the HP Image Library functions.
- An internal image. This is an image in memory that is created and managed by the HP Image Library functions.
- X Drawable. An X window image such as an HP terminal window or an X pixmap such as an icon symbol.

Fig. 1 summarizes the relationship between these four forms of images and the HP Image Library.

While the image library functions support TIFF image files, the extensible file support library addresses support of non-TIFF image files. The extensible file support functions operate on image files in an object-oriented approach, independently of the file type. Therefore, if the programmer defines a new type of image file, the existing client programs continue to work on the image file without modification.

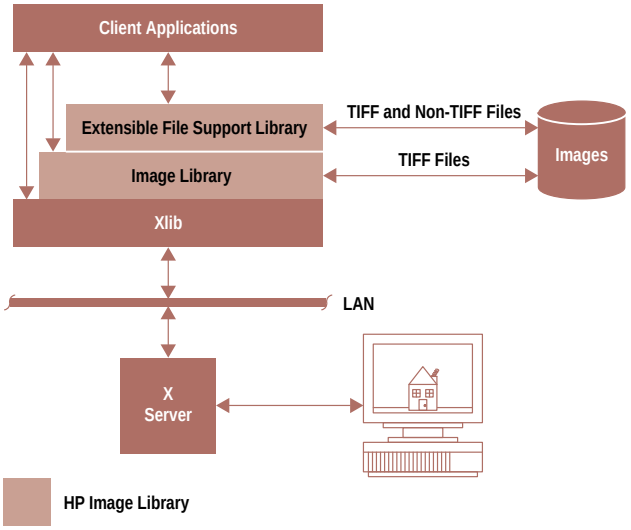

Fig. 2. The image library and extensible file support library in the client/server architecture.

Although the extensible file support library also supports TIFF, that support is not as extensive as the image library support of TIFF. The extensible file support of TIFF files enables applications to treat TIFF files as merely another image file.

The HP Image Library runs on the HP-UX* operating system and uses the X Window System to display images (see Fig. 2). Because the image functionality is layered on standard X, the user gains all of the client/server advantages of X.

Using Pipes to Access Images

To simplify the programming task, the image team designed the image library and the extensible file support functions around the concept of an image pipe. An image pipe is a series of calls to functions, beginning with a producer (a function that reads an image from a source such as an image file) and ending with a consumer (a function that writes an image to a destination such as a display). Fig. 3 shows the steps in a sample image pipe.

While only one producer and usually only one consumer function are allowed per pipe, the pipe can also contain functions that manipulate the image. These filter functions perform operations such as decompressing, scaling, rotating, or converting the image.

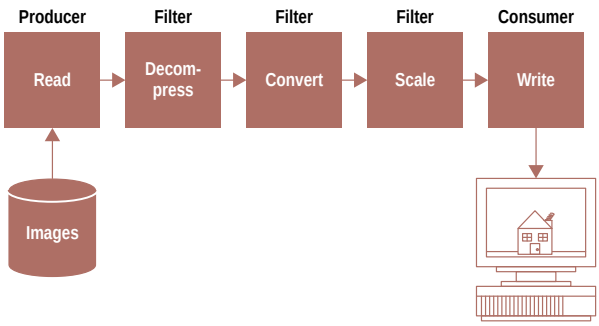

Fig. 3. The steps in an image pipeline.

A producer can access any form of image (Xwd, file image, client image, or internal image). Filters operate on the image data supplied by the producer. The consumer receives the filtered image and writes it to the display, a file, or an internal or client image.

When the pipe is executed, it processes the image in horizontal slices, which are called strips. Processing strips requires smaller buffers between filters and less memory than processing an entire image. Because processing strips can be located in cache memory, image processing performance benefits. For complex image processing, images begin to appear quickly because there is no need to wait for the entire image to be processed before displaying it.

Industry Standard File Types

Since there are several sources for images (e.g., scanners, PCs, fax machines, LAN, and so on) multiple types of image file formats have emerged.

To enable applications to handle these different file types, the HP Image Library provides reading and writing support for the image file types listed in Table I.

The following are definitions for file types not defined earlier.

- LZW. Lempel-Ziv and Welch compression format (described below)
- Xwd. X window (an image created by storing the contents of an X window in a file)
- GIF. A common format for palette images
- JPEG. Joint Photographic Expert Group compression, a lossy compression format
- JFIF. JPEG compression that does not conform to TIFF specifications.

Compression and Decompression

Image compression is a process of storing an image in a way that uses less disk space than is used by the uncompressed image. In developing the compression and decompression routines for the HP Image Library, the image team had to account for three factors:

- The quality of the displayed image
- The speed of decompressing and displaying the image
- The amount of compression required.

The quality of the displayed image partly depends on whether lossless or lossy compression is used. Lossless compression involves techniques that allow the image to be perfectly reconstructed. This method of compression stores the information about the repetitive patterns of pixels rather than storing every pixel. For example, for a sequence of 25 pixels each having a value of 92, lossless compression could store two items: one pixel of value 92, and the information that this pixel repeats 25 times.

Lossy compression is a method that uses different techniques to achieve even lower storage requirements than lossless compression. Lossy compression is recommended for photographic images. For the HP Image Library, lossy compression is based on discrete cosine transform (DCT) techniques. These are numerical techniques that transform complex color or grayscale information into less complicated information. Although the original color or grayscale values are lost, the image normally appears identical when it is decompressed and displayed.

Table I
File Types Supported by the HP Image Library

File Type	Reads	Writes	Typical File Contents	Version
TIFF Types				5.0, 6.0
Bitonal	Yes	Yes	Line art, text	
Grayscale	Yes	Yes†	Black and white photos	
Palette	Yes	Yes	Color photos and X screen dumps	
RGB	Yes	Yes	High-quality color photos	
YCbCr	Yes	Yes	High-quality color photos	
Compressed TIFF Types				5.0, 6.0
PackBits	Yes	Yes	Compression for the bitonal images	
LZW	Yes	Yes	General-purpose compression	
CCITT Group 3	Yes	Yes	Common fax images	
CCITT Group 4	Yes	Yes	Common fax images	
JPEG	Yes	Yes	High-quality color photos	
X Image Types				X11
Xwd	Yes	No	Pixmap image from Xwd (Z format)	
Xbm	Yes	No	Bitonal X bitmap images	
Xpm	Yes	Yes	Color ASCII X pixmap images	
JFIF	Yes	Yes	Color photos and continuous tone images	8-R8
GIF	Yes	No	Xv/Xgif network images, dialup services	87a
Starbase Pixmap	Yes	No	HP Starbase pixmap images	1

† Writing grayscale images is supported only in 8-bit format.

To address user concerns about disk space, the HP Image Library supports a full range of compression and decompression methods. Application programs can access the compression and decompression methods listed in Table II through the HP Image Library.

Except for JPEG and JFIF compression, all compression methods are lossless. Although JPEG is lossy, that loss is normally unnoticeable unless the image is compressed by more than a factor of 20 times.

Typically, lossless compression methods reduce the storage requirements to 50% of the original disk space. However, by

Table II
Compression and Decompression Methods Provided in the HP Image Library

Compression and Decompression Method	Image Types
JPEG	Grayscale, RGB, and YCbCr
JFIF	YCbCr
Group 3	Bitonal
CCITT Group 3	Bitonal
CCITT Group 4	Bitonal
LZW	Bitonal, grayscale, palette, RGB, and YCbCr
PackBits	Grayscale, palette, bitonal

using the lossy JPEG compression technique, a typical compression reduces the storage needed to 5% of the original disk space. However, the image library still displays the photograph with high-quality resolution and no noticeable change in display time.

All the compression methods provide fast display and high-quality image appearance. For JPEG, the fastest display implementation is provided in the version of the HP Image Library supplied with HP MPower. With JPEG, the programmer can choose the amount of compression by trading off the amount of compression with the image quality desired.

Before compressing a color image file with JPEG, the application can save additional space by first converting the file to YCbCr format and subsampling it. Subsampling is a technique that reduces the color information to be compressed, without noticeable change in the image quality. Subsampling is really a form of compression, because it stores fewer pixels than exist in the image. Because the human eye perceives degrees of brightness with much higher resolution than exact shades of color, only a fraction of the chrominance samples are stored. After subsampling, the subsampled bits can be replicated before displaying the image. The replication process is called upsampling. The upsampled image looks identical to the original image.

If the application needs to send files to a fax machine, it can use the CCITT Group 3 and Group 4 compression methods. However, most fax machines support only Group 3 format. The HP Image Library can convert any type of image it supports to any other type it supports. Therefore, a photographic image in RGB format can be converted to CCITT Group 3 format so that the file can be sent to a fax machine.

To compress images, end users must experiment with compression methods on various types of images. The following general guidelines are helpful in determining which compression method to use for reducing the storage requirements of images:

- For photographic images when some of the image detail can be sacrificed, choose JPEG (choosing the desired compression amount) or JFIF.
- For Xwd (X window) screen dumps, or other computer-generated images, use LZW.
- For image files being sent to a fax machine, use CCITT Group 3.
- For fax files to be stored or sent to HP MPower systems, use either CCITT Group 4 (if the content is mostly white space, such as text) or either Group 3 or CCITT Group 3 (if the content is highly detailed).
- For images being ported to various non-HP-UX systems, choose PackBits.

Display Quality

Because the resolution of an image on paper can be 20 to 100 times higher than that of the computer screen, end users are concerned about image display quality. To optimize the appearance of displayed images, the image team included the following capabilities in the HP Image Library:

- Simultaneous display of different color images by sharing a common color palette for 8-plane displays
- Gamma adjustment of colors, such as changing the brightness of an image
- Programmable choices for the type of dithering technique used
- Automatic dithering of images on 8-plane displays and remapping of pixel tones
- Automatic conversion of color images to grayscale format on monochrome displays
- Programmable conversion of bitonal images to display as grayscale images.

Dithering

Dithering is a technique that trades screen resolution for more colors or gray levels. While the resulting image may be more grainy than the undithered version, the contrast between the greater range of colors or gray levels provides improved appearance. Dithering is accomplished by modulating the color values between two adjacent color tones. From a moderate distance, the human eye automatically blends these regions of color together and perceives the average intensity.

Two options are available in the HP Image Library for dithering: error-diffusion dithering, using what is called the Floyd-Steinberg method, and area-based dithering. Compared to error-diffusion dithering, area-based dithering provides a more matted appearance in the image (see Fig. 4a on page 42). However, if speed is the primary concern, an image can be displayed faster by using area-based dithering.

For area-based dithering the HP Image Library performs the following two steps:

1. In one area at a time, it applies an 8-by-8 array of positive and negative values to the color or grayscale values. This step preserves the average color or grayscale level in the area, because the sum of the values in the array is always zero.

In a simplified example, this step might subtract 20 from the value of half the pixels and add 20 to the value of the other half of the pixels. Thus, the average value of the pixels in the area remains the same.

2. The values from step 1 are reduced to a number of values that the screen can display. For example, in a grayscale image values 1 through 256 need to become values 1 through 32 for a monochrome screen. In this example, each value is divided by 8 (256/32). Pixels of gray level 1 through 8 become value 1, values 9 through 16 become 2, and so on.

In error-diffusion dithering, the HP Image Library changes the values of pixels one at a time rather than working on areas of pixels. After assigning a dithered value to one pixel, error diffusion records what the difference was in that pixel's value. That difference is then added to the next pixel's value before it receives its dithered value.

For instance, suppose there are two neighboring pixels with values of 90 and 140. If the pixel with the value 90 receives a dithered value of 100, the difference is 10. So, 10 is added to the next pixel's value of 140, making a pixel of value 150 before it receives its dithered value. Perhaps the 150 is dithered to 170, making the current difference 20. So, 20 is added to the next pixel, and the process repeats.

Error diffusion is the dithering method used by the HP MPower ImageView tool,† whereas area-based dithering is the default dithering method.

Before displaying an image, the HP Image Library checks the characteristics of the display device to determine if dithering is necessary. The HP Image Library automatically dithers the image when displaying an RGB or YCbCr image on a pseudo color display device or displaying a grayscale image on a bitonal device.

For RGB images on a pseudo color device, the HP Image Library uses area-based dithering. For grayscale images on a bitonal device, the library uses error diffusion. When a grayscale image is displayed on a bitonal device, the HP Image Library dithers it to a bitonal image.

By default, the HP Image Library chooses the dithering method by using the following criteria:

- If the screen is a pseudo color (palette) image, area-based dithering is used. The image is first converted to RGB and then to palette.
- If the screen is monochrome (bitonal), error diffusion is used. The image is first converted to grayscale, then to bitonal.

Image Conversion for Space and Readability

For reasons such as saving space and improving the readability of text-oriented images, an application can convert images from one image type to any other image type.

Space and Color Images. The high quality of a 24-bit RGB image is only visible on a 24-plane system. When displayed on an 8-plane display system, an RGB image is automatically converted to palette. However, an uncompressed 24-bit RGB image occupies more space than it would as a palette image.

To save space on 8-plane display systems, the application can convert the 24-bit RGB images to palette or YCbCr format. YCbCr format is preferable if JPEG compression is also used. Beyond the additional compression possible, the image space required can also be reduced through the subsampling technique.

Space and Text Images. Concerning monochrome images, the programmer can save space by converting grayscale images of text to bitonal images. A grayscale image requires eight times the space required by a bitonal image.

† ImageView is a tool for displaying, manipulating, saving, and printing images of different image types.

HP Image Library Scaling Functions

The HP Image Library scaling capability performs three types of scaling according to which option is used: scale to gray, area-sample scaling, and simple scaling. In each case, the scaling algorithm accounts for the type of image involved, whether any image type conversion is needed, and whether the program is requesting that the image be scaled up or down.

Simple Scaling

Simple scaling gives the fastest scaling performance but the lowest image quality. This method uses pixel replication if the image is enlarged and pixel decimation if the image is being reduced. No image conversion is performed.

Area-Sample Scaling

Area-sample scaling gives the highest-quality scaling results. This method only applies to scaling an image down in the current release of the HP Image Library. In scaling the image down, sample scaling uses area sampling techniques based on the image type:

- For a color palette image, the image is first converted to an RGB image. The RGB image is scaled by area sampling.
- For bitonal images, the image is temporarily converted to an averaged grayscale type and then the image's pixels are set to on or off, based on a threshold value set by the application. (High-threshold values darken the image and low-threshold values lighten it.) The resulting grayscale image is then scaled by area sampling.

Bitonal-to-Gray Scaling

Bitonal-to-gray scaling applies only when scaling down bitonal images. It converts bitonal images to grayscale and uses area-sampling techniques to scale the image down. To convert a bitonal image to a grayscale image, the HP Image Library assigns each black or white pixel a gray level between 0 and 255.

While a scanned photographic image should be a grayscale image to retain the different levels of gray, a scanned text-oriented document or other text image can be bitonal and still be read from the screen by users.

Readability and Text Images. When a grayscale image of text is stored as bitonal, it can be converted back to grayscale if it was scaled down for display. The result is text that is easier to read. Thus, the image can be stored as a bitonal image, using less space, and displayed as a grayscale image.

Image Manipulation

To manipulate images, the HP Image Library includes a number of functions for scaling, rotating, mirroring, cropping, and changing image colors.

Scaling Images. The scaling function maps one image into an image of a different resolution. It allows the application to account for the differences in resolution between the display screen and the image captured by devices such as scanners. It also provides a mechanism to resize images to different window resolutions (see "HP Image Library Scaling Functions," above).

An application can scale an image up or down. For example, scanners are typically 300 dpi and display monitors are only 90 dpi. Therefore, images scanned at 100 dpi or higher resolution must be scaled down to the screen resolution.

The scaling function includes options for faster scaling or high-quality scaling. Faster scaling replicates or removes pixels, depending on whether the image is being scaled up or down. High-quality scaling converts bitonal images to

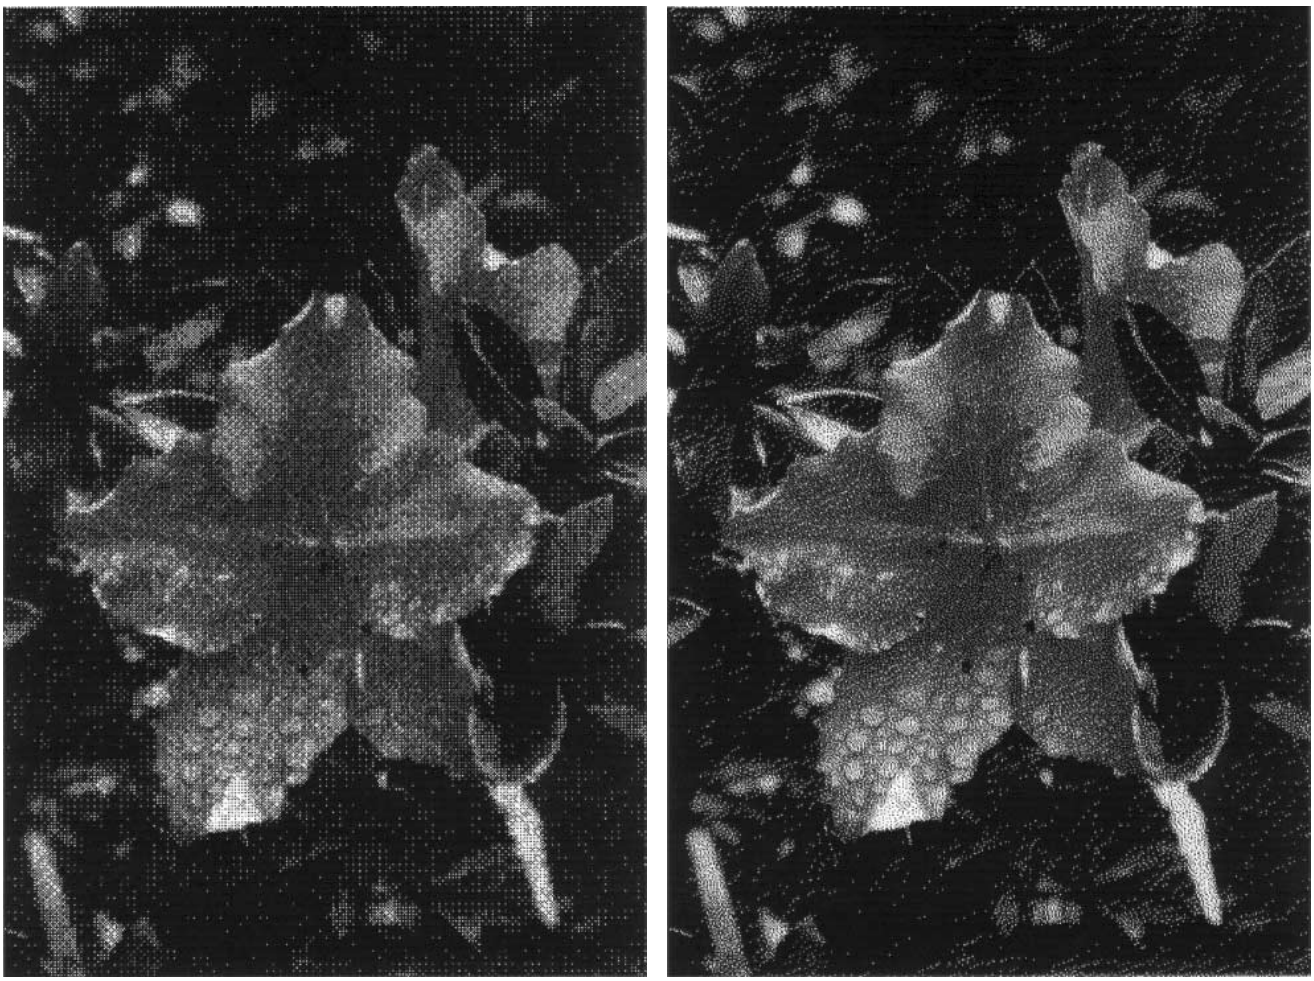

(a) (b)

Fig. 4. A comparison between (a) area-based dithering and (b) error-diffusion dithering.

grayscale or scales color images by area sampling. Area sampling replicates pixels based on the average color values of an area of pixels, producing a clearer image.

Rotating and Mirroring Images. The rotation function rotates an image at integer multiples of 90 degrees. The rotation can be clockwise or counterclockwise. If the image is rotated by 90 or 270 degrees, the width and height of the image are reversed. Otherwise, the image retains the same dimensions as before rotation. The mirroring function mirrors the image about the x or y axis.

Cropping Images. The cropping function extracts a rectangular section of an image, creating a cropped image that is either the same size or smaller than the original image. An application might combine the scale and crop functions to implement features such as panning and zooming.

Changing Image Colors. The color mapping function changes the image colors by mapping the color values of the image pixels into different color values. This function can be used by applications to change the colors, brightness, or contrast of the image.

Custom File Types and Custom Functions

To allow programmers to extend the HP Image Library, functions are included for defining new types of files and creating custom functions.

Custom File Types. For unsupported types of image files, the programmer can extend the extensible file support library. The programmer defines a new type of file using an extensible file support function and then rebuilds the extensible file support library.

Thus, if an existing application accesses files through the extensible file support library, the newly defined file type can be accessed as just another extensible file support file. This object-oriented approach allows the programmer to create routines without worrying about what type of file is being manipulated.

Custom Image Functions. For capabilities not available in the HP Image Library, the application can define a new function. An application-defined function (or custom function) can be a producer, a filter, or a consumer function.

Fig. 5. An HP ImageView screen.

Application-defined producers and consumers might be used to read or write to an image-capable device. The application-defined producer outputs a pipe image, ideally in strips, that may have an input image type that is unknown to the image library but known to the application-defined consumer. The application-defined consumer can be defined to accept a standard or nonstandard input image type that is unknown to the HP Image Library.

An application-defined filter can be created to operate on a client image that is not in a format supported by the HP Image Library. For example, the client image can be created to support a color image in CMYK (cyan, magenta, yellow, and black) format. While standard HP Image Library functions can read and write this client image, operations to manipulate the image require an application-defined function.

End-User Image Tools

Based on the image and extensible file support libraries, end-user tools have been created by the image team and several other HP teams. For example, the HP online help facility, and the HP MPower components fax, DeskScan/UX, Whiteboard, ImageView, and HP SharedPrint use these libraries for image display and manipulation.

A central part of the HP MPower image display is ImageView, an OSF/Motif-based image display application (see Fig. 5). A basic version of this application is built into the standard run-time applications on HP 9000 Series 700 and 800 systems.

As a client application of the image libraries, ImageView displays all the file types supported by the extensible file support library. The user displays an image by double-clicking on a file icon in HP VUE. The user can then zoom in on arbitrary areas of specific interest and resize images by dragging a corner of the window.

In the HP MPower version of ImageView, users can also compress an image file, print images, adjust contrast, brightness, and orientation, and save those changes. Also, users can display images without dithering and fix the image scale during display.

Image Developer's Kit

The image team created an image developer's kit that programmers can use to create applications built on the HP Image Library. The applications developed by programmers can be clients of the X server, or they can work solely with image files without using X. Most applications access both image files and the X server because it is the only means for displaying images.

The toolkit can be installed on HP 9000 Series 700 systems and consists of the following components:

- Header files for the HP Image Library functions
- A range of sample image files, such as TIFF, GIF, and Xpm files
- Man pages that describe all HP Image Library function calls
- Source files for sample applications that contain calls to the HP Image Library functions
- Makefiles that convert sample source files into executable programs
- Source files that show how to extend the HP Image Library, adding support for other types of files
- A utility called imageutil that provides command line options for image viewing and manipulation.

Acknowledgments

The authors would like to thank the image library engineering team for their valuable contributions to the image project. Enrique Santos and Melinda Shebell led the two phases of the product development. Peter Kaczowka designed the overall architecture of the image library and the extended file support library. Derek Larsson created functions that manipulate, dither, and display images and ImageView enhancements in HP MPower and he helped with creating this article. Lynn Miller and Joel Stave developed the original ImageView application. Peter Voegelin developed compression and decompression routines, and Noelie Alito developed routines for reading and writing TIFF images. Daniel Bass developed compression and decompression routines including JPEG routines.