

Chapter 11: Light, Sound, Magic: Representing Multimedia Digitally

Fluency with Information Technology Third Edition

by
Lawrence Snyder



Copyright © 2008 Pearson Education, Inc. Publishing as Pearson Addison-Wesley

Digitizing Color

- RGB Colors: Binary Representation
 - Giving the intensities for the three constituent colors—red, green, blue—specifies color on monitor
 - Color intensity is represented as a quantity (0 through 255)
- Binary Numbers Compared with Decimal Numbers
 - Number of digits is the base of numbering system
 - Binary is two digits: 0 and 1
 - Decimal is 10 digits: 0 through 9
 - Hexadecimal is 16 digits: 0 through 9, A through F

Copyright © 2008 Pearson Education, Inc. Publishing as Pearson Addison-Wesley

11-2

Place Value in a Decimal Number

- To find the quantity expressed by decimal number, the digit in a place is multiplied by the place value, and the results are added

Table 11.1 The decimal number 1,010 representing one thousand ten = $1,000 + 10$

10^3	10^2	10^1	10^0	Decimal Place Values
1	0	1	0	Digits of Decimal Number
1×10^3	0×10^2	1×10^1	0×10^0	Multiply place digit by place value
1,000	0	10	0	and add to get a decimal 1,010

Copyright © 2008 Pearson Education, Inc. Publishing as Pearson Addison-Wesley

11-3

Place Value in a Binary Number

- Works the same way except that the place values are successive powers of 2

Power	Decimal	Binary
0	$1 = 10^0$	$1 = 2^0$
1	$10 = 10^1$	$2 = 2^1$
2	$100 = 10^2$	$4 = 2^2$
3	$1,000 = 10^3$	$8 = 2^3$
4	$10,000 = 10^4$	$16 = 2^4$
...

Copyright © 2008 Pearson Education, Inc. Publishing as Pearson Addison-Wesley

11-4

Place Value in a Binary Number

- Given binary representation, we can find decimal equivalent value by multiplying the digit times the place value and adding the results

Table 11.2 The binary number 1010, representing the decimal number ten = 8 + 2

2 ³	2 ²	2 ¹	2 ⁰	Binary Place Values
1	0	1	0	Bits of Binary Number
1×2 ³	0×2 ²	1×2 ¹	0×2 ⁰	Multiply place bit by place value
8	0	2	0	and add to get a decimal 10

Converting a Binary Number to Decimal

- Add the decimal values for the places in the binary number with 1's

Table 11.3 Binary representation of the decimal number one thousand ten = 11 1111 0010

2 ⁹	2 ⁸	2 ⁷	2 ⁶	2 ⁵	2 ⁴	2 ³	2 ²	2 ¹	2 ⁰	Binary Place Values
1	1	1	1	1	1	0	0	1	0	Bits of Binary Number
1×2 ⁹	1×2 ⁸	1×2 ⁷	1×2 ⁶	1×2 ⁵	1×2 ⁴	0×2 ³	0×2 ²	1×2 ¹	0×2 ⁰	Multiply place bit by place value
512	256	128	64	32	16	0	0	2	0	and add to get decimal 1,010

Black and White Colors

- A byte is allocated to each RGB intensity
 - The smallest intensity is 0000 0000
 - The largest is 1111 1111 in binary
 - This is 255 in decimal
 - This is FF in hex
- Black (#000000) is no color; white (#FFFFFF) has full intensity for each RGB color

Changing a Decimal Number to a Binary Number

- If the number being converted is smaller than the place value below it, copy the number into the next cell to its right; enter 0 as the binary digit.
- If the number being converted is equal to or larger than the place value below it, subtract the place value from the number and copy the result into the first cell of the next column; enter a 1 as the binary digit.

Number being converted	200									
Place value	512	256	128	64	32	16	8	4	2	1
Subtract										
Binary Number										

Lighten Up: Changing Color by Addition

- What color does this represent:
1100 1000 1100 1000 1100 1000
- Each byte contains the decimal value 200. The color is RGB(200,200,200).
 - In HTML, write in hexadecimal #C8C8C8
 - Equal amounts of red, green, and blue, closer to white than black (medium gray)
 - All colors with equal RGB values are black, white, or gray

To Increase Intensity: Add in Binary


- To make a lighter color of gray, change the common values to be closer to white (larger numbers)
 - For example, add 0001 0000 (decimal 16) to each color:
1101 1000 1101 1000 1101 1000 RGB(216,216,216)

1100 1000	binary representing decimal number	200
+ 1 0000	binary representing decimal number	16
1101 1000	binary representing decimal number	216

Figure 11.1. Adding 16 to an RGB value.

Lighter Still: Adding with Carry Digits

- Binary addition is similar to decimal addition
 - Work from right to left, adding corresponding digits in each place position
 - Sometimes we can add the two numbers and the result is expressed as a single digit (1+0=1)
 - Sometimes the sum cannot be expressed in one digit and we must carry to the next highest place (1+1=10, put down 0 and carry 1)

		Illustrates the "no carry-in" cases	
1	0000	←	Carry, shown explicitly
	1100	←	A
	+ 1010	←	B
	1 0110	←	Sum

↓ ↓ ↓ ↓ Illustrates the "carry-in" cases

$$\begin{array}{r}
 1 \ 1110 \leftarrow \text{Carry, shown explicitly} \\
 0 \ 1011 \leftarrow A \\
 + 0 \ 0111 \leftarrow B \\
 \hline
 1 \ 0010 \leftarrow \text{Sum}
 \end{array}$$

$$\begin{array}{r}
 11 \quad 11 \leftarrow \text{Carry} \\
 110 \ 1001 \leftarrow A \\
 + 110 \ 0011 \leftarrow B \\
 \hline
 1100 \ 1100 \leftarrow \text{Sum}
 \end{array}$$

Table 11.4 Summary of the rules for binary addition. The carry-in is added to the two operands, A and B, to give the place digit and the carry-out.

Carry-in	0	0	0	0	1	1	1	1
A	0	1	0	1	0	1	0	1
B	0	0	1	1	0	0	1	1
Place digit	0	1	1	0	1	0	0	1
Carry-out	0	0	0	1	0	1	1	1

Overflow

- Because computers use fixed-size bit sequences, what happens when there are not enough bits to represent the correct result of a binary addition?
- Called *overflow exceptions*
 - Computers report them when the computation they're told to perform overflows; programmer has to find way to recover

Computing On Representations

- Changing the Colors of a Moon Photo
 - Imagine a scanned black and white photo of the moon
 - In the computer, the pixels of the photo form a long sequence of RGB triples. What values do they have?
 - They are all black, white, or gray
 - Suppose you want a colorized version?

Removing the Smudge

- You have very dark gray values (28,28,28) from a smudge on the scanner's glass
- To remove the smudge and transform the pixels into the colors you remember seeing
 - Anything very close to black (first 2 digits are 00) can be changed to black (00xx xxxx) (0-63)
 - Algorithm for this is: Any three RGB bytes, each of whose first 2 bits are 00, are replaced with all zeros

Making the Moon Orange

- To shift color of moon to orange, change the white pixels:
 - Pick a shade of orange, say (255,213,132)
 - Change all the white pixels (255,255,255) to this shade
 - This will not change the gray of the craters
- To change Light Gray into Orange Tint:
 - Red byte: Leave unchanged
 - Green byte: reduce green slightly (subtract 42)
 - Blue byte: reduce blue significantly (subtract 123)

Boosting the Red

- You decide the gray parts of the moon need to be more luminous
- Boost the red
- Shifting the red in the orange pixels to 255 is too red
 - Split the difference. Add half the difference between the current value and pure red

Digitizing Sound

- An object creates sound by vibrating in a medium such as air
 - Vibrations push the air
 - Pressure waves emanate from the object and vibrate our eardrums
 - The *force*, or intensity of the push determines the volume
 - The *frequency* (number of waves per second) is the pitch




Figure 11.2. Moon photos. (a) The original black-and-white picture, (b) tinted version of original, (c) with boosted highlights.




Figure 11.3. Sound wave. The horizontal axis is time; the vertical axis is sound pressure.

Analog to Digital

- To convert continuous information into discrete information, convert it to bits
- From zero line on graph, record with binary number the amount by which the wave is above or below it (positive or negative sound pressure)
- At what points do we measure? We can't record every position of the wave

Sampling

- Take measurements at regular intervals
- Number of samples in a second is the *sampling rate*
 - The faster the rate, the more accurate the recording




Figure 11.4. Two sampling rates; the rate on the right is twice as fast as that on the left.

How Fast a Sampling Rate?

- Sampling rate should be related to the wave's frequency
 - Too slow a rate could allow waves to fit between the samples; we'd miss segments of sound
 - Guideline is *Nyquist Rule*: Sampling rate must be at least twice as fast as the fastest frequency
 - Human perception can hear sound up to 20,000 Hz, so 40,000 Hz sampling rate is enough.
 - Standard for digital audio is 44,100 Hz (44.1 KHz)

ADC, DAC

- Digitizing Process:
 - Sound is picked up by a microphone (called a *transducer*)
 - The signal is fed into an *analog-to-digital converter* (ADC), which samples it at regular intervals and outputs binary numbers to memory
 - To play the sound, the process is reversed
 - Numbers are read from memory into *digital-to-analog converter* (DAC), which creates an electrical wave by filling in between the digital values
 - Electrical signal is output to speaker, which converts it to a sound wave




Figure 11.5. Schematic for analog-to-digital and digital-to-analog conversion.

How Many Bits per Sample?

- How accurate must the samples be?
 - Bits must represent both positive and negative values
 - The more bits, the more accurate the measurement
 - The digital representation of audio CDs uses 16 bits (records 65,536 levels, half above and half below the zero line)




Figure 11.6. (a) Three-bit precision for samples requires that the indicated reading is approximated as +10. (b) Adding another bit makes the sample twice as accurate.

Advantages of Digital Sound

- We can compute the representation
- MP3 Compression
 - One computation is to compress the digital audio (reduce number of bits needed)
 - Remove waves that are outside range of human hearing
 - MP3 usually gets a compression rate of 10:1
 - Lower bandwidth requirements, popular for Internet transmission
- Reproducing the Sound Recording
 - Bit file can be copied without losing any information
 - Original and copy are exactly the same

Digitizing Images and Video

- It would take 51 minutes to display an 8 x 10 color image scanned at 300 pixels per inch (21.6 MB) with a 56kb/s modem
- How can we see screen-size pictures in second while surfing the web?
- Typical computer screen has under 100 pixels per inch
 - Storing picture digitized at 100 ppi saves a factor of 9 in memory (reducing resolution)
 - This would still take 5 1/2 minutes to send at 56kb/s
 - Solution: JPEG Compression scheme

Compression

- Changing the representation to use fewer bits to store or transmit information
 - Example: fax is a long sequence of 0's and 1's encoding where page is white or black. Run length encoding is used to specify length of first sequence of 0's, following sequence of 1's, etc.
 - Lossless compression—original representation can be perfectly reproduced

JPEG

- Used for still images
- Our eyes are not very sensitive to small changes in hue (gradation of color), but are sensitive to small changes in brightness
 - Store a less accurate description of hue (fewer pixels)
 - Gets a 20:1 compression ratio without eyes being able to perceive the difference



Figure 11.7. Detail from an image compressed using JPEG.
(a) 14:1 compression (b) 140:1. Check images at www.aw.com/snyder

MPEG Compression Scheme

- Same idea as JPEG, applied to motion pictures
- JPEG-like compression is applied to each frame
- Then "interframe coherency" is used
 - MPEG only has to record and transmit the differences between one frame and the next
 - Results in huge amounts of compression

Optical Character Recognition (OCR)

- Reading license plate to deduct toll from car's account
- What are the difficulties?
 - Computer must capture image of license plate but camera will see other highway images
 - *Frame grabber* recognizes when to snap image and send to computer for processing
 - Computer must figure out where in the image the plate is
 - Scans groups of pixels looking for edges where color changes
 - Looks for *features*
 - *Classifier* matches features to letters of alphabet

OCR Technology

- Enables computer to "read" printed characters
 - Business applications: Sorting mail and banking

Virtual Reality: Fooling the Senses

- Creating an entire digital world
- Applies to all senses and tries to eliminate the cues that keep us grounded in reality
- *Haptic devices*
 - Input/output technology for sense of touch and feel
 - Haptic glove enables computer to detect where our fingers are. When we bring our fingers close enough together, gloves stop their movement so we feel like we're holding something

The Challenge of Latency

- The challenge is for the system to operate fast and precisely enough to appear natural
- *Latency* is the time it takes for information to be delivered
- Too long latency period ruins the illusion
 - Absolute limit to how fast information can be transmitted—speed of light

The Challenge of Bandwidth

- How much information is transmitted per unit time
- Higher bandwidth usually means lower latency

Bits Are It

- *Bias-Free Universal Medium Principle:*
 - Bits can represent all discrete information, but have no inherent meaning
- Bits: The Universal Medium
 - Everything that can be represented in a sensible way, can be manipulated
- Bits: Bias-Free
 - The meaning of bits comes entirely from the interpretation placed on them through programs
- Bits are Not Necessarily Binary Numbers
 - Bits can be interpreted as binary numbers, or not, depending on use


0000 0000 1111 0001 0000 1000 0010 0000 = 15,796,256 interpreted as a binary number
=  interpreted as an RGB(241,8,32) color (last 3 bytes)
= ADD 1,7,17 interpreted as a MIPS machine instruction
= N U S ñ ð interpreted as 8-bit ASCII—null, backspace, n-tilde, blank
= L: +241, R: +280 interpreted as sound samples
= 0.241.8.32 interpreted as an IP address
= 00 F1 08 20 interpreted as a hexadecimal number

Figure 11.8. Illustration of the principle that “bits are bits.” The same 4 bytes shown can be interpreted differently.