Lecture 14: Multimedia Systems I (Mar 24, 2005)
Yap

March 24, 2005

1 ADMIN

- Looking Ahead (Rest of Semester)

  Mar 24, 29, 31: Chapter 7, Multimedia

  Apr 5, 7, 12: Chapter 8, Multicomputer

  Apr 14, 29: Chapter 9, Security

  Apr 31: Chapter 12, OS Design Issues

2 Review

- Q: Describe how disk thrashing can happen in the maintenance of the free list?

  A: The free blocks are of two kinds: those used as a "free list node", and those that are "unused".

  Each "free list node" points to a fixed number of free "unused" blocks, and the last pointer points to another free list node.

  We only need to keep the FIRST free list node in main memory. When a block are needed, an unused node is removed from this node. When a block is freed, it is stored in this node.

  We change this node when (1) it becomes full or (2) it becomes empty.

  Thrashing happens when the node is alternately full/empty.

  EXERCISE: write out the algorithm for this non-thrashing solution!

3 Introduction to Multimedia

- Examples: digital movies, video and music files

  – literally, multimedia means more than one media

  – the "traditional" media are text (traditional), images (possibly colored)

– the "new media" might be called **continuous media** and includes sounds (audio) and continuously changing images (video).

– we use "movies" (or "video clips") to refer to the combination of audio and video.

- These have VERY different characteristics than traditional text files.

  For instance, the NORMAL way to use these files is NOT by using a text editor, but by using a media player.

- Operating system issues:

  –Storage of these large files

  –Scheduling of media player have realtime constraints

  –Both issues (large storage requirement, realtime constraints) are characteristic of continuous media.

- Examples:

  – Size: An uncompressed 2-hour HDTV movie need a 570 GB file.

  – Realtime: a NTSC video (standard in The Americas and Japan) must deliver 30 frames per second. A PAL and SECAM video (standard in Europe) must deliver 25 frames per second.

  – Quality of Service parameters: average bandwidth, peak bandwidth, minimum and maximum delay, bit loss probability.

- Other kinds of media?

  – An interesting example is **interactive games**. Is it another media? They include video clips, text as well as interaction.

  – Video on demand. This is a delivery issue but has significant OS issues. In fact, a "video server" is just a specialized OS to handle this media.

- Each "movie file" has several subfiles

  – one video file

  – multiple audio files (for the sound effects, voices in multiple languages, multiple text (caption) files)

  – E.g., DVD can store up to 32 language and subtitle files.

4 AUDIO Basics

- AUDIO ENCODING

  – Audio frequency 20 Hz to 20,000 Hz

  – **Amplitude** is measured in terms of pressure. A pressure of 0.0003 dyne/cm^2 is the lower limit of audibility

– **Loudness** is a measure of relative pressure, in units of decibel (dB). The measre is relative to 0.0003 dyne/cm^2, and is taken on a logarithmic (base 10) scale.

– E.g., a sound with amplitude of 0.3 dyne/cm^2 would be 1000 times the lower limit. This would be $\log_{10}(1000) = 3$ decibels.

– 120 dB is the threshhold of hearing pain.

– To encode sound waves, we **sample** at some frequency (Ideal: the Nyquist rate is twice the highest frequency component of the sound)

– **Quantization** refers to the discrete levels of sampled amplitude. (e.g., 8-bit samples give 256 possible values).

– E.g. Telephone systems uses PULSE CODE MODULATION which samples 8000 times per second using 7- (North America, Japan) or 8-bits (Europe).

This gives a data rate of 56,000 or 64,000 bps. Frequencies above 4 kHz are lost.

– E.g., Audio CDs has sampling rate of 44,100 Hz at 16-bits. This captures all of the audible frequencies. The bandwidth is now 705.6 Kbps (monaural) or 1.411 Mbps (stero).

Note that a Fast Ethernet line has bandwidth of 100 Mbps.

– Audio Compression (exploiting psychoacoustic properties) can achieve a 10 times compression.

– E.g., a 10 minute stereo file would require 1.4 x 600 = 840 Megabits or just over 100 MB. Using audio compression, we only need 10 MB of space.

5 VIDEO Basics

- Video encoding

  – A video is a sequence of **frames**

  – A rate of 50 frames per second will create an illusion of continuous motion.

  – To display a black-and-white frame, the electron beam scans across the screen (horizontal trace) and slowly move down to the bottom of screen, and then returns to the top of screen (vertical trace) for the next frame.

  – At a sampling rate of 25 frames per second, the video appears to flicker for some. This is especially true of older people because the image on their retina decays faster.

  – To effectively achieve 50 frames/sec, the **interlacing** technique is used: scan only the odd horizontal traces in first “half-frame”, and the even horizontal traces in the second “half-frame”.

– Many monitors scan at a rate of 75 times per second. In this case, interlacing is not needed, and we say the scanning is **progressive**.

– To display a color frame, we need three components: But these 3 components are not uniquely determined. In fact, the perception of color is a highly complex subject (below is a very much simplified account).

– Let us look at bandwidth requirements. VGA uses 640×480 pixels per frame. SVGA uses 800×600 pixels per frame. XGA uses 1024×768 pixels per frame.

– E.g., At 24 bits per pixel, 25 frames/sec, a XGA device requires bandwidht of 472 Mbps.

- COLOR

  – Visible Light occupies a small range of the electromagnetic spectrum, from 4.3×10^{14} Hz (red) to 7.5×10^{14} Hz (violet).

  – Color perception is how our eye responds to a mixture of these visible light frequencies.

  – Color perception can be measured along 3 dimensions:

  Hue or Color (E.g., read, yellow, purple)

  Saturation or Purity (E.g., red is highly saturated, pink is unsaturated)

  Intensity or Luminance (E.g., how bright is the color)

  – One objective measure of the above is:

  Hue is wavelength.

  Saturation is the ratio of light in the dominant wavelength relative to the amount of white light.

  Intensity is the energy per unit time per unit solid angle.

  – Next, we discuss several ways to parametrize color space:

  – CIE Chromaticity Diagram (XYZ) space

  The international standard (1931) where X, Y, Z are imaginary color components.

  – Screen Raster Graphics often use the RGB system.

  R=red, G=green, B=blue

  – TV (North America, NTSC) Screens uses the YIQ

  Y=luminance (=CIE's Y), I and Q encodes Chromacity.

  (REMARK: B&W Television just uses the Y component!)

  – OTHER COORDINATE SYSTEMS: YUV system (used by PAL in France) CMY system (used in color printers), HLS sytem (from Tektronix), HSY system.

  – WHAT IS IMPORTANT to note is that there are empirical formulas for transforming among these various coordinate systems.

6 Video Compression

- Compression is a necessity.

  Lossy versus non-lossy compression.

- Encoding (once) and decoding (many times) algorithms needed.

  Thus slow encoding is acceptable, but decoding must be fast.

- JPEG (Joint Photographic Experts Group) standard for still pictures

  This compression often produce 20:1 compression.

- JPEG ENCODING ALGORITHM:

  Assume RGB encoding of a 640 by 480 pixel image with 24 bits/pixel.

  – STEP 1. First convert to YIQ coordinates.

  Corresponding to Y,I,Q, we have 3 matrices of size 640 by 480, 8 bits per pixel.

  – STEP 2. The I and Q matrices are converted into 320 by 240 matrices by averaging each 2×2 subblock.

  – STEP 3. Each matrix is next subdivided into 8×8 subblocks. Apply the DCT (discrete cosine transform) to each subblock.

  Now, MOST of the values of the transformed subblocks will be close to zero.

  – STEP 4. Apply quantization based on a user supplied **weight table**. This table gives a larger weight to the LESS important coefficients. (Weight of 1 means most important)

  E.g., if a coefficient is x and the weight is w then we replace it by $\lceil x/w \rceil$.

  – STEP 5. The $(0, 0)$ coefficient is the DC Component. We replace this by its DIFFERENCE from the previous block. Most of the time, this difference is a small number (or 0).

  – STEP 6. Linearize the 64 coefficients of each block, using a anti-diagonal zigzag pattern, and THEN apply run-length encoding of this list.

  – STEP 7. Use Huffman encoding on the result.

- JPEG DECODING: run the process backwards (except for quantization)

- MPEG Standard (Motion Picture Experts Group)

  – MPEG-1 (up to 1.5 Mbps, CDROM applications) MPEG-2 (1.5 to 15 Mbps, Digital TV and DVD applications) MPEG-4 (1.5 to 15 Mbps, object-based compression)

  – Here, we treat only MPEG/video (ignore MPEG/audio).

  The popular MP3 is MPEG-1 Level 3 audio compression.

  – Exploit spatial and temporal coherence.

– 3 kinds of frames

I (**intracoded**): JPEG frame

P (**predictive**): Block-by-block difference from last I- or P-frame

B (**bidirectional**): Differences with last or next frame (I- or P-frames)

– E.g., $(IB_1B_2P_1B_3B_4P_2)$

Thus, B_1 and B_2 can refer to I or P_1.

– Typical bit rates: I-frames (1 bit per pixel), P-frames (0.1 bit per pixel), B-frames (0.015 bit per pixel).

– REORDERING OF FRAMES for transmission: this is needed to ensure that the decoder can do this with at most 3 frame buffers. In the above example, the transmission order would be:

E.g., $(IP_1B_1B_2P_2B_3B_4)$

– I-frames appear are inserted every half second or so.

– P-frame are based on **macro blocks** (16 by 16 for luminance, 8 by 8 for chrominance). Each macro block will try to find a "close enough" macro block in previous frame. If found, we just encode the displacement vector AND the difference from the previous block. OTHERWISE, if nothing is "close enough", just JPEG encoding.

– B-frames can refer to macro blocks in a previous OR succeeding I- or P-frames.

7 Real Time Scheduling

- Example: Processes A, B, C.

  A is every 30 msec (e.g., NTSC 30fps) for 10 msec,

  B is every 40 msec (e.g., PAL 25fps) for 15 msec,

  C is every 50 msec (e.g., slow NTSC 20fps) for 5 msec.

- Check schedulability (recall)

- Figure of independent scheduling (Fig 7-11 from text)

- Allow preemption in multimedia applications.

- RMS (Rate Monotonic Scheduling):

  – independent processes

  – realtime constraint for periodic processes

  – nonperioding processes have no deadlines

  – same CPU time needed by each process

  – preemption has no cost

- RMS ALGORITHM:
  - Each process gets a priority equal to its frequency (1/period)
  - Always run the highest priority ready process (use preemption)
  - Liu/Layland 1973: this is optimal for static scheduling