Unit 1. Introduction to Images

1.1 Images

Images and Pixels. An image is a rectangular array of dots called *pixels* (picture elements). The number of rows M and number of columns N of dots in an image are specified. At each row-column intersection (m,n) there is a pixel, or picture element. The point (m,n) is the *location* of the pixel, while the *pixel value* at that location is designated by p(m,n), or sometimes by f(m,n) or f(x,y). Figure 1.1 shows the location (m,n), where $0 \le m \le M-1$ and $0 \le n \le N-1$. Note that in the figure the downward vertical direction is x and y is the the horizontal rightward direction. The origin is in upper left corner.



The pixel values in an image may be *grayscale* or *color*. We first deal with grayscale because it is simpler and even when we process color images we often process the intensity part, which is grayscale, and then put the color back into the processed image. Grayscale usually has a range from 0 (no light intensity, or full black) to 255 (full light intensity, or full white), in integer steps. Thus the 256 grayscale values for pixels are 0, 1, 2, ..., 255. Each of these takes one byte of storage in a computer or in a file, so if an image has MxN = 256x256 = 65,536 pixels then that image takes 65,536 bytes of pixel storage in computer memory or on a storage device. An MxN = 1280x1024 image has 1,310,720 pixels. A pixel value p(x,y) or p(m,n), where $0 \le m \le M$ and $0 \le n \le N$, is a byte (0 to 255 in binary) in grayscale.

Fig. 1.1 The Image Coordinates.

Color images most often take one of two different forms. The most common method is called *true color* and uses one byte for each of red, green and blue. Thus a single pixel value requires 3 bytes of memory or disk storage: one for each of red (\mathbb{R}) , green (G) and blue(B). From these values we can form (256)x(256)x(256) = 16,777,216 discrete colors, which is about the maximum number of different colors that humans can distinguish. An MxN = 256x256 color image of three bytes per pixel would then require 3(65,536) = 196,608 bytes. For an MxN = 1280x1024 color image the requirement is 3(1,310,720) = 3,932,160. It is clear that color and more pixels are more costly in computer storage and time to send on the Internet.

The older VGA format for color is to allow only 256 colors at any one time on the screen. A byte indicates a color but it is actually the address from 0 to 255 of one of 256 color registers, each of which contains 18 bits for 6 bits each of R, G and B. The 256 color set, called a *pallette*, must be loaded into the registers before the image can be displayed, or else the pallette can be read from the image file. This is suitable for many types of images, but the trend is toward even more colors than true color, which may be a waste of resources due to the fact that such fine resolution of color is wasted on humans. Color images are covered in detail in a later unit.



Figure 1.2 shows the grayscale pixel values as a function f(m,n) of the pixel locations at rows m and columns n. Thus we can picture an image as a 3-D surface that has elevation (gray level) as range above the image plane that is the domain. The gray levels are discrete values, so the surface is made of discrete steps at the pixels.

Figure 1.2. Display of a pixel value.

Image Files. An image is stored in a particular file format. The most popular formats nowadays are GIF (Graphics Interchange Format), JPEG (Joint Photographic Experts Group), PNG (Portable Network Graphics), TIFF (Tagged Image File Format), PGM (Portable Gray Map) and PPM (Portable Pixel Map). Examples of the file names for the image *lena256* are *lena256.gif*, *lena256.jpg* (or *lena256.jpg*), *lena256.ppg*, *lena256.tif*, *lena256.pgm* and *lena256.ppm*. The PGM format is strictly grayscale and is the simplest file format for processing (there is no compression nor decomposition of the data into file sections. We examine this format later in this unit.

1.2 Displaying Images with the Matlab Tool

The Tool. We assume here that the reader has a PC (personal computer) available or else has a computer account and a user directory, which we will call *user*. If the machine has *MS Windows, Linux, Solaris* or *Unix* installed as the operating system, then it should also have *Matlab* 6 (or later) installed. If the computer has *Linux, Solaris* or *UNIX* as its operating system, then it may have the image program *XView* installed (shareware). Our tool will be *Matlab* 6 (or later) with the *Image Processing Toolbox*.

The program *XnView* is a valuable free program because it permits the conversion between dozens of different image formats. One can download it for the PC for free (the basic version). It can be found by bringing up *Netscape* (or *Internet Explorer*) and performing a search on *XnView*. Lview Pro can be downloaded for the PC and used free for a trial period. The network address <u>www.lview.com</u> can be used for this. It can also be purchased for a reasonable price with complete documentation. Our main tool will be *Matlab* for either *MS Windows*, *Linux*, *Solaris* or *UNIX* because it is the most powerful and is the only tool that permits us to process images in the frequency domain. We will include some simple fundamental programs in C for processing PGM image files. The public domain images *lena256.tif*, *building.tif*, *shuttle.tif* and *peppers512.tif* can be downloaded from the authors Web site at <u>http://ultima.cs.unr.edu/index.php</u> by clicking on the little square under the heading "Downloads" where the caption reads "Some Images for Processing!" To save each image, right click on it and then select "Save Image As" in the pop-up menu.

Displaying the First Image with Matlab. The directions for displaying *Lena* with the *Matlab* tool are given below.

1. From an MS Windows PC: click on the *Matlab* icon. If there is no such icon, then click on the *Start* button (this is usually located on the left side of the bar across the bottom (or top) of the screen, but the bar could be on the right side of the screen as a vertical bar). Next, click *Programs* (or *All Programs* for Windows XP) in the window that pops up and then click *Matlab Release XX* on the menu of programs that appears. The basic Matlab window will come up on the screen as in the example shown in Figure 1.3. Click in the *Command Window* frame on the right hand side and the cursor will appear (the symbols ">>" indicate where your typed commands will show). Along the top of the Command Window are the words *Current Directory*, followed by a text entry field and an arrowhead pointing down. Click on this arrowhead and then select the directory where your images are stored. We assume hereafter that the directory is set to where the images are that we will process. The Matlab command window is shown in Figure 1.3.

2. From a *Linux/Solaris/UNIX* **Machine**: from a terminal window (you must be in X-windows), type *Matlab* at the cursor.

> matlab

It may take several seconds to load if *Matlab* is being served up by a network server machine. But when it comes up, it will look the same as on a PC.

3. Display the Image: whether you are in Windows or Linux/Unix, type the following at the command line cursor in the Command Window on the right hand side in the window.

>> I = imread('lena256.tif');
>> imshow(I);



Figure 1.3. Lena256 displayed by Matlab.

Figure 1.3 shows how to read in the file *lena256.tif*. The first command uses the function *imread()* to read in the file and put the raw image data in memory under the variable name *I*, which could be *Im1* or *myImage* or any other variable name that starts with an alphabetical character. The second function command shows the image I on the screen by writing it from computer memory into the graphics memory. Figure 1.3 shows the results. To move the image, click and hold wn on its top frame bar and drag to the location desired.

1.3 Changing Contrast: An Example of Processing and Saving Images

We will use the three tools demonstrated above to process images. Here we use the image *lena256.tif* to show the processing of an image and the saving of it to a file. The processing changes the contrast of the image, that is, changes the range of image gray levels.

Changing Contrast with Matlab. First we run *Matlab* and then read in an image as Im1. Then we use the function *imadjust()* to provide the input interval and output interval to which the input interval is to be mapped. The intervals are standardized so that [0, 255] is [0, 1] in the function. The commands are given below where Im1 is the variable into which lena256.tif is read and Im2 is the variable that contains the adjusted image.

>> Im1 = imread('lena256.tif'); >> Im2 = imadjust(Im1, [0.2, 0.8], [0, 1]);

Here, all standardized shades of gray from 0 to 0.2 are mapped into 0 and all shades from 0.8 to 1 are mapped into 1. The interval from 0.2 to 0.8 is stretched linearly to 0.2 to 1.0. Figures 1.4 and 1.5 show the respective *before* and *after* images. Images can have the contrast increased (stretched) or decreased (compressed) in this manner. The general format is

>> image2 = imadjust(image1, [low-in high-in], [low-out high-out]);





Figure 1.5. Matlab contrasted lena256.gif.



Saving and Printing with Matlab. To save the processed image, click on the processed image to bring it into focus and then click on *File* in the upper left corner (see Figure 1.5). In the menu box that pops up, click on *Save As* to save in the same format or on *Export* to transform to another image file format. Click the down arrow at the bottom to bring up a list of options (EPS is encapsulated postscript, but there are JPEG, TIF, PNG, PGM, PPM and others). When everything is selected as desired, then click *OK*.

To print the processed image, click on *File* and then on *Print* (select the *Print Preview* instead of *Print* to see what the printed image will look like).

1.4 An Overview of Image Processing

Public Domain Images. The image Lena of Figure 1.3 is 256x256 and has been in the public domain for almost half a century (it supposedly appeared in *Playboy* magazine originally). It is often used by image processing researchers around the world for comparison of their methods and algorithms with other ones and appears in most journals on image processing. We will perform certain image processing algorithms on it as well as on other images in the public domain. Another image that is often used is *Peppers512*, which is 512x512 pixels. It is shown in Figure 1.6. Other images that we will use often are the public domain images *Shuttle* and *Building*, which are both NxM = 320x240 in size. These images are shown respectively in Figures 1.7 and 1.8. These images are quite different and are useful for demonstrating different types of image processing.

Figure 1.6. The original *Peppers*.



Fig. 1.7. The image Shuttle.



Fig. 1.8. The image Building.



In what follows, we describe two ways of considering the field of image processing. The first is by the needs, or purposes, of the processing. The second is by type and method of processing. These are listed below.

The purposes of image processing. The list below gives important processes, but not necessarily all.

1. Image Enhancement

- improve image or prepare for some other purpose
- smoothing, sharpening, contrast reduction/increase, color improvement, edge enhancement, edge detection
- adjust histograms
- enlarging/reducing image size, interpolation for values of new pixels
- combining multiple images to show details better
- preprocessing before applying some other methods
 construction of complete image from multiple partial ones

2. Image Restoration

- removal of speckle (dot) noise, short line noises or scratches
- filtering out dark and light bands, elimination of artifacts
- removal of motion blur
- unwarping of spatial distortions
- elimination of glint, glare and haze

3. Image Analysis

- obtain numerical or graphical parameters from image properties
- segment image into similar parts and extract features
- obtain regional features
- measure object parameters
- detect cracks in material or foreign materials in packaged products, etc.
- "see" to detect and track objects in scenes (computer vision)

4. Image Compression

- reduce image file size in number of bytes of images for transmission/storage
- lossless or lossy give respectively low or high ratio of original to compressed size in bytes (e.g., 2-to-1 or 20-to-1)

5. Computer Vision

- detect objects and extract features from images
- use features to make decisions to control or respond to scene activity
- enhance, analyze a scene, measure or track objects in images

The Methods/Types of Image Processing. These are the kinds of processes done on image files.

1. Images

- images, grayscale and color pixels
- image file formats, PGM, PPM, TIFF, GIF, PNG, JPEG, BMP

2. Point Processes

- threshold gray levels
- contrast stretching and contraction
- pixel transformations, histogram adjustment
- map pixel values to hide or expose certain objects
- histogram equalization, dual histogram equalization

3. Area Processes

- transform pixel values according to the values of its neighbors
- smooth, sharpen, detect edges
- filter out noise, scratches (despeckling)
- trimmed mean and convolution filtering

4. Frame Processes

- registration of two or more images
- combining images via pixelwise sums, subtraction, multiplication, division
- combining images via pixelwise boolean logic (or fuzzy logic) functions
- combining by minimum or maximum operations at each pixel
- fuse images by other mathematical methods
- multispectral methods

5. Geometrical Processes

- expand/shrink lighter areas relative to darker areas to smooth boundaries, fill holes and remove noise
- affine/linear transformations to rotate, translate and scale images
- interpolation, transformed pixel locations, downsampling and upsampling, zooming in and out
- nonlinear transformations to remove distortion, or mirror or flip images
- segmentation, clustering of pixels, labeling pixel blobs

6. Frequency Domain Analysis

- Discrete cosine transforms (DCT's)
- Fast Fourier transforms (FFT's)
- lowpass, bandpass, highpass and bandstop filters in the frequency domain
- Gaussian filters in the frequency domain
- Convolution and frequency filters in the spatial domain, spatial Gaussian filters
- Deconvolution, blind deconvolution
- registration using frequency features

7. Color Image Processing

- capturing color images, color cameras, color scanners
- human color perception
- the color models RGB, CMY, HSI, CMYK and color model transformations
- intensity levels in color images, processing color intensity
- pseudo-color for image enhancement

- color image file formats: PPM, GIF, PNG, TIFF, JPEG, BMP

8. Image Compression

- lossless and lossy encoding
- run length encoding
- LZW compression, GIF and PNG
- discrete cosine transforms (DCT's) and JPEG

9. Special Topics

- stereo vision
- image synthesis
- hyperspectral analysis

1.5 The Field of Image Processing

The Origin of Image Processing. Imaging began in the 19th Century with photography and continued with x-rays, television and electronic scanning in the 20th Century. Image processing as a field of study began in the 1950s with pictures of the earth' surface from high flying "spy" airplanes and then with pictures taken from orbiting satellites. Electronic sensors were sent into space to probe the surfaces of the planets and their moons in the 1970s and 1980s. The newer infrared and optic sensors, and additionally synthetic array and high range resolution radars create images that require intensive processing to reveal details for detection and classification of man-made objects, crops and other foliage and of minerals. These are captured from ground stations, unmanned aerial vehicles, airplanes and satellites.

Applications of Image Processing. Nowadays, image processing is used in

medical diagnostics, forensics, biological microscopy, inspection of parts and materials, crop yield estimates, foliage types and area estimates, minerals, defense intelligence, topographic maps (a type of stereo vision), ocean temperatures, meteorology, fire danger monitoring, and other areas

An important developing area that is based mainly on image processing is computer vision. It includes enhancing images, selecting objects, identification and recognition of objects, monitoring the behavior of the objects, tracking objects, and related areas. The goal here is to mimic huiman vision, perception and analysis.

1.6 PGM Grayscale Image Files

P2 and P5 Types of PGM Image Files. The simplest file format for grayscale raw image data is the PGM format. It has multiple subformats that we are interested in: i) P2 that stores pixel values as the ASCII characters for the digits, delimited by spaces between consecutive pixel values; and ii) another known as P5 that writes the bytes without delimiters as binary numbers from 0 to 255.

Each PGM file has a filename, such as *lena256.pgm*, and a header of three lines. The first line contains either the characters *P2* or *P5*. The second line starts with the character # and is a comment that provides some information about the image. The third line contains three integers that are respectively the number of columns, number of rows and the maximum gray level (usually 255).

Viewing PGM Files with an Editor. Table 1.1 shows a P2 PGM file for *lena256.pgm*. The datum 158, for example, is three ASCII characters for the respective digits 1, 5 and 8, followed by a space. Table 1.2 shows the P5 PGM version where each pixel is stored as a byte that is a binary number from 0 to 255 (text editors show 0 through 127 as ASCII characters and 128 through 255 at whatever code the editor uses for those values).

Table 1.1 The P2 PGM file lena256.pgm.

P2 # Created with LVew Pro 256 256 255 158 165 158 158 158 155 158 155 161 155 150 155 155 158 151 150<End-of-File>

Table 1.2. The P5 PGM file lena256.pgm.

We see that the P5 type of PGM is *packed* (without delimiters) so there is a single byte for each pixel. This is also called the *raw data* format. The size of this file is 65,576 bytes. The P2 type of PGM uses a byte for each numerical symbol (digit) and therefore requires three bytes for each number greater than 99 and also uses a space character after each pixel value. Thus the file is nearly four times as large. This P2 file is 245,724 bytes. However, humans can read the P2 type of PGM file, whereas they can not read the characters of the packed bytes, which can appear different in different editors (characters 128 to 255). Appendix 1.A gives the ASCII code.

To access the data in either type of PGM file, we read the first three lines, set the the number of columns and rows, and then read the rows into a data structure. Then we can process this raw data and write the new pixel values out to a new PGM file. We do this this when we write programs to process the data. Matlab converts images to raw standardized (0 to 1) data for processing.

1.7 Converting to and from PGM File Formats

The quickest and easiest way to convert an image such as *lena256.tif*, *lena256.gif*, *lena256.png*, or *lena256.jpg* to *lena256.pgm* is to use Matlab. It can be done with either the *MS Windows* or *Linux* (including *Solaris* and *UNIX*) operating system.

Converting with Matlab. Run *Matlab* first and then load the image into memory and show it with the command lines given below. Then click on the image to bring it into focus. Next, select *File* in the upper left corner of the image frame, then click on *Export* in the menu that pops up and select the directory where the file is to be saved. Find the textbox with the label *Save as type:* to its left and click on the option arrow at the right of the textbox to see the options. Select the type of file to convert to from these options (TIF, PNG, JPEG, GIF, BMP, PGM, etc.). Finally, click on *Save*. The loaded file will be saved in the selected directory in the file format choosen.

To load and display the image *building.tif*, for example use the following commands.

>> Image1 = imread('building.tif');
>> imshow(Image1);

To convert it to PGM, click on *File, Export*, then select Save as Type, choose PGM and save.

1.8 Exercises

1.1. Suppose a color image file has 1024 rows and 1024 columns, uses *true color* of 3 bytes per pixel (one byte for each of R, G and B). The display is to use 76 frames per second. How many bits per second (bps) is required for the video signal from the graphics card to the monitor?

1.2. Write an algorithm in pseudo-code (high level English statements that indicate what is to be done and what computations are to be made) that reads the *lena256.pgm* data and computes the average pixel gray level over all image pixels. Now do the same for an algorithm that computes the mean-square error (variance) of the gray level over the image.

1.3. Convert *lena256.tif* to the file *lena256.pgm* by use *Matlab*.

1.4. Use a text editor and look at the data in the *lena256.tif* file. What can be surmised from this display? Can you detect the image data and the formatting/compression data?

1.5. Use *Matlab* to map the gray levels 0 to 50 to black, shades 200 to 255 into white and 51 to 199 into the interval [1, 254] on the image *lena256.tif*. Show the before and after images.

1.6. Use *Matlab* to convert lena256.tif to lena256.pgm in the binary format (P5). Now open the output file in an editor program and examine the data. Convert this to ASCII format (P2) and compare the data with a text editor. What can one surmise? Which is best for human consumption and which is best for machine consumption? Also locate the file type, the comments, the number of columns and rows and the maximum gray level.

1.7. Write a message on a piece of paper with black ink and scan this as an image. Use a tool to lower the contrast until the message can not be read. This file can be displayed but no one can read it. Now use a tool to stretch the contrast to include black and white and all shades of gray in between. What is the result? Use a threshold to convert to black and white. How does this compare with the original?

Decimal	Hexadecimal	Symbol
0	0	NULL
1	1	SOH (Start of Heading)
2	2	STX (Start of Text)
3	3	ETX (End of Text)
4	4	EOT (End of Transmit)
5	5	ENQ (Enquiry)
6	6	ACK (Acknowledge)
7	7	BEL (Bell)
8	8	BS (Backspace)
9	9	HT (Horizontal Tab)
10	А	LF (Linefeed)
11	В	VT (Vertical Tab)
12	С	FF (Formfeed)
13	D	CR (Carriage Return)
14	Е	SO (Shift Out)
15	F	SI (Shift In)
16	10	DLÈ (Data Line Escape)
17	11	DC1 (Device Control 1)
18	12	DC2 (Device Control 2)
19	13	DC3 (Device Control 3)
20	14	DC4 (Device Control 4)
21	15	NAK (Negative Acknowledge)

Appendix 1.A - ASCII Code

22	16	SYN (Synchronous Idle)
23	17	ETB (End of Transmit Block)
$\frac{1}{24}$	18	CAN (Cancel)
25	19	EM (End of Medium)
26	1 Δ	SUB (Substitute)
20	1R	FSC (Escape)
27	1D 1C	ESC (Escape)
20		rs (rile separator)
29		GS (Group Separator)
30		RS (Record Separator)
31	IF	US (Unit Separator)
32	20	(Space)
33	21	!
34	22	"
35	23	#
36	24	\$
37	25	%
38	26	&
39	27	1
40	28	(
40	20	
42	2)) *
42 12	2A 2D	_
43 44	2D 2C	Т
44	20	, (D 1)
45	2D	- (Dash)
46	2E	. (Period)
47	2F	/
48	30	0
49	31	1
50	32	2
51	33	3
52	34	4
53	35	5
54	36	6
55	37	7
56	38	8
50	20	0
J / 50	39 2 A	У
38 50	5A 2D	
59	3B	;
60	30	<
61	3D	=
62	3E	>
63	3F	?
64	40	(a)
65	41	Ā
66	42	В
67	43	С
68	44	D
69	45	Ē
70	46	F
71	47	G
72	+ / 18	U Ц
12 72	40 40	П I
13	49	L T
/4	4A	J
/5	4B	K
76	4 <u>C</u>	L
77	4D	Μ

78 79 80 81 82 83 84 85 86 87 88 89 90 91 92 93 94 95	4E 4F 50 51 52 53 54 55 56 57 58 59 5A 58 59 5A 5B 5C 5D 5E 5F	N O P Q R S T U V W X Y Z [$\$ (Care (Unde
96 97	60 61	a
98	62	b
99	63	с
100	64	d
101	65	e
102	66	f
103	67	g
104	68	h
105	69	1
106	6A	j
107	6B	K 1
108	6C	l
109	6D	m
110	0E 6E	n
111	0F 70	0
112	70 71	p
113	71 72	y r
115	72	I S
115	73 74	t s
117	75	11
118	75 76	v
119	77	Ŵ
120	78	x
121	79	V
122	7A	Z
123	7B	{
124	7C	Ì
125	7D	}
126	7E	~
127	7F	DEL (Delete)

(Caret) (Underline)

Appendix 1.B - The Display and Capture of Images

Cathode Ray Tubes for Display. The inside surface of the screen of a monochrome monitor is coated with a phosphoric material that converts the energy of colliding electrons into light emission on the outer side. This material is uniformly coated so that the light emitted is of a single color such as white, amber, or green. Any spot where there has been an absence of colliding electrons for a short time appears dark to the viewer. The screen is the flattened end of a large vacuum tube that contains an electron gun at the other end. The gun's cathode is heated electrically so that it emits a stream, or beam, of electrons toward the inside of the screen. Such a device is called a *cathode ray tube* (CRT).

Figure 1.B.1 shows the scheme. Electrons pass through two different pairs of parallel metallic plates, of which one pair is horizontal and the other is vertical. A voltage difference across each pair pulls the electrons up or down, left or right. The *control signal* voltages on the plate pairs are designed to force the electron beam to move across the screen in rows that are imperceptibly slanted downward and then return across to the left at a position one row down. This continues until all rows on the screen have been traced, at which time the beam tracing system is reinitialized to repeat the process. Such reinitialization causes a very short time delay. This fashion of painting the screen is call *raster scanning*.

Figure 1.B.1. The Picture Tube Scheme.



The denser the electron beam on any dot of phosphor, the brighter is the light emitted from that dot on the screen. If the beam density is sufficiently low, the dot appears black, while if it is at maximum level, the dot emits the maximum intensity of light and is white. The *intensity signal* f(t) is determined by the image data and f(t) controls the electron beam density at each instant that corresponds to a position in the raster scan. The timing of the raster scan control signal is such that a small dot, or area, say, of 0.26 mm diameter, is excited by electrons during a very small time interval Δt . Such a dot is called a *pixel* (for "picture element"). When the entire screen has been *painted*, that is, all pixels have been excited in order during the raster scan, we say one *frame* has been executed. Many frames are painted each second.

A graphics interface card connects into the computer bus (data bus, control bus, and power bus, where *bus* denotes a set of lines that carry signals). It converts binary values that represent the intensity level of the pixels into a voltage signal that controls the intensity of the electron gun at the specific times that particular pixels are being painted via the raster scan. In this manner, a temporary image is painted as a frame. The *persistence* is a property of the phosphor in that it keeps emitting light for a short time after excitation stops. The first pixel in a frame is still emitting light when the last pixel in the frame is being excited. Such light emission must decrease and be significantly less perceptible in a fraction of a second so that it does not garble symbols in the next frame.

A digital screen is a matrix of tiny units, each of which is excited by voltages across the respective column and row lines at the matric point (m,n) of intensity f(m,n). This is a more natural way to display an image on a screen.

Image Data and Data Rates. Rates greater than 44 frames per second are necessary to avoid the human perception of flickering. To achieve this rate, previous display systems traced alternate rows to the screen and then on the next scan wrote the rows in between. Thus 44 frames of half the number of rows were traced each second to avoid the flickering, although the actual rate was 22 full frames per second (a trick borrowed from the cinema industry). Such a method of scanning is called *interlaced*. Nowadays graphics systems are mostly

noninterlaced and can display more than 44 frames per second, 75 Hz, 90 Hz, 120 Hz or higher (*Herz*, or cycles per second, which here means frames per second).

A 1280x1024 screen has 1,310,720 pixels. Let each grayscale pixel have an intensity value from 0 to 255 (one byte, or 8 bits). Then a file of 1,310,720 bytes is needed to store an image. A stream of bytes is read from a file and written to the graphics memory on the graphics interface card and the values are used on the next scan. The image may not appear on the screen instantaneously. Some UNIX systems wait until all data are in graphics memory and then put it all on the screen in a single frame scan, so the image appears on the screen instantaneously. At 60 frames per second of 1,310,720 values of 8 bits each, this requires a graphics system that has a rate of 629,145,600 bits per second (629.145 *Megabits per second* or *Mbps*).

Display of Color Images. The display of a color image requires three different phosphors that respectively emit three independent colors of light when excited. It is known that any color can be obtained by adding the correct proportions of red (R), green (G) and blue (B). Color video display systems use this property. A color video monitor contains closely packed dots of three types of phosphor on the inside of the screen. Each pixel has an area composed of three dots: one each of R, G and B light emitting phosphor. Three electron guns are used of which one excites each of the red, green and blue phosphor dots. The R gun, for instance, is set to scan along the R phosphor dots, while the others each trace along their particular set of dots. The pixel color is determined by the combination of intensities of each of R, G and B. Equal intensities of R, G and B make a gray shade, from black to white.

A color image file consists of binary values for the pixels as does a monochrome image file. The difference is in the interpretation of these values. In one method, a binary value for a color image is broken into three parts to produce a binary value for each of the R, G and B electron guns. For example, in *true color* a pixel value is 3 bytes long. The first byte represents R, the second is for G and the third is for B. These files are quite large. As an example, an 800x600 image would take 480,000x3 = 1.44 MB (Megabytes). However, each pixel can take one of $2^{24} = 16,777,216$ colors. Each of R, G and B can take $2^8 = 256$ intensities. For example 11111111 00000000 000000 is a pure red of the highest intensity, while 00001111 00001111 00001111 has equal intensities of 15 for each of R, G and B and so appears middle gray (equal values for R, G and B is always a shade of gray). A value of 24 bits of all 1's gives the highest intensity of gray, which is white.

Another common method is to use a single byte for each pixel value so that there are 256 values. In this scheme each pixel value is actually the address of one of 256 registers of 18 bits (6 for each of R, G and B). Functions can be used to put a set of 256 RBG colors in the 256 registers, so that 256 different color images can be displayed in a single image. While a single image can use only one set of 256 colors, the color set (pallette) can be changed for another image (but not part of the way through a raster scan to put a single image on the screen with colors from another pallette). Because each 256 colors set can be selected from $2^{18} = 262,144$ different colors, the number of such pallettes is (262,144)!/[256!(262,144 - 256)!], which is the number of ways 256 things can be selected from 262,144 unique things (a very large number). Thus one image can show 256 shades of various yellows for flowers and some greens for the stems and leaves, while another image may show various shades of blues and purples. When the R, G and B values are all equal in the 18-bit color registers, then there are 256 shades of gray.

Figure 1.B.2 shows the registers. If one image is displayed in 256 colors and another is to be displayed after that in a different set of 256 colors, then the color registers must be given new color values (rewritten by a computer program function). Each color register contains the three R, G and B parts of 6 pixels each that control the three respective color gun intensities for exciting the respective set of R, G and B phosphor dots in the pixels. While the color values of 6 bits each are digital values, they are transformed by a *digital-to-analog* (D/A) converter into analog (continuous) signals to control the intensity of the color guns.



Figure 1.B.2. VGA 256 Color Scheme.

Different file formats exist for image color information. Before we process an image we will convert it into a PPM color file format (3 bytes per pixel in the raw packed data file) or a PGM grayscale format. After processing, it can be translated into another file formats if compression is needed to reduce the size.

Capturing and Storing Images. An image is often captured by passing a photograph through a *scanner*, which may have color capability. It may also be captured directly from a video camera that digitizes the light signal into groups of bits. In the first case, a set of small optical sensors, or photovoltaic diodes, one of which is shown in Figure 1.B.3, is arranged to form a row as shown in Figure 1.B.4. Light rays from the tubular light source reflect from a horizontal line along the picture into the small optical sensors. Each single device receives the light reflected from a small dot area on the picture. Its average dot intensity is converted into current by the device and then captured into memory.

The timing signals of the scanner cause the currents associated with the dot intensities along a horizontal row to be *captured* into digital values by shifting and latching (writing) to registers, which are sequences of bit-memory devices (*flip-flops* that can be set to 0 or 1). The picture then moves a slight amount in a direction perpendicular to the row of sensors and then another row of dots is captured similarly, and so forth until the entire page has been captured as digital data.

The size of the detectors and their closeness together determine the sampling interval along the horizontal (row) spatial dimension. The resolution is measured in *dots per inch* (dpi). Many inexpensive scanners capture 600 dpi horizontally, but because they can move the picture a slight amount, they can obtain 1200 dpi (or more) vertically. The horizontal dots are often interpolated to a resolution of 1200 dpi (or more) by inserting a dot between each pair in a row. More expensive scanners interpolate to thousands of dpi. However, high resolution color images are slow in printing, displaying and transmitting on the Internet and take up large amounts of storage space.

A video camera signal is fed into an interface card (a printed circuit card with logic chips installed) that is connected to the computer busses (data, control, address and power busses), which then captures the sampled values as a stream of binary values by storing them in memory on the card. As the memory on the interface card fills with data, an associated computer program transfers the data into a buffer in computer memory for writing to a file on disk. The data goes from the hard disk to the graphics card for display.







Figure 1.B.5 gives a higher level overview of image capture and display. Figure 1.B.6 shows a head writing electromagnetic pulses on a moving medium that is coated with ferrous material (disk or tape).



Figure 1.B.5. Image capture/display.

Figure 1.B.6. Reading/writing pulses.