

Chapter 20: Multimedia Systems





Chapter 20: Multimedia Systems

- What is Multimedia
- Compression Techniques
- Requirements of Multimedia Kernels
- CPU Scheduling
- Disk Scheduling
- Network Management
- An Example: Cineblitz





Objectives

- To identify the characteristics of multimedia data
- To examine several algorithms used to compress multimedia data
- To explore the operating system requirements of multimedia data, including CPU and disk scheduling and network management





What is Multimedia?

- Multimedia data includes
 - audio and video clips (i.e. MP3 and MPEG files)
 - live webcasts
- Multimedia data may be delivered to
 - desktop PC's
 - handheld devices (PDAs, smart phones)





Media Delivery

- Multimedia data is stored in the file system like other ordinary data.
- However, multimedia data must be accessed with specific timing requirements.
- For example, video must be displayed at 24-30 **frames** per second. Multimedia video data must be delivered at a rate which guarantees 24-30 frames/second.
- **Continuous-media data** is data with specific rate requirements.





Streaming

- **Streaming** is delivering a multimedia file from a server to a client - typically the deliver occurs over a network connection.
- There are two different types of streaming:
 1. **Progressive download** - the client begins playback of the multimedia file as it is delivered. The file is ultimately stored on the client computer.
 2. **Real-time streaming** - the multimedia file is delivered to - but not stored on - the client's computer.





Real-time Streaming

- There are two types of real-time streaming:
 - (1) **Live streaming** - used to deliver a live event while it is occurring.
 - (2) **On-demand streaming** - used to deliver media streams such as movies, archived lectures, etc. The events are not delivered in real-time.





Multimedia Systems Characteristics

- Multimedia files can be quite large.
- Continuous media data may require very high data rates.
- Multimedia applications may be sensitive to timing delays during playback of the media.





Compression

- Because of the size and rate requirements of multimedia systems, multimedia files are often compressed into a smaller form.
- MPEG Compression:
 - (1) MPEG-1 - 352 X 240 @ 30 frames/second
 - (2) MPEG-2 - Used for compressing DVD and high-definition television (HDTV)
 - (3) MPEG-4 - Used to transmit audio, video, and graphics. Can be delivered over very slow connections (56 Kbps)





Operating Systems Issues

- The operating system must guarantee the specific data rate and timing requirements of continuous media.
- Such requirements are known as **Quality-of-Service (QoS)** guarantees.





QoS Guarantees

- Guaranteeing QoS has the following effects in a computer system:

(1) CPU processing

(2) Scheduling

(3) File systems

(4) Network protocols





Requirement of Multimedia Operating Systems

- There are three levels of QoS

(1) Best-effort service - the system makes a best effort with no QoS guarantees.

(2) Soft QoS - allows different traffic streams to be prioritized, however no QoS guarantees are made.

(3) Hard QoS - the QoS requirements are guaranteed.





Parameters Defining QoS

- Throughput - the total amount of work completed during a specific time interval.
- Delay - the elapsed time from when a request is first submitted to when the desired result is produced.
- Jitter - the delays that occur during playback of a stream.
- Reliability - how errors are handled during transmission and processing of continuous media.





Further QoS Issues


- QoS may be **negotiated** between the client and server.
- Operating systems often use an **admission control** algorithm that admits a request for a service only if the server has sufficient resources to satisfy the request.





Figure 20.1

Resources on a file server





CPU Scheduling

- Multimedia systems require **hard realtime** scheduling to ensure critical tasks will be serviced within timing deadlines.
- Most hard realtime CPU scheduling algorithms assign realtime processes static priorities that do not change over time.





Disk Scheduling

- Disk scheduling algorithms must be optimized to meet the timing deadlines and rate requirements of continuous media.
- Earliest-Deadline-First (EDF) Scheduling
- SCAN-EDF Scheduling





Disk Scheduling (cont)

- The EDF scheduler uses a queue to order requests according to the time it must be completed (its deadline.)
- SCAN-EDF scheduling is similar to EDF except that requests with the same deadline are ordered according to a SCAN policy.





Deadline and cylinder requests for SCAN-EDF scheduling

request	deadline	cylinder
A	150	25
B	201	112
C	399	95
D	94	31
E	295	185
F	78	85
G	165	150
H	125	101
I	300	85
J	210	90





Network Management

- Three general methods for delivering content from a server to a client across a network:
 - (1) **Unicasting** - the server delivers the content to a single client.
 - (2) **Broadcasting** - the server delivers the content to all clients, regardless whether they want the content or not.
 - (3) **Multicasting** - the server delivers the content to a group of receivers who indicate they wish to receive the content.





RealTime Streaming Protocol (RTSP)

- Standard HTTP is stateless whereby the server does not maintain the status of its connection with the client.





Figure 20.1

Streaming media from a conventional web server

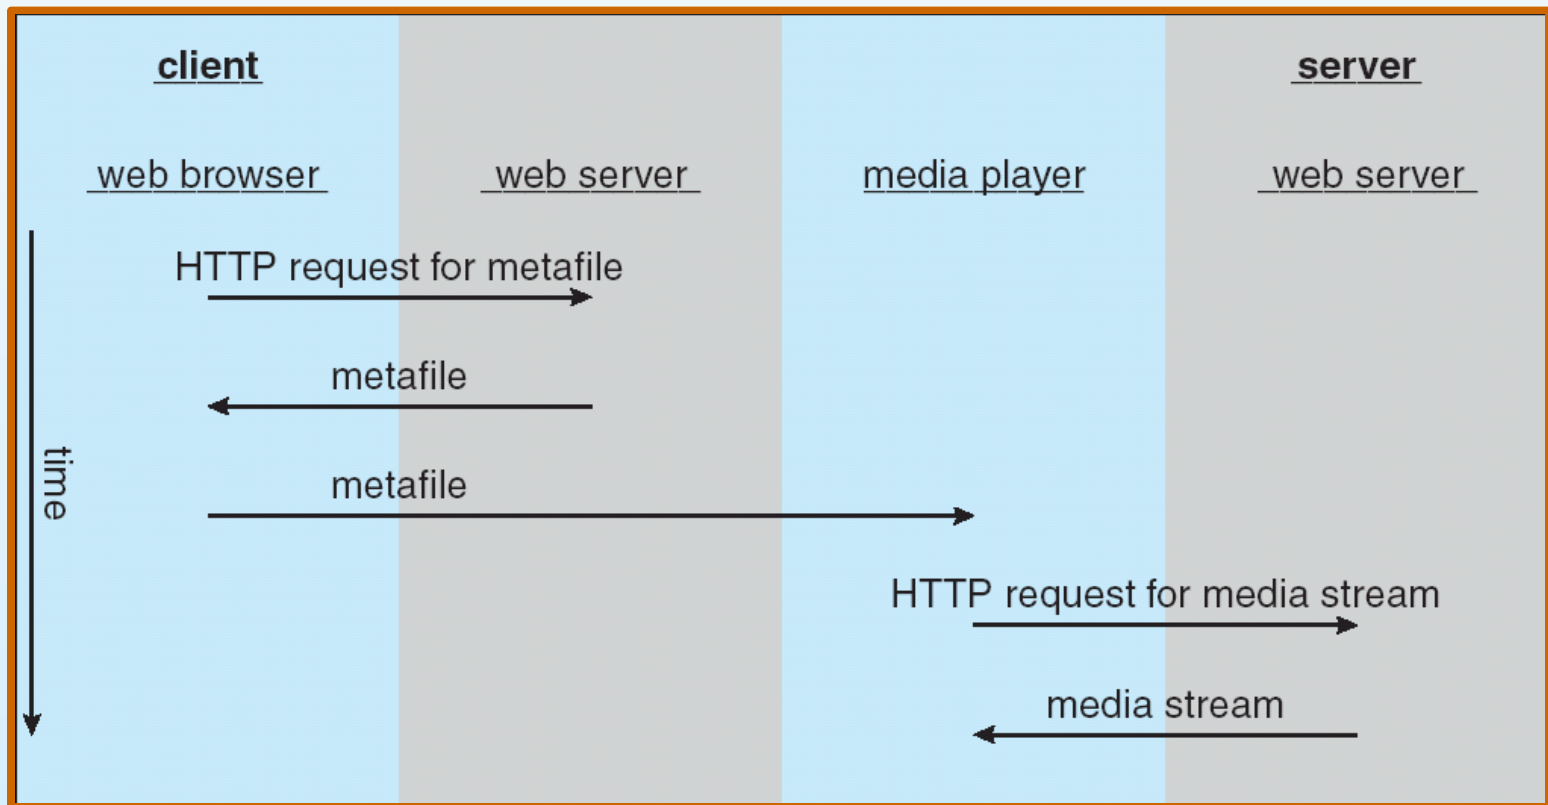
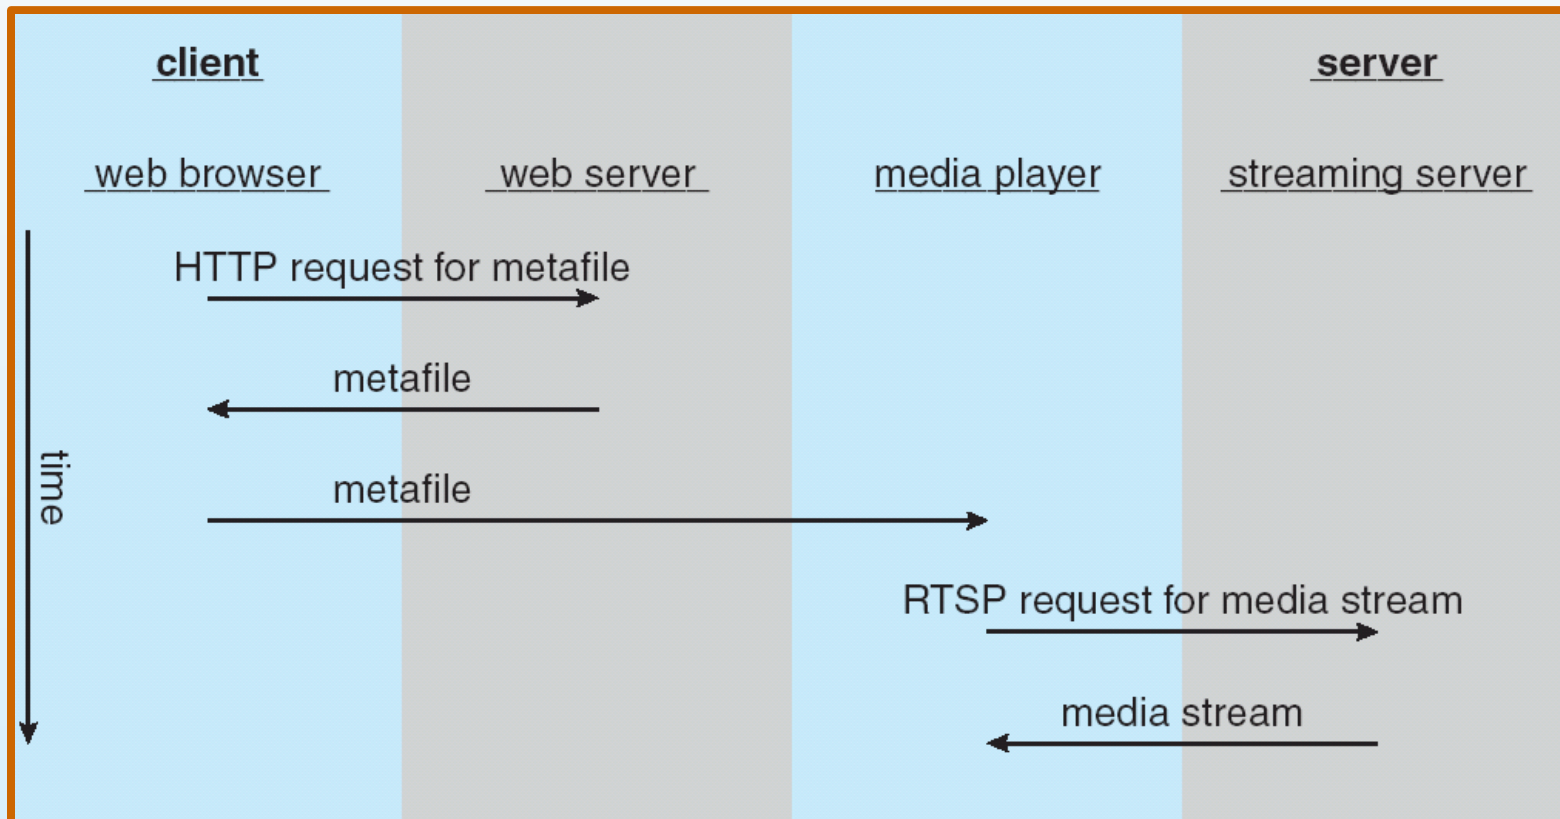




Figure 20.3

Realtime Streaming Protocol





RTSP States


- SETUP - the server allocates resources for a client session.
- PLAY - the server delivers a stream to a client session.
- PAUSE - the server suspends delivery of a stream.
- TEARDOWN - the server breaks down the connection and releases the resources allocated for the session.





Figure 20.4

RTSP state machine





CineBlitz Multimedia Server

- CineBlitz supports both realtime and non-realtime clients.
- CineBlitz provides hard QoS guarantees to realtime clients using an admission control algorithm.
- The disk scheduler orders requests using C-SCAN order.





CineBlitz Admission Controller

- Total buffer space required for N clients where client has rate requirement of r_i


$$\sum_{i=1}^N 2 \times T \times r_i \leq B.$$





Figure 20.05

Double buffering in CineBlitz





CineBlitz Admission Controller (cont)

- If t_{seek} and t_{rot} are the worst-case seek and rotational delay times, the maximum latency for servicing N requests is

$$2 \times t_{seek} + \sum_{i=1}^N \left(\left\lceil \frac{T \times r_i}{b} \right\rceil + 1 \right) \times t_{rot}.$$





CineBlitz Admission Controller (cont)

- The CineBlitz admission controller only admits a new client if there is at least $2 \times T \times r_i$ bits of free buffer space and the following equation is satisfied

$$2 \times t_{seek} + \sum_{i=1}^N \left(\left\lceil \frac{T \times r_i}{b} \right\rceil + 1 \right) \times t_{rot} + \sum_{i=1}^N \frac{T \times r_i}{r_{disk}} \leq T.$$





In.20.1

request	deadline	cylinder
A	150	25
B	201	112
C	399	95
D	94	31
E	295	185
F	78	85
G	165	150
H	125	101
I	300	85
J	210	90





Exercise 20.10

request	deadline	cylinder
R1	57	77
R2	300	95
R3	250	25
R4	88	28
R5	85	100
R6	110	90
R7	299	50
R8	300	77
R9	120	12
R10	212	2



End of Chapter 20

