

Introduction to image technology

by R. M. Helms

Business today is wrestling with mountains of paper that must be moved, filed, located, and moved again from person to person. Often the paper must be stored for extended periods of time, sometimes as long as seven or more years. The long-term storage of paper records is becoming more and more costly. An image system not only makes the document capture, retention, and retrieval process more cost efficient, but also makes it a faster service to the users. The purpose of this paper is to explain the basic concepts of image processing in business.

Image processing in business is used to replace paper documents with an electronic facsimile or likeness, and then store, retrieve, display, print, and distribute the electronic image as we have done with other media since the practice of keeping business records began. Archaeologists have discovered warehouses full of dinner-roll-shaped clay invoices and bills of sale from diggings in ancient Babylonia, representing a 5000-year record retention.

Image as a form of computer data is new to most computer installations. This paper presents an introduction to image technology and briefly discusses some of the effects it has on computer systems. The focus in this paper is on image as it applies to business applications. First, the concept of image is introduced. This is followed by an introduction to the use of image in business.

An image is a collection of *picture elements* that transform the original picture into an electronic image. Picture elements are analogous to the tesserae of which a mosaic is composed. Another way of thinking of image is that a digital image is a page of numbers usually arrayed as a matrix, wherein the position of the number represents the position of the picture element, and the value represents its brightness or amplitude. The position of each picture element is its location within a grid pattern, as expressed by the two coordinates x and y . An array of picture elements is shown in Figure 1. The picture elements are also known as *pixels* or *pels*.

Each element of an image can have only one discrete numeric position and brightness value. Thus we call the image a *digital image*. The picture on a computer display is composed of these picture elements. The eyes cannot resolve the square or circular picture elements from a distance, but close examination of the display reveals these image building blocks. These circles or squares are the visual manifestation of the picture elements, which shortened form the

© Copyright 1990 by International Business Machines Corporation. Copying in printed form for private use is permitted without payment of royalty provided that (1) each reproduction is done without alteration and (2) the *Journal* reference and IBM copyright notice are included on the first page. The title and abstract, but no other portions, of this paper may be copied or distributed royalty free without further permission by computer-based and other information-service systems. Permission to *republish* any other portion of this paper must be obtained from the Editor.

Figure 1 Picture elements





Figure 2 Evolution in information processing



word *pixel*. Pixel, in turn, became *pel* because some older computer languages required the names of variables to be less than five characters long. Although some people have noted differences, they are not commonly accepted. Thus, all three terms—picture element, pixel, and pel—are interchangeable in this paper. Another term often used for pixel is “dot.” Dot however, is reserved for use with printer pels.

The way an image is built up by a series of rows of pixels is also called *raster format*, a term that comes down to us from the scanning pattern of the cathode-ray tube. Raster also refers to the pattern of scan lines that constitute a television display; these are drawn by an electron beam as an array of closely spaced parallel lines called *raster lines*.

Historical background. One of the first uses of digital image was the Bartlane Cable Picture Transmission System introduced in the early 1920s. This system connected London, England, with New York and could transmit a newspaper photograph in three hours with five shades of gray.

The digital computer made many things possible. The general progression of applications is suggested by Figure 2. During this whole period, image processing made substantial progress, even though any given period may have been dominated by another application. For example, the Jet Propulsion Laboratory (Pasadena, California) in 1964 used image processing to correct and improve pictures of the moon sent to earth by Ranger 7. This required a large, fast, and expensive computer with a very sophisticated display. These image devices were usually complex and expensive.

Since that time, several developments have converged to simplify, reduce cost, and make these services generally more accessible. The most dramatic development is the introduction of the personal computer. With the PC, users had access to a system with large amounts of memory, fast processing, and mass storage. Another line of convergence is the television industry, which has undergone its own evolution by way of the invention of home video systems. Television cameras have decreased in price and increased in quality. We also note such other technological advances as optical storage, new high-resolution gray and color displays, high-speed processors, parallel processors, and floating-point processors. These and many other advances have created the situation wherein image is no longer used only in the world of science. Image is now cost effective in business. The focus of this paper is the use of image technology in modern business. Here we look at the technological concepts needed for line-of-business image as well as such other applications as a database of photographs.

It is interesting to contemplate the development of communication of a person and the evolution of computers. When a baby is born, the first form of

data acquired and learned is image—the mother's facial expression and tone of her voice, a visual and aural image of the baby's mother. Next toddlers learn to draw graphics. Text follows, and numbers are last. Computer evolution has been just the opposite: numbers, text, graphics, and finally image.

Image versus graphics. In today's world of super special effects in science fiction movies, it is becoming difficult to separate image from graphics. There is, however, a basic distinction: images are captured, whereas graphics are created.

If a user begins with a capture device such as a scanner and "takes a picture" into the computer, the result is an image. Graphics are drawn, a simple example being a bar chart. A good analogy for illustrating the difference is the difference between a painting and a photograph. A painting is a graphics design because, even though it can appear very realistic, it has been created. A photograph, on the other hand, is the captured image of reality. Even if a photograph were to be shot through a special-effects lens or filter and even if it does not look realistic, it still is a photograph—an image. Thus we see the basic distinction between graphics and image.

Another key point is that images are stored electronically in raster lines of pixels, and graphics are stored as series of commands. However, graphics can be converted to *artificial images* through filling and shading. The word "artificial" refers not to whether it is a true image but to its origin. *Synthetic image* is another term used for artificial image. However, once created, an image's origin cannot be determined. This fact has no effect on the principles of image technology as discussed in this paper.

A look of realism does not guarantee a true image; artificial images can be extremely realistic. Many of today's graphics pictures are actually too realistic. Usually, reality is not as perfect as computer-generated graphics appear, which is the usual tip-off to an artificial image that is created graphically, rather than an image that was captured. It is common practice now that special effects for movies are often created in a computer as artificial images, rather than built as models and photographed.

Noncoded information. Image is one form of noncoded data. In coded data, each unit of information, often a byte or word, has meaning. In noncoded data, there is no unit of information. The bytes represent only light levels on an area. Digital voice

or audio, such as on compact discs, is another example of noncoded data. Graphics would be an example of coded data because the unit of information is a graphic command. The usual sources of

Up to 95 percent of all information used in an enterprise is noncoded.

images in an enterprise are the writing or printing on pieces of paper, photographs, or three-dimensional scenes.

Up to 95 percent of all information used in an enterprise is noncoded or, said another way, only 5 percent of the information in many businesses is coded in nature. Phone calls, handwritten communications, letters, and such are potential sources of noncoded data.

Today only 5 percent of business data is coded into a computer. That is, only 5 out of 100 items of data are stored. Each of the remaining 95 pieces takes about 10 times as much space per piece because of its noncoded nature. Thus it takes about 200 times as much space to store that 95 percent as the original 5 percent. That is, 20 times as many pieces require 10 times more space per piece.

A scanned page of text is a picture of that text, a fact that becomes important in two areas. The storage requirement of an image is about 10 times as great as that for coded information. More importantly, if you wish to revise the text, you cannot use a text editor for an image. Text editors are built to work on coded text where each byte represents a character. On the other hand, an image of text is merely a picture. Optical character recognition (OCR) is the process of converting a picture of the image back to the ASCII or EBCDIC byte that represents it. When writing a document with a word processor, you may use OCR to convert a printed paragraph back to ASCII, so that it can be imbedded in the document. Also, once the page has been converted back into bytes, a search of the contents can be performed. An image cannot be searched, except by the human visual process.

Vision. To complete our discussion of image, we look at another term used in image technology—*vision*. Vision means not only to see but also to understand the contents of an image. For example, television is an image device in that it captures, stores, and displays image. A robot, on the other hand, may use a television camera and a vision system to locate a part in an assembly line. This locating operation is broken down into two steps. First, the image is captured into the computer using the camera. Second, the image is analyzed to under-

An important use of vision in manufacturing is defect detection.

stand its contents in order to locate a desired part. This is vision in the form of image processing, and as such it can be used in the office and in manufacturing.

An image of a printed page is simply a picture of that page. Although it may look like text, the text cannot be edited using a word processor. To do this you must convert the picture of the text into actual coded text, which may be done by optical character recognition (OCR). OCR is another manifestation of the understanding of an image. Here a page of typed or written text is scanned, then the OCR function recognizes the images of the characters and converts them back into a coded representation. The result of this process is an image of the page and a text file of the content.

Vision is used to examine the contents of an image and to extract features. In a manufacturing line, a vision system may be used to locate a part on the assembly line, so that a robot arm can pick it up. Locating solder joints on a circuit board so that a probe can be placed is another common use of vision. An important use of vision in manufacturing is defect detection. A vision system may be used to examine a solder joint and look for cracks. These types of jobs are good uses for vision systems because the repetition of the job makes them very tedious for people to do on a continuing basis.

Image systems

There are a variety of reasons for an enterprise to buy an image system. One of the most common is to reduce the cost of handling paper. This reduced cost comes from several factors. There is no need to file paper documents, and this saves the cost of filing space, the cost to perform the filing, and the cost of retrieving the paper. Another cost-saving factor is the avoided cost of locating lost paper. When a folder of paper is in use, no one else can use it. The folder can be misfiled when it is returned, thereby further confusing the matter.


A benefit of an image system is speed of retrieval. Active documents that are not in archive can be found in seconds. When talking to a client on the phone, speedy document retrieval can make the difference between a quick answer and playing telephone tag, while trying to get back with the needed information. Speed can also be a factor in sending a document to a remote location. The FAX machine has affected the overnight courier business by providing electronic delivery in seconds. This applies to image systems also. Local communications in image systems are very fast; remote access is often available at a longer response time. This difference in time is a result of the speed of the communications line.

Image activities. The basic image activities include capture, display or print, file, and communicate. Advanced image activities include edit, vision, enhance, and analyze. A basic business image system is illustrated in Figure 3.

Capture. The process begins with the capture of the image. To capture an image is to get the data into a computer from the original source document, photograph, or other hard copy. The two basic types of image capture are scanning and video digitizing.

Scanning. In the office, scanners are perhaps the most common way to capture images. A scanner looks similar to a photocopier. A piece of paper is inserted into the machine, and the image is scanned. The key distinction here is that the original object is a piece of paper. This paper could be a photograph or a print on plastic, but it cannot be a three-dimensional object. This same restriction applies to most photocopiers. The scanner speed varies greatly with the model and the kind of image the scanner is building. A scanner that can work in both bilevel and gray usually runs much slower in gray because eight-times as much data must be sent to the com-

Figure 3 Typical image system




puter. The communications connection is another speed limiter. A scanner like the one used in the IBM ImagePlus™ system can work at one page per second, whereas a lower-cost scanner, such as the IBM PageScanner™ device, takes about thirty seconds for the same page and image. Color scanners require a three-fold increase in the amount of data recorded compared to a gray scanner.

Video digitizing. Another approach is to use a video camera attached to a video digitizer. The video camera creates a television signal that the video digitizer then converts to digital information in the computer. Video lends itself to three-dimensional originals. There are many types of video digitizer cards, which vary by resolution and number of shades of gray or colors. A slow card may take up to eight seconds to do a capture. This is called a slow-scan video digitizer. A real-time video digitizer can capture an image in one-thirtieth of a second. It is important to

use a real-time digitizer when taking pictures of people, because a slow scan unit requires that a person remain still for as long as eight seconds. There are two popular resolutions in videos, 640 by 480 and 320 by 240, wherein the unit of the first number is pixels per line and the second number is the number of lines in the image. A full, high-quality video signal can give a 640 by 480 image, which is called a *full-frame* of video. A frame of video is made up of two fields, each of 240 lines. A field grabber will take a 320 by 240 image in one-sixtieth of a second.

Other capture devices. Other image-capture devices not often thought of as such include satellites, medical scanners, and seismic readers. Images derived from infrared, X-rays, and ultraviolet radiation are outside human visual capabilities. One purpose of an *image display station* is to process and show these invisible images to the user.

Figure 4 Overlapping pels



FAX. A FAX (or facsimile) machine may be compared to a pair of photocopiers connected by a telephone line. The original flat document is fed through one machine, which captures the image of the page. The image data are then sent via the telephone line to the other machine, which prints the image. One of the major factors leading to the success of the FAX machine is that there is one standard way of describing the data. Almost all FAX machines conform to the International Telegraph and Telephone Consultative Committee (CCITT) FAX standards.

A FAX modem is used to connect to the phone lines. A FAX modem is unique in that the modem and system transmit and receive only images and do not communicate with regular data modems. The text that the FAX system prints on the page is not true text but an image of the text. For a computer to send text to a FAX machine, the computer must change the text into image, i.e., it must rasterize the text.

The FAX machine has become such a widely used image device because almost all FAX machines talk using the CCITT Group 3 standard. This makes it possible for almost all brands of machines to exchange images. FAX cards for personal computers are also quite common. With such a card, a FAX machine can become a remote scanner/printer. Because FAX machines talk only image, performance is a major limiting factor in extending the usage further.

Image terms

Resolution and picture elements. The term *resolution* is used in the discussion of video to refer to the dimensions of the image in pels per line and by the number of lines (e.g., 640 by 480). This is one form of resolution. Many definitions of resolution exist, depending on the way it is used. It is very important

to understand exactly how the term is used to determine the meaning. In principle, all definitions of resolution are based on the ability of a system to distinguish closely spaced lines or points. A telescope that cannot distinguish a double star is said to be unable to resolve a double star, for example. In video, we use resolution to mean the number of pels in the width and the number of pels in the length of an image, independent of the size of the television screen. This gives us the total number of pixels in an image but does not tell us the size of the original. Thus, the greater the number of pixels per image, the higher is the resolution of that video system.

Another definition of resolution is pels per unit of measure, such as 200 pels per inch (ppi). With this definition, if we know that the image is 200 ppi and the image has a width of 1000 pels, we can calculate the width of the original to be 5 inches. When resolution is expressed in this way, we also need to know the maximum number of pels in each direction or the maximum size of the original. This approach is most common with scanners. An example is the IBM PageScanner device that supports 300 ppi over an 8.5-inch by 11.7-inch area. Here it is considered important to preserve the size so that letters on a page are reproduced at the same size. When reading specifications on scanners from countries that use metric measurements, resolution may also be given in microns per pel. Occasionally we see the number of bits per pel or number of shades of gray also called resolution. This, however, is not generally accepted terminology. Also, dots per inch (dpi) is used to describe the resolution of printers.

To *enhance* an image is to make it more useful. An example is to sharpen the focus on an image to bring out hidden details. If an image has poor focus or is blurred due to motion during the capture, image enhancement can be used to bring the picture back to a more useful state. Sharpening the focus is very useful when looking for hidden details but any noise in the image will be brought out also. Softening the image removes noise at the sacrifice of detail. Often combinations of techniques are used, such as a median filter to remove noise followed by a high pass filter to enhance the sharpness or focus of the image.

One last comment on resolution concerns a distinction between addressable and discrete resolution. Consider the pixels illustrated in Figure 4. With 50 percent overlapping pels, if this were a 200-ppi pattern, the *addressable resolution* is 200 ppi, whereas the *discrete resolution* is only 100 ppi. This difference

in resolution between addressable and discrete is the result of the pel overlap. Discrete resolution refers to the number of pels per inch as defined by the diameter of the pel. In this example a pel is 0.01 inch.

Almost all specifications quote only the addressable resolution. The discrete resolution can be calculated from the pel diameter, if quoted, and from this we can calculate the amount of pel overlap. A device with high addressable resolution but low discrete resolution shows smooth edges on lines, but it cannot produce fine detail.

The *aspect ratio* of a pel is the ratio of width to height. If the width is the same as the height, that is, the width-to-height ratio is 1, then the pel is said to be a *square pel*.

A *square pel* is very desirable. Both the geometric square and circle fall within the square-pel definition. If a pel is not square, a correction must be made when the image is rotated in order for the image to not appear as squashed or skewed. Most printers, displays, and capture devices now give a square pel. However, a square pel does not mean that the width and height of the total image are the same. As an example, a video image with 640 by 480 pels has a square pel, but the overall image is not square.


Although many people have heard of a pixel, few have heard of a *three-dimensional pixel (3D pixel)*. A 3D pixel is one with not only width and height but also depth. We usually think of a pixel as a square, but a 3D pixel is a cube of image data. Such a pixel is created by an underground scanner, which is used when looking for oil, or by a body scanner when doing medical analysis. A 3D pixel is termed a *voxel*, meaning volume picture element.

Types of images. Three major types of image are dealt with in this paper: bilevel or two-level image, gray, and color.

Bilevel image. A black-and-white image is the most common and simplest example of a bilevel image, although any two-color image falls within the definition. A newspaper is an excellent example of a bilevel image. Each area of a page is either black ink (assuming the newspaper is not printed in color) or white paper. In image terms, each pel is either black or white.

Photographs in a newspaper appear to be composed of areas of gray. This is an illusion of the eye. If a

Figure 5 Halftoning



Photograph by Gordon Giddings with permission

newspaper photograph is examined under a magnifying glass, the arrays of black dots of which the photograph is composed become visible. When these dots are sufficiently fine and tightly packed so that the eye cannot see (resolve) the individual black pels, the eye averages the areas to shades of gray. This is called *halftoning* and is the basis of most image printing technology. Halftoning is illustrated in Figure 5.

Most bilevel image original documents are pieces of paper with handwriting, typing, or printing on them. When working in bilevel, the color of the ink and the paper are ignored, with dark blue and black ink both showing up as black pels on a white field. This is the same principle as is used by a photocopier. Also, the bilevel image does not truly indicate black ink on white paper. The bilevel image could have been blue ink on pink paper. Therefore we refer to "bilevel," rather than "black-and-white" images.

In FAX, most documents for the office are printed, typed, or drawn with black lines on a white paper background. Because the FAX system is looking for dark ink, a 1 (true) is used for ink and a 0 (false) is used for paper. No indication of the true color of the ink or paper is given, just that the ink was dark and the paper light. To this day, these are the standards in bilevel image. As will be noted later in this paper, this is the opposite of the gray and color standards, in which 0 means black and 1 means white.

Gray image. Black-and-white television is classified as a gray visual display device. Each pel on the screen is a shade of gray, not black or white as in halftoning.

Most gray systems use 256 shades of gray as their granularity or shade resolution. This is due both to the historical use of a byte per pel and to the fact that the untrained human eye cannot see more than 256 shades of gray. (Actually most persons cannot distinguish more than 128 shades of gray.) Medical and scientific imaging systems tend to need more levels, not because scientists can necessarily see more, but because—by changing the way people look at the data—they can sense details that are beyond human vision. Although a person looking at a 256-shade image and a 1024-shade image probably cannot detect any difference, the computer can artificially change the intensity scale to focus attention on part of the scale and make tiny changes perceivable.

Both gray and color images are referred to as *continuous-tone images*, because most gray and color images are derived by the digital sampling of continuous (analog) information. Most gray images start as photographs or 3D objects.

As mentioned previously, bilevel image is based on a FAX-and-scanning-paper model. Thus 0 (false) means paper (white) and 1 (true) means ink (black). Gray image is based on television, therefore 0 (false) is the absence of light (black), and 1 (true) is the presence of light (white). To avoid the confusion of 255 being white on an 8-bit scale and 1023 being white on a 12-bit scale, gray image work often refers to 0 to 1 with fractional values for the gray scale.

People see neither colors nor gray-scale information equally. This knowledge is invaluable when assessing priorities in the design of image systems. For example, changes in dark gray are perceived more easily than changes in light gray. If you view a scale with equal steps of gray as measured scientifically, your mind will tell you the dark shades are changing more quickly than the bright shades. This creates an interesting problem in that most image-capture devices break down gray to equal steps of reflected light. If you look at a scale of gray values and see it as equal steps, you are viewing a gray logarithmic scale. A logarithmic scale changes black slowly. The rate of change from black to white accelerates in these devices as they measure more and more white.

It is a basic problem that computers process image data differently than the brain and eye, a fact that

has to be accounted for in a system for creating realistic-looking images. It so happens that perception by all the senses is logarithmic. We can see in

**The purpose of halftoning is
to create the illusion of gray
by means of bilevel pel patterns.**

very dim light and hear the brush of grass against a foot. At the same time, we are not overwhelmed by desert light or thunder, both of which on a linear scale would be painful.

Halftoning. The purpose of halftoning is to create the illusion of gray by means of bilevel pel patterns. Halftoning is often used with printers, but it can also be used with bilevel displays. Pictures in newspapers and magazines, as mentioned earlier, are printed using the halftone process.

As an example, assume a matrix of 4 by 4 pels, which is used to create the illusion of gray shades. To show solid black, all pels within the matrix are black. To show dark gray, most pels are black. For light gray, most pels are white, and for white, of course, all the pels are white. Figure 5 shows a very simple form of halftoning. There are many more advanced techniques that give superior results. However, all are based on the same basic principle that the more black pels in an area, the darker it seems.

Discrete resolution plays a significant part in halftoning. To create the illusion of gray, the patterns of black pels should be small enough that the human eye cannot resolve them individually. An excellent example of this is black-and-white photography. Most people assume that a photographic print shows continuous shades of gray. Actually, it does not. A black-and-white photograph is a collection of metallic silver grains on white paper. These grains are very small and spaced in random patterns. They become obvious only when a negative is enlarged greatly, and the visible grains are referred to as photographic grain. High-speed film uses larger grain size in part to achieve the higher speed.

Color image. Except for filling in areas with colors of a single hue, the standard color process is known as the red-green-blue (RGB) process. This process mimics the three color-sensitivity bands of the human eye. The RGB process is used for all color imaging applications: photography, printing, televi-

**The purpose of compression is to
reduce the size of the data used to
represent the image.**

sion, and other forms of color imaging. The color printing process is a four-color process because black is added and, for technical reasons, printer colors are the complements of RGB: magenta (RGB minus G), cyan (RGB minus R), and yellow (RGB minus B). The underlying principle, however, is the same. One can observe this process by looking closely at a television screen where one can see the red, green, and blue dots. Similar to gray images, 256 intensities for each of the red, green, and blue are available. The red, green, and blue are called bands of the image, and all color images have contributions from each of the three bands. Gray and bilevel images have one band each. The blending of the intensities of RGB gives the various colors. All three bands at full brightness are seen as white.

As in gray, the medical and scientific communities often need 12 bits per pel per band rather than eight. With eight bits per pel per band, there are 16 777 216 possible colors, and 12 bits yields over one billion colors. Human perception is even poorer for color than it is for gray. The untrained eye can distinguish only about 128 hues of green, 64 hues of red, and 16 hues of blue. Whereas the exact number of hues per band can be argued, these are the correct ranges. This limitation on the ability of humans to discern color is taken into account in image processing, because it is what is perceived that is important.

As with gray, most color-image originals are photographs and three-dimensional objects. Color images are captured by scanning the original three times,

once each through red, green, and blue filters. The actual capture may be done in one pass using a beam splitter or three sensors, but the process is the same.

If we see a leaf on television and the leaf is not the same green color as the real leaf on the tree, we cannot know this until we compare the television leaf with the real leaf. The reason is that the eye has no point of reference with respect to the leaf, and the brain assumes that the color as viewed is correct. This would not work if the televised leaf were too unrealistic in color. However, the eye is not easily fooled with an inaccurate flesh tone, unless it mistakenly assumes that too red a face, for example, represents anger or embarrassment. Because the eye can detect a correct flesh tone in a face without having to see the real face, the flesh tone is a reference color. The television industry goes to great lengths to ensure that flesh tones are correct. Attention to reference colors in an image product that works with color is critical to human perception.

System considerations

Images are constructed from large amounts of data, a fact that has some profound impacts on system design.

Compression. The purpose of *compression* is to reduce the size of the data used to represent a digital image for purposes of storage or communication. Compression does not change the physical size or detail of the resulting image. That is, a scanned 3-by 5-inch original after compression, transmission, and revisualization, results in a 3- by 5-inch output digital image. The idea of compression is to reduce the amount of data used to represent the image primarily by techniques that reduce redundant information. To make an image physically smaller is called *reduction*; to make an image physically larger is called *enlargement*.

There are many different types of compression schemes, and these are tied to the type of image being compressed. They also assume a model of the characteristics of the image. An example is the compression used for FAX, which assumes a scanning rate of 200 ppi on a page that is mostly white with black lines and characters. This sounds quite reasonable until we consider the problem of compressing halftoned photographs, which consist of short runs of white and black. The use of FAX compression on a halftoned image can create more data than are present in the uncompressed data image. Codes are

used to represent the sequences of white and of black. These codes are of variable lengths, with very short codes given to very common things and very long codes given to uncommon things.

Lossless versus lossy systems. At present, there are two classes of compression: lossless and lossy. In *lossless compression* the decompressed image is ex-

The choice requires tradeoffs between fidelity and transmission efficiency.

actly the same as the original image. In *lossy compression*, the decompressed image is almost the same. The purpose of a lossy compression is to minimize the amount of data while maintaining visual integrity. Lossless compression is better when high visual fidelity is required, and lossy is better when more compression is desired. Often it is not necessary for the image to be exactly the same as the object. The choice requires tradeoffs between fidelity and transmission efficiency. Perhaps all that is needed is to be able to read the text in the image, not to study the exact configurations of the letters. Lossy is most common in gray and color compression. An exact compression (another way to say lossless compression) of color may yield only a 2-to-1 or 3-to-1 reduction in quality of image data, whereas a lossy compression, with loss that is beyond most human perception, can yield a reduction of 20 to 1. When we look at the amount of data we deal with in color, this can make a significant difference in storage or communications requirements.

Target of the image. When choosing between lossy and lossless compression, the target of the image becomes a key consideration. If the target is a person who will casually view a picture, a lossy-compression method with its reduced data size is preferred. On the other hand, if the person is a doctor looking at a digital X-ray, the highest possible fidelity is vital to the interests of the patient.

When defining a pixel, we noted that one's eyes do not notice the squares. This is an extremely important concept. Human vision is the key to most business image processing. It is not what is there that is important, but what people perceive as being there. This is a generalization that has exceptions, but it is still a very useful concept. Human perception has an effect on how images are displayed, printed, compressed, and enhanced. Medical imaging uses computers to enhance digital X-rays to help locate areas of interest, thereby enhancing the resulting human perception of the image.

Machine vision is an exception to the rule of human perception. In machine vision, the target of the image is a machine rather than a human being. Machines do not have the perception weaknesses of a human being, and machines can locate things in an image that are beyond human perception.

Compressed image sizes. Certain rules of thumb exist for storing the image of a page. One thing that is very constant is that images require more storage to represent a page than the coded version of the same page. A page of printed text, stored as text, requires about 3K bytes of storage. That same page stored in image form requires about 25K bytes of storage. This is an eight-fold increase in size. Also the image format cannot be interrogated as can coded text. An 8.5- by 11-inch bilevel business letter at 200 ppi compresses to about 25 kilobytes (KB) of storage per page. At 300 ppi that same page requires about 40 KB of storage. This assumes a page with a few paragraphs of typing. A very complex page, such as an insurance form, requires at least double the storage required for straight text. In engineering work, a page of 33 by 44 inches is known as an E-size drawing. With this increased physical size comes an increase in compressed-data size. Data for an E-size page can range from 250 KB to 500 KB. These are not hard and fast rules, but rules of thumb.

Size and complexity of the original object contribute to the size of the compressed image. Compressing gray and color photographs is also affected by the content. For a video-resolution image (640 by 480) of a face with lossy compression and excellent visual integrity, the size is 25 KB in gray and 40 KB in color. Faces are relatively simple to compress. However, a busy picture like a house with shrubbery and many details might require 50 KB in gray and 80 KB in color.

Object interchange architectures. A document is a collection of pages that can include text, graphics,


and images. A very important consideration in image work is data-stream selection. Image data are not easy to convert to alphanumeric data. Also, because of the large size of the databases that enterprises are building, data conversion may be impractical. These considerations are illustrated in Figure 6.

The IBM Systems Application Architecture™ (SAA™) document data stream is Mixed Object Document Content Architecture-Presentation (MO:DCA-P). MO:DCA-P is a carrier data stream that consists of objects and the layout information that specifies how a document is to be printed or displayed. Objects may be in line as part of the document, or they may reside in an external library.

Objects that can be carried in MO:DCA-P include text, image, graphics, and fonts. In an image system, these are the most common objects to carry. However, MO:DCA-P can carry other types of objects, as required. Image Object Content Architecture (IOCA) is the data stream for images, the purpose of which is to put the image data into a format that can be understood by another program. Things like type of image, width, length, resolution, compression type, and such must be carried in the data stream. Then an image becomes self-contained, that is, other programs can interpret the image without referencing a separate database for such data. Often image data cannot be interpreted without the associated data. IOCA can handle bilevel, gray, and color images. MO:DCA-P is used to describe where on a page an image is to be placed. Although having an image data stream is obvious, it is not as obvious why there is a text data stream in an image product. The answer comes from storage requirements. The purpose of a document system in a business is to store documents, and an image is just one type of object that a document may carry.

It is common in an enterprise for customer correspondence to be on preprinted forms. If the completed forms are scanned and stored in a database, the result would be to store multiple copies of the form. Also, the printed materials would be stored as image rather than as text, meaning that image is at least eight times larger. Instead, the idea is to store the printed data as text along with its font description and a reference to an image of the form stored in a sample-form library. Then, when the document is requested, the system presents the text information in the correct font and overlays it on the form. The operator sees the same document, with at least an eight-fold reduction in storage. A comparison of the

Figure 6 IBM Object Interchange Architecture




methodologies of (A) preprinted forms and (B) images is shown in Figure 7. The text data stream is Presentation Text Object Content Architecture (PTOCA), and the font data stream is Font Object Content Architecture (FOCA). Graphics, such as for annotating an image, could be represented with Graphics Object Content Architecture (GOCA).

Types of storage. Because of the size and volume of image data, traditional storage media are often impractical. Three classes of storage become important in image processing: short-term, near-term, and long-term storage. There are two key factors that separate one from the other—speed of access and cost.

The highest-speed storage is usually the highest-cost storage. Therefore, it must be used and reused on a short-term basis. Magnetic disk is the common form of short-term storage for storing images. Large enterprises keep the documents that are being processed in short-term storage for quick retrieval. Access to short-term, high-speed storage is measured in units that are less than a second.

A near-term storage system used for data that are needed in seconds is shown in Figure 8. Here the cost is expected to be lower. Stacks of optical disks analogous to the records in a juke box are used as near-term storage in most image systems. In an optical juke box, the optical disks are stored in slots.

Figure 7 Comparison of text-on-preprinted-forms and image methodologies



The turntable is one or more drives. When data from a disk are required, the arm fetches the selected platter from its slot and places it in an empty disk drive. When use of the disk is complete, the arm returns the disk to its slot.

Optical disks use lasers analogously to magnetic read-write heads of magnetic disk drives. In a pure laser system, the laser is used to create pits in the disk material. There are three types of optical disks: read-only-memory (ROM), write-once-read-many (WORM), and write-many.

A ROM disk cannot be written on. The disk is stamped with all of the data prewritten. The ROM

disk is used for very stable data, such as manuals. Music compact disks are an example of a ROM disk. One of the advantages is cost. Once a master disk is tooled, producing many disks is relatively cheap.

WORM disks can be written only once. The drive is able to burn a pit but cannot unburn one. WORM disks range from 500 megabytes (MB) of data to multiple gigabytes (billion) per side of the platter. They are much lower in cost than magnetic storage, but more expensive than ROM disks. Image systems use WORM disks for two reasons: They provide large amounts of storage at low cost, and, because of their unchangeable nature, they have served as evidence in court for storage of documents. One of the advan-

tages of an image system is the elimination of paper movement. Storing and moving paper is a big expense in business that an image system can remove. Much of the financial advantage comes from destroying the original paper and relying on the electronic form only. In previous systems, paper was photographed onto microfilm and the paper destroyed. If the document was required in court, the microfilm version was presented as the best available evidence. Details of legal aspects of optical storage are still being worked out.

Another interesting feature of WORM is the ability to destroy data stored there. It is not possible to change data because bits cannot be undone, but it is possible to destroy data by setting all bits on. In many environments it is necessary to destroy data rather than just flag it for deletion.

A write-many optical disk is just like a magnetic disk in use, but the technology is different. Write-many disks can be updated, although they are somewhat slower than magnetic disks at this time. It will be interesting to see how write-many disks are supported in the courts. Optical disks, because of their low cost, have made image applications practical.


Long-term storage is data that are archived outside the juke box. Long-term storage requires an operator to mount the platter either in a drive or in a shelf of a juke box.

Impact of communications. As was shown, the large volume of image data has quite a large impact on data storage. Because of the large volume of data involved, communications is also affected. When running an application that uses numbers and letters for data, an average transaction is a screen of 2000 bytes. Allowing for overhead and such, even a 2500-byte transaction load would be high. If an image transaction requests a page of image, a normal average would be between 50 to 60 KB. This is consistent with the typically 20-fold increase in data volume.

This increase is not only seen between systems, but it also occurs to and from the input and output of a system. Consider printing an image as an example. At 300 ppi an 8.5- by 11-inch page of print in bilevel is 1 MB of data. If a laser printer is rated at 10 pages per minute, this 1 MB of data would have to travel to the printer in a maximum of 6 seconds.

From the discussion so far, we can conclude that image is a new data type to computers. With its size

Figure 8 A juke box near-term storage system



and characteristics come some considerations, but complete new areas of business automation are now possible. Image can now be justified in a business environment by providing new levels of service and support. With these new opportunities come new challenges.

Document management systems

Document management primer. Understanding problems in the area of document management is the first step to understanding an image document management system.

The challenge of document management can be summed up in three Rs: repositories, retention, and reference. All businesses have repositories of documents that are retained for a period of time and then purged from the active files. These documents are then archived for a legal retention period. Before the document's end of life, it may be used or referred to by professionals. It may have been a piece of mail that was received, such as a claim against an insurance policy. The claim document may have then

been sent to a processing department, where the claim was processed by one or more individuals. Processing this claim may have required reference to historical documents in a central file or repository. When processing is completed, the claim request is filed in a folder. When the customer inquires about the status of the claim, a service person may request the file for review and then respond to the inquiry.

All of these files, folders, and such are stored in the "paper house." If one were to take a tour of a company's paper house one would get a feel for the operation and for its application of the three Rs. One might ask one's self a number of related questions:

- Where are the repositories? What is the volume of documents they contain? Who maintains them?
- How long does the company retain the documents? Does the company have a purge strategy?
- Who refers to the document and how often?

Consider the following hypothetical paper house:

- Central files of 1 000 000 folders
- File requests per week averaging 30 000
- Probability of 25 percent that the requested folder is out of the file
- File growth rate of 5 percent per year (i.e., new folders minus purged folders)
- Suspense rate of 15 percent for claims processing, i.e., further processing delayed
- Files missing for up to 45 days

The following is a discussion of the above information gathered on our hypothetical tour of the paper house. The magnitude of the document management challenge is difficult to comprehend. The processing, storage, retrieval, and use of 1 000 000 of anything is difficult to manage. We start with the space required to hold this many folders, and we are confronted with the fact that storage space is costly. The next issue is the size of the staff to maintain this repository. There are probably 30 to 50 people doing filing, retrievals, and chasing out-of-file folders. Let us assume that there are 30 000 requests per week by processing personnel for folders. That means 30 000 slips of paper come to the paper house each week, and 30 000 times per week a clerk seeks out a folder in the paper house. Out of these requests, 7500 times per week nothing is found, and a slip stating "out-of-file" is returned to the requestor.

Business executives know they have these problems, and in the following sections we review some manual

systems that are used today to solve some of them. We also look at the impact of operational image. A study, therefore, includes at least three areas of interest: the mail room, a service area, and a filing area.

Manual systems. Assume that the manual system under consideration here includes the special color coding of folders, special numbering schemes, and sometimes special equipment. As an example of special equipment, consider that the 1 000 000 folders may be stacked on movable shelves so that aisles between the shelves appear and disappear as the clerks crank the shelves back and forth to find folders. This scheme reduces the amount of floor space reserved for shelves. Other equipment that may be used are power files, in which the shelves are mounted on an assembly much like the seats of a Ferris wheel. The operator advances the power file shelf assembly with the assistance of a motor. Then the file is retrieved manually from the shelf.

The next level of advancement in filing systems might involve computer assistance to cross reference policy numbers to shelf locations. The clerk then takes the file request to the terminal and enters a policy number to get a shelf location.

Another enhancement in manual systems might be to automate the request process. For example, the requestor might place requests for files, using a data processing terminal. An application in the host might print picking lists four times per day in a central files department. The lists might be printed on address label stock together with the requestor's return address. This list may be printed in a sorted order to optimize the clerk's retrieval task. The application software that automates the requests eliminates the mail time for the request slip to arrive at the repository location. For the return trip, the clerk puts the package together and drops it in the internal mail.

Microform systems. Microform technology as used by businesses generally falls into one of three forms: roll film (also called microfilm), jackets, and fiche.

All of these systems involve the photographing of documents on a film stock, the best known of which is the microfilm system. This system involves creating rolls of film with pictures of documents on them in arrival sequence. A variation of this is the jacket system, which is an attempt to organize microfilm into folders of related documents. To do this, the rolls of microfilm are cut up and filed in jackets that

are equivalent to a paper folder. The microfiche alternative results in similar filing sequence as the jacket. However, the fiche is created from a light-sensitive stock.

The advantage of microform systems is that the film equivalent takes up considerably less space than the paper. Microfilm indexing software is available on the market to aid in finding images on roll film. Some vendors have racks that will hold many rolls of film, and the retrieval is automated by robotics. After loading the roll the view station will fast forward to the desired document. Retrieval for jackets and fiche is done manually with perhaps the assist of color schemes for the top edges of the jackets and power files to hold the repository.

These systems do little or nothing for the person who is trying to process a claim. The requestor still has to wait for internal mail to return the needed documents for review. The file clerk sees the benefit of having a significant reduction in out-of-file situations because the film never leaves the repository. The clerk makes copies of the film for the requestor, and a view station has a printer option. The clerk makes copies of the film and places the document copies in the mail. The central files department gains benefit from microform systems, and requestors receive fewer out-of-file responses. All this increases the productivity of the clerical staff, but it does little for the end users of the documents. They must handle physical papers as they would in a manual system.

The image system. An image system, as illustrated in Figure 9, is designed to change the way companies handle business records. We now follow a claim through an image system. The claim is opened in the mail room and digitized by a scanner or camera. The digital image of the claim is indexed in the system and dispatched to a processing department as electronic mail. The processing department receives images at electronic speeds at an image workstation. The claims processor may need to see historical documents from the repository, which in this example are digitized images on optical media. The documents can be retrieved at electronic speeds. There is no waiting for the mail cart to deliver the request, for the clerk to pull the file, or for the clerk to return the folder. The claims processor can complete the action item in one pass. The folder is requested and transmitted electronically. The out-of-file response is eliminated, because the documents never leave the system. Multiple access to a folder is

possible, because many actions to an account may be required at one time. With an image system, many people can reference the folder, which results in productivity gain for the end users, a reduced clerical staffing level, and a likely increase in service levels throughout the organization. Queries are handled in seconds.

Scope of image applications. Scope is an important factor in determining and managing business opportunities. A business should understand that the benefits of image systems come from establishment-level solutions that tie multiple departments together.

Department. The most commonly discussed image solution is one that solves a problem within a department. Most image solutions implemented today are departmental solutions. A company might want to apply image to the central files department or at the agency level. Each of these areas may be called departmental in scope and may have different problem definitions.


The central files department would like a way to pull and print file requests faster. Such departments receive several thousand requests per week and would like to simplify the processing. One complicating factor is a 25 percent out-of-file response to file folder requests from the central filing area. Central files requires a system that improves the maintenance of central files. The scope of the problem definition restricts the solutions to improving clerical productivity for the central files department.

An agency might view image processing as the ability to electronically deliver new business applications to the company's underwriters. The agent is defining his problem as the company's ability to deliver the application more quickly to the underwriter and thereby improve turnaround time in the processing of new business applications.

Establishment. An establishment view of image processing would include several departments for which the solution would require a comprehensive, integrated approach to applying image technology. This scope requires a company to examine the integration of image with existing data applications. The company has to look beyond the limited payback within a department and to look at the possible savings across the establishment.

An establishment-level solution will help documents flow within and between departments. It will bring

Figure 9 An image document system



the establishment closer together through electronic image delivery. This step is similar to electronic mail except that instead of keying in data, scanners are used to create an electronic image for delivery.

Enterprise. The enterprise is the largest scope for an application and includes multiple establishments in the solution architecture. The establishment could have several companies under a holding company or perhaps be in several other businesses. Enterprise solutions may be studied today for volume procurement purposes.

Integrating these separate departmental systems could be very costly without a document management strategy for the organization. One's position should be that the data processing system is the logical place to begin the integration of data and document management systems. The business is processed on the data processing systems, and this investment should be leveraged further to include document management. The installation of a departmental system will result in the duplication of information on two systems and the necessity to synchronize those systems. The problem is compounded if there are multiple departmental systems

and even further if there are multiple vendors providing these systems.

Operational image

An operational image system is shown schematically in Figure 10. Operational image is the use of image technology in applications that have the following characteristics:

- High volume of documents to be processed
- Paper intensive
- Access to documents required to complete work items
- Management of the paper flow critical to the customer's business


Many vendors are attempting to provide a solution with high enough performance and function to meet

the demanding requirements of an operational image environment.

In the past, users required efficient archival storage and wanted to avoid having folders out-of-file for long periods of time. These requirements were satisfied with cost-effective microform solutions. As image technology has developed, the ability to retrieve the stored document more speedily than with microform systems has become an additional user requirement. The end user wants electronic access to the stored document, as well as rapid retrieval of a printed copy. The need for multiple departments to share access to their stored documents is an additional requirement.

A user's perception of how image technology will work is that first paper is converted to digitized images at the point of mail delivery. The digitized

Figure 10 Operational image system



images are then electronically dispatched to professionals for processing. When all work groups have completed their tasks, the electronic folder contains all pertinent information regarding actions taken.

We now discuss the major components of an operational image system.

Components of operational image. The major software components of operational image are:

- Repository
- Folder management subsystem
- Work-flow management subsystem
- User's application

Repository. The purpose of the repository is to store objects. In a system such as the one we have been describing, these objects are documents with images in them. To avoid confusion with the term "object" as used in document architectures, we call them simply "documents."

Documents can be stored in several classes of storage. As explained earlier, the three classes of storage we deal with are short-term, near-term, and long-term storage. The repository must keep track of where the documents are stored. It is normal for documents to reside in multiple classes of storage. If a document was in long-term storage, it would not be moved to short-term storage, then near-term storage. To do so would involve another copy written on optical storage in the juke box. The document would be moved to short-term storage, then deleted when done. If a new copy is needed, it is reloaded in short-term storage.

Document and folder management subsystem. The repository is the place where documents are stored. A document-and-folder-management subsystem is a place where the document profiles and the folder structures are kept. A *folder* is a collection of related documents. As an example, a folder might be all documents related to a particular customer. Folder schemes can range from the very simple to the very complex. A document and folder management subsystem must be capable of dealing with such questions as the following:

- Can a folder contain a folder?
- Can a document reside in multiple folders without multiple copies of the document?
- Are multiple repositories covered by the folder manager?

Each document has a profile that describes the document and the kind of data that it contains, which might include owner of the document; date and time of its creation; class of storage; security; description of document; retention parameters; audit trails; name the document is known by in the repository, and so forth.

Such fields as class of storage and retention parameters are called the management policies, an example of which might be the following:

- Initially store on magnetic medium and optical juke box.
- After 30 days, delete from magnetic storage.
- After 90 days, move optical platter to shelf (archival storage).
- After 7 years, dispose of the optical platter.

A storage and retention policy for a more temporary document might be the following: Initially store on magnetic medium. After 30 days, and once the document has been handled, delete it from magnetic storage. A carefully designed set of management policies is a very important part of an effective image system.

Work-flow management subsystem. Work-flow management software performs two major functions. One is to get documents from one point to another; this is the highway system for the delivery of images. The other function is the dispatching of images to action queues. For example, in the mail room, a document might be digitized by a scanner. A work-flow management subsystem allows for assigning an action and designating a processing department or individual. The action request is queued. A more complex case might require multiple actions. These actions may be sequence dependent, or they might be queued for parallel processing. Work-flow management software is a critical enabling component to achieving the maximum benefits from an operational image implementation.

Work-flow software can be very simple or very complex. In a simple system, the user would manually direct the piece of work to the next station. In a complex system, the processing of a document would be divided into a series of steps, based on the type of work to be done and such parameters as dollar limits. A fully automated system would direct the document through the steps, with decisions on where to send it next being made by the system. This is a very complex system, and it is seldom done that way. A hybrid

system incorporating features of a manual and a fully automated system is usually used, with the exception cases being handled manually and the system directing the standard cases.

Long-term storage of documents. A critical factor in image storage is that of standards. When considering the storage of images over periods as long as years or decades, there are questions that must be answered: Will the medium survive? Will there be equipment that can read the medium? Will the data stream on the medium be understood?

The first point is one of the length of time a medium will last before errors begin to appear in the data. The length of time magnetic media can survive is well understood, because magnetic media have been around for a long time. Optical media technology, on the other hand, is not mature and has an uncertain shelf life.

While life of a medium is the most critical issue, it is the easiest of the three issues to deal with. Availability of equipment is often a greater problem. In optical media, for example, there are currently no accepted standards, which probably means that the most prudent choice is to read an optical disk in the brand of drive that created it (i.e., wrote it). As optical storage matures, this will probably change. With this change will come the abandoning of many of the media types supported today. If this happens, one will have to keep the older reader drive in working order at least long enough to copy images onto the new media. As an example in the magnetic audio recording technology, technicians are still occasionally called upon to transfer old magnetic wire recordings onto magnetic tape. Although it may seem easier to maintain an obsolete technology, it is often the more difficult solution to the problem of long-term storage of documents. There is a point at which it is cheaper to replace equipment than to maintain it.

Copying to new media may be the only acceptable alternative. New media pose the problems of cost of the new media and the cost of the time to transfer from one medium to another. Unlike film-based storage, a digital copy is an exact duplicate of the original. Thus a digital image database is the preferred means for long-term storage. However, when considering a data-stream technology, it is well to look to the future and consider whether the data stream is well accepted and whether the data stream is well documented.

If a data stream is well accepted, documentation is less of a concern, because it will be understood in the future. If a data stream becomes less popular but is well documented, a conversion transform can be written to a new standard. If, on the other hand, one does not have good documentation on a data stream, we may put future generations in the position of being unable to transform a data stream.

Concluding remarks

Many factors have held back image development in the past, but almost all can be summarized in one word—cost. The main costs were in storage and CPU cycles. The introduction of optical storage has also given us large-scale storage that, when compared to the cost of storing the original paper, is cheaper. Cost of memory and communications also delayed the development of image. Memory is dropping dramatically in price, and raw CPU cycles, particularly at the workstation, have jumped. An example is the Image Adapter/A used in many of the ImagePlus workstations. This PS/2® card has a 32 bit 8.5 MIPS RISC processor on it that only five years before would have been compared to a mainframe. Image work takes many CPU cycles. As an example, to decompress a typical page of a business document takes over five million cycles. The capabilities of the output devices have also jumped. High resolution displays and printers are now becoming common and are reasonably priced.

Acknowledgments

Several individuals have helped me over the course of writing about this diverse topic. I would like to thank Gregory Dominato of Toronto Information Development, Bill Linhart of Image Market Development, Peter Somerville of Image Strategy and Architecture, Dennis Spokany of the Image Systems Center, and George Stierhoff of the *IBM Systems Journal*.

ImagePlus, PageScanner, Systems Application Architecture, and SAA are trademarks, and PS/2 is a registered trademark, of International Business Machines Corporation.

General references

- Architectures for Object Interchange*, GG24-3296-00, IBM Corporation; available through IBM branch offices (1988), pp. 20-67.
- G. A. Baxes, *Digital Image Processing*, Prentice-Hall, Inc., Englewood Cliffs, NJ (1984), pp. 3-31.
- K. B. Benson, *Television Engineering Handbook*, McGraw-Hill Book Company, Inc., New York (1986).

M. F. Cowlshaw, "Damped Error Diffusion in Binary Display," IBM Technical Disclosure Bulletin **28**, No. 3, 1290-1291 (August 1985).

H. J. Durrett, *Color and the Computer*, Academic Press, Inc., Orlando, FL (1987).

R. W. Floyd and L. Steinberg, "Adaptive Algorithm for Spatial Grey Scale," *SID International Symposium Digest of Technical Papers* (1975), pp. 36-37.

R. W. Floyd and L. Steinberg, "An Adaptive Algorithm for Spatial Greyscale," *Proceedings SID* **17**, 2, 75-77 (1976).

R. C. Gonzalez and P. Wintz, *Digital Image Processing*, Addison-Wesley Publishing Company, Reading, MA (1977), pp. 1-12.

Image Object Content Architecture (IOCA) Reference, SC31-6805-00, IBM Corporation; available through IBM branch offices (1990).

J. F. Jarvis, C. N. Judice, and W. H. Ninke, "A Survey of Techniques for the Display of Continuous-Tone Pictures on Bilevel Displays," *Computer Graphics and Image Processing* **5**, 13-40 (1976).

H. Scholfield, *Monochrome & Full Color Thermal Imaging*, paper presented at Electronic Imaging East '88, Oct 3-6, 1988, Institute for Computer Graphic Communication (IGC), Boston, MA **1**, 268-270.

P. Stucki, *MECCA—A Multiple-Error Correction Computation Algorithm for Bilevel Image Hardcopy Reproduction*, Research Report RZ-1060, IBM Research Laboratory, Zurich, Switzerland (1981).

G. Wyszecki and W. S. Stiles, *Color Science*, John Wiley & Sons, Inc., New York, NY (1982).

Richard M. Helms IBM Canada Laboratory, 844 Don Mills Road, North York, Ontario M3C 1V7 Canada. Mr. Helms joined IBM in 1976. He is currently part of the Image Systems Management group, where he is the lead image architect. Since joining IBM, Mr. Helms has worked in Canadian marketing in both sales and support roles. He joined the IBM Canada Laboratory in 1982 where he cofounded the Toronto Image Systems Centre. Mr. Helms received his B.A. degree in mathematics from Clarion University, Clarion, Pennsylvania, in 1971.

Reprint Order No. G321-5402.