CS 414 – Multimedia Systems Design

Lecture 28 – Media Server (Part 2)

Klara Nahrstedt

Spring 2008

Administrative

- MP3 is out – deadline April 4
- Discussion Section – Monday, March 31

Outline

- Multiple Disks and Multimedia
  - RAID
    - Data Striping
    - Group Creation
- Disk Management
  - Data Interleaving
  - Disk Scheduling
    - EDF
    - SCAN-EDF
    - Group Sweeping
    - Mixed Scheduling

Need for Multiple Disks Solutions for Media Server

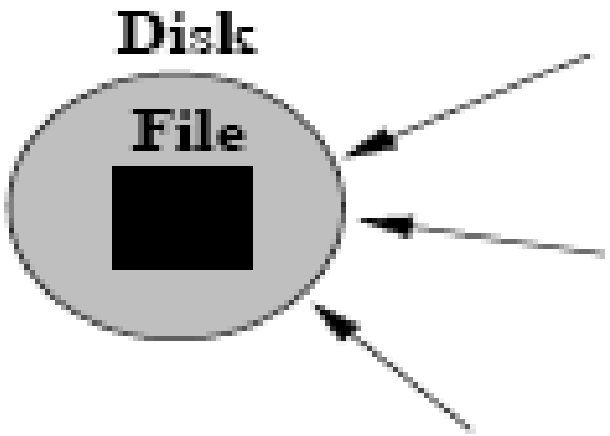

- Limitation of Single Disk: Disk Throughput
- Approach: 1 Maintain multiple copies of the same file on different disks
  - Very expensive
- Approach 2: Scatter multimedia file across multiple disks

Approach: Data Striping

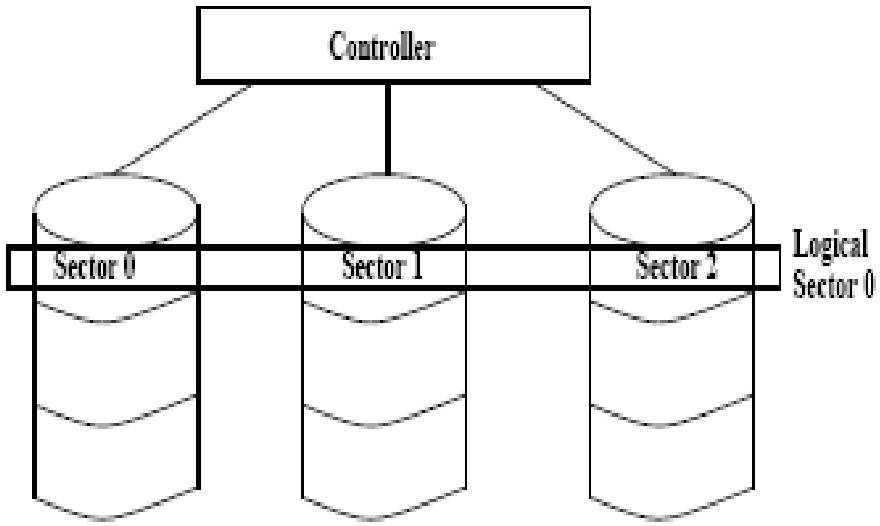

- RAID (Redundant Arrays of Inexpensive Disks)
  - Addresses both performance and security
  - (0-6) RAID levels – different approach at combining performance enhancements with security/fault-tolerance enhancements
- Disks spindle synchronously
  - Operate in lock-step parallel mode
- Striping improves BW, but does not improve seek or rotational delay

Data Striping – Group Creation

Multiple RAID: Creation of Subgroups of disks into independent logical disk arrays; limits # of disks per file

Declustering: Groups are not made up of complete disks ; # of disks for any stripe is fixed and of same size, but disks on which stripe is located differs

Dynamic Declustering: Non - static strip allocation to disks

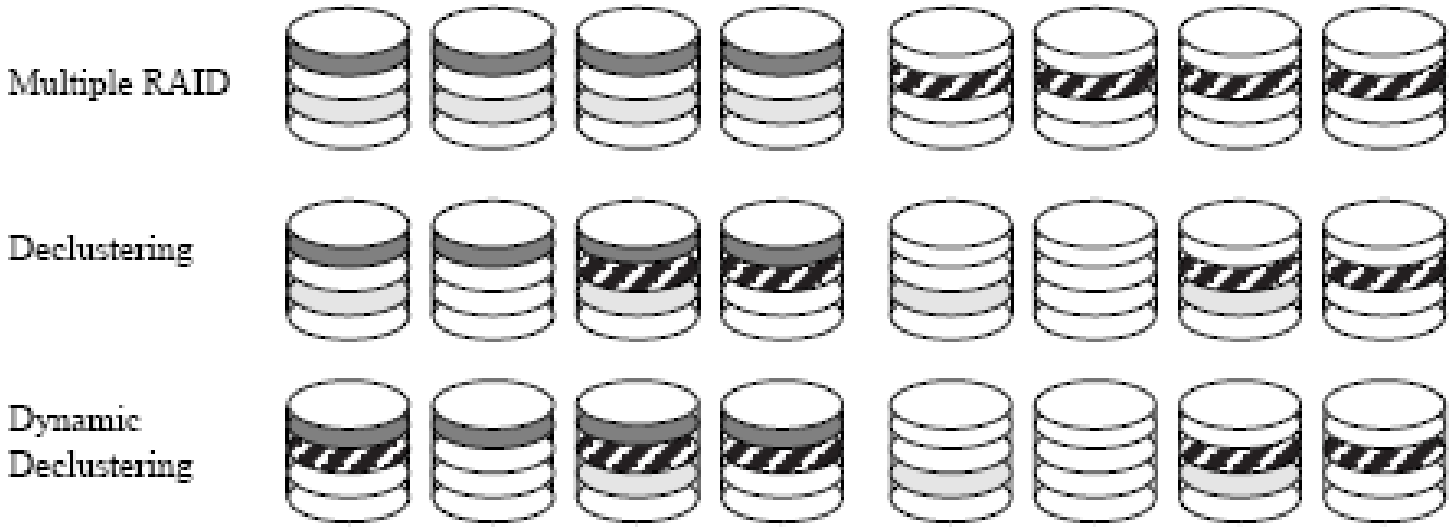

Storage/Disk Management

- Disk access – slow and costly
- Reduce disk access
  - Use block caches (anticipate future reads or writes)
  - Reduce disk arm motion
    - Blocks accesses in sequence (continuously) , place together on one cylinder
    - Interleaved vs non-interleaved storage

Data Interleaving

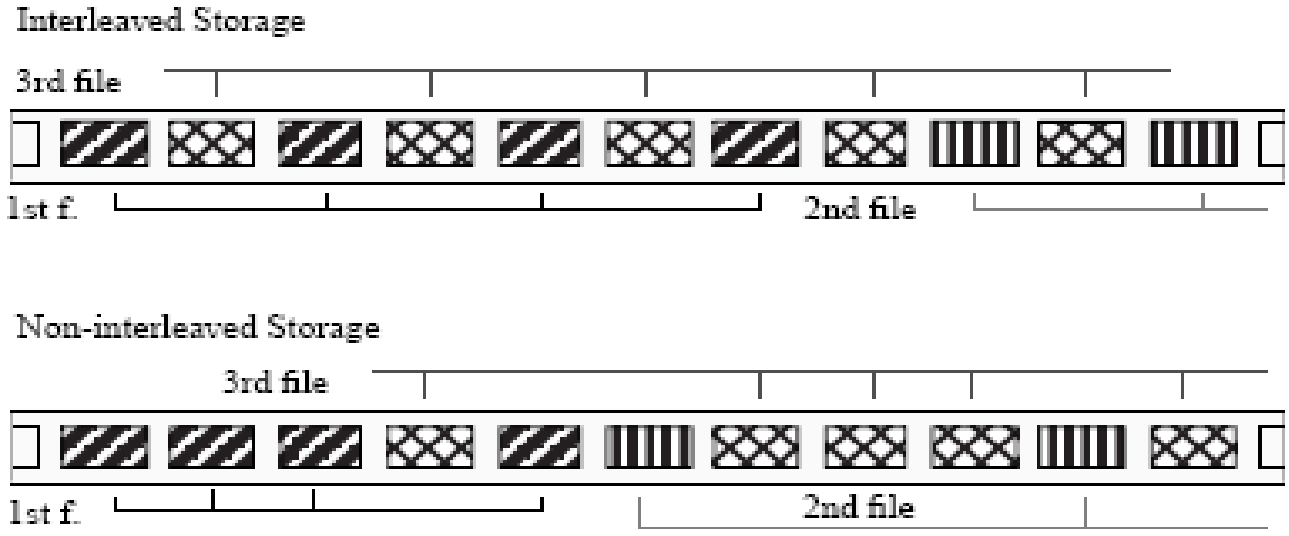

Data Interleaving On single disk
(consecutive blocks are placed on The same cylinder But in interleaved way

Round k	Disk 1	Disk 2	Disk 3
1	File A, block 1	File B, block 1	File C, block 1
2	File C, block 2	File A, block 2	File B, block 2
3	File B, block 3	File C, block 3	File A, block 3
4	File A, block 4	File B, block 4	File C, block 4

Data Interleaving On Multiple Disks
(Disks are not Synchronized)

Disk Scheduling Policies

- Goal of Scheduling in Traditional Disk Management
  - Reduce cost of seek time
  - Achieve high throughput
  - Provide fair disk access
- Goal of Scheduling in Multimedia Disk Management
  - Meet deadline of all time-critical tasks
  - Keep necessary buffer requirements low
  - Serve many streams concurrently
  - Find balance between time constraints and efficiency

EDF (Earliest Deadline First) Disk Scheduling

- Each disk block request is tagged with deadline
- Policy:
  - Schedule disk block request with earliest deadline
  - Excessive seek time – high overhead
  - Pure EDF must be adapted or combined with file system strategies

EDF Example

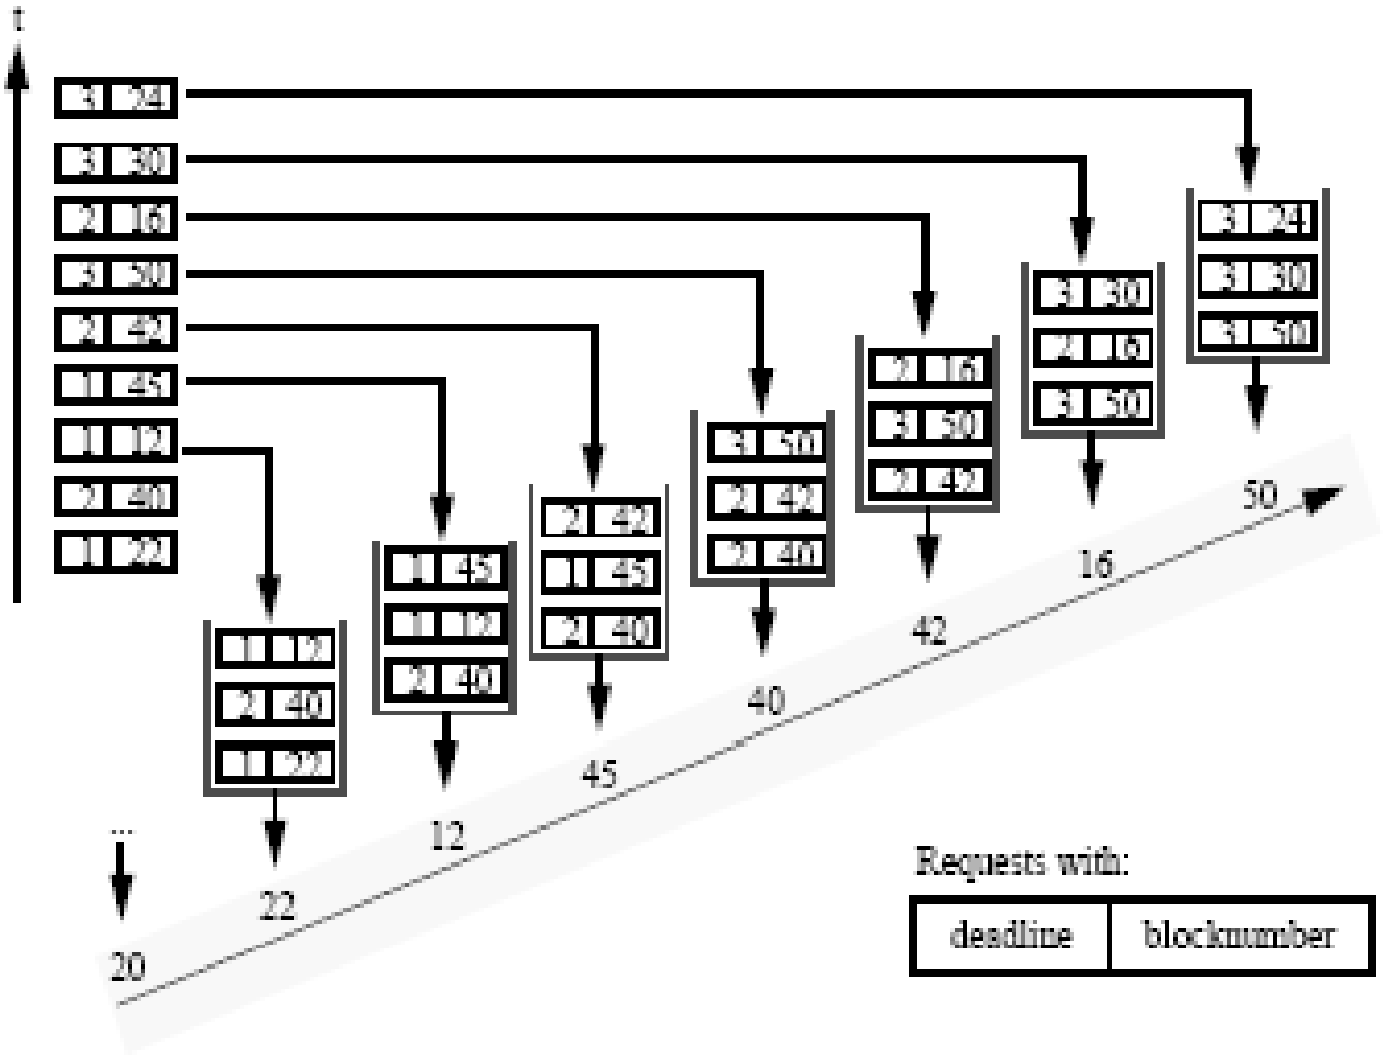

SCAN-EDF Scheduling Algorithm

- Combination of SCAN and EDF algorithms
- Each disk block request tagged with augmented deadline
  - Add to each deadline perturbation
- Policy:
  - SCAN-EDF chooses the earliest deadline
  - If requests with same deadline, then choose request according to scan direction

Implementation of SCAN-EDF

- Notation:
  - D_i be deadline of disk block request '*i*'
  - N_i be track position on disk
  - N_{max} be maximum number of tracks
- Deadline Modification:
  - $D_i + f(N_i)$
  - $f(N_i)$ converts track number of '*i*' into a small perturbation of deadline
  - Perturbation small enough so that
    - **$D_i + f(N_i) \leq D_j + f(N_j)$ for $D_i \leq D_j$**
    - **Possible $f(N_i) = N_i/N_{max}$**

SCAN EDF Example (N_{max} = 100)

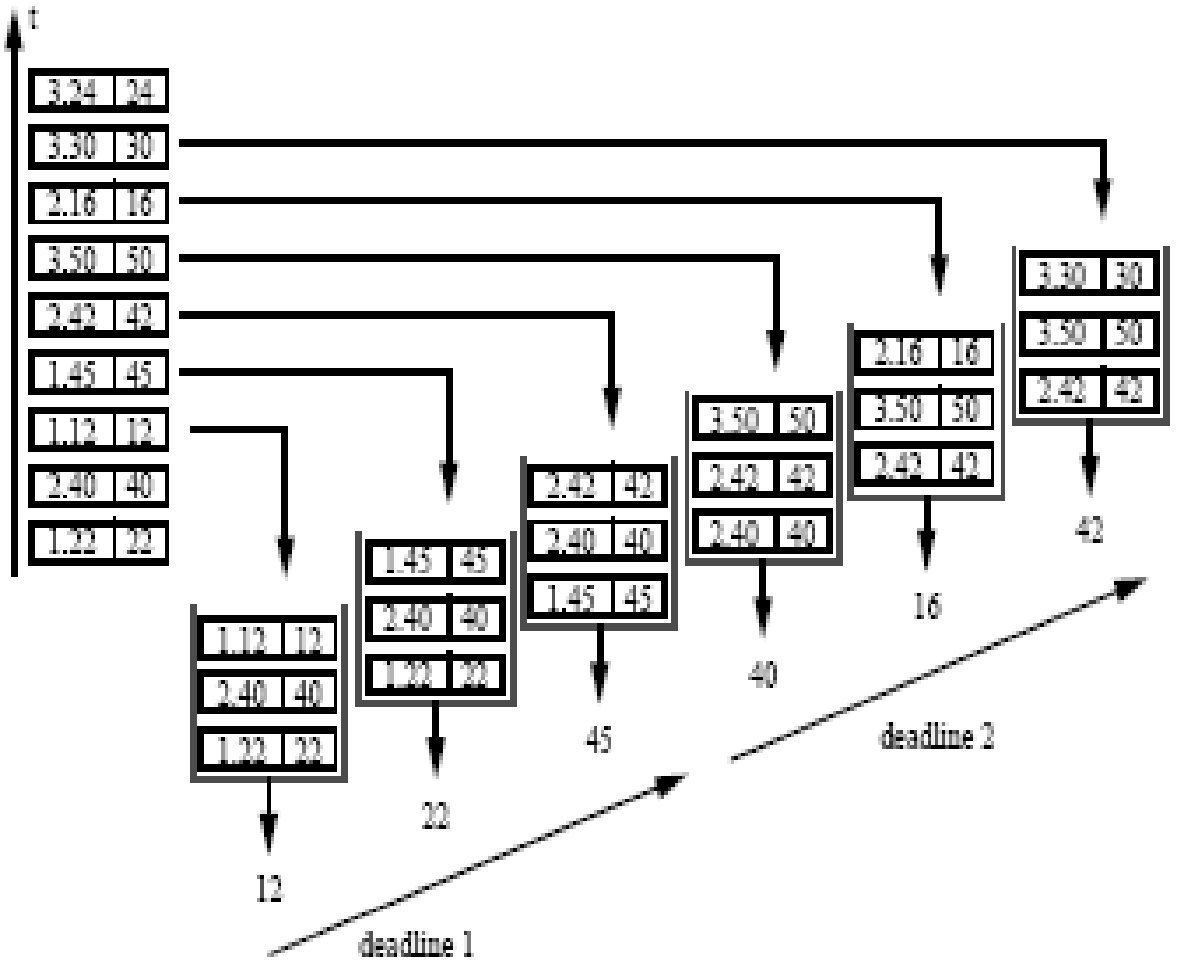

Enhanced SCAN-EDF (1)

- Use more accurate perturbation of deadline
- Consider
  - Actual track position of disk head *'N'*
  - N_{max} – max number of disk tracks
  - N_i – next track to be considered

Head Moves Upwards

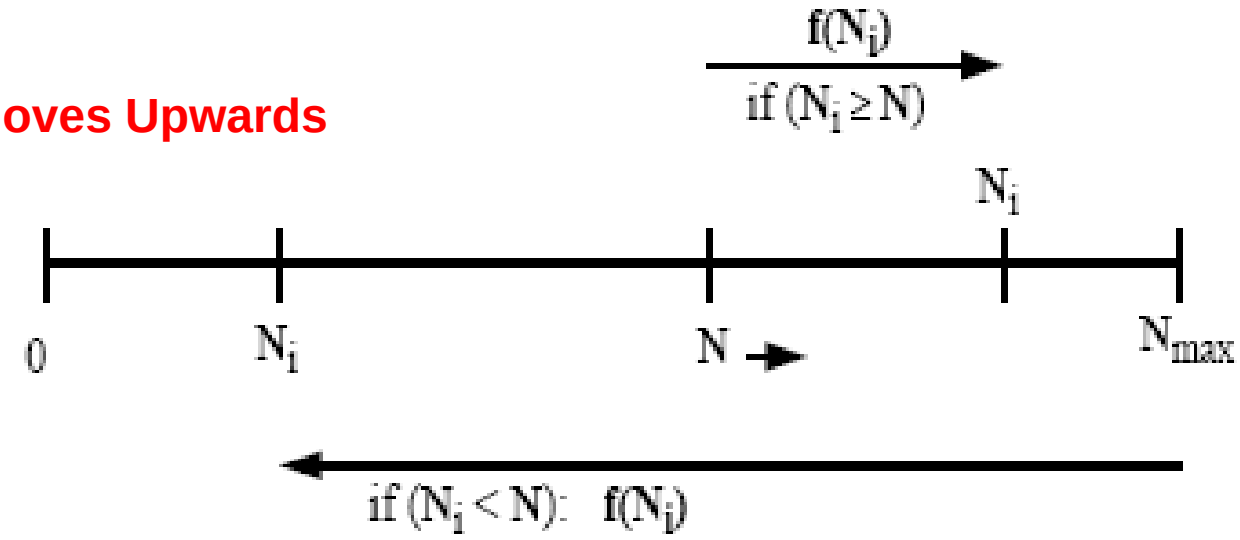

Enhanced SCAN-EDF (2)

- Algorithm:
  - If head **moves upwards** (towards N_{max}), then
  - (a)

$$\forall N_i; N \le N_i \le N_{\max}, f(N_i) = \frac{N_i - N}{N_{\max}}$$

  - (b)

$$\forall N_i; 1 \le N_i < N, f(N_i) = \frac{N_{\max} - N_i}{N_{\max}}$$

Enhanced SCAN-EDF (3)

- If head **moves downwards** (towards 1), then

(a) $\forall N_i; N < N_i \leq N_{\max} : f(N_i) = \frac{N_i}{N_{\max}}$

(b) $\forall N_i; 1 \leq N_i \leq N : f(N_i) = \frac{N - N_i}{N_{\max}}$

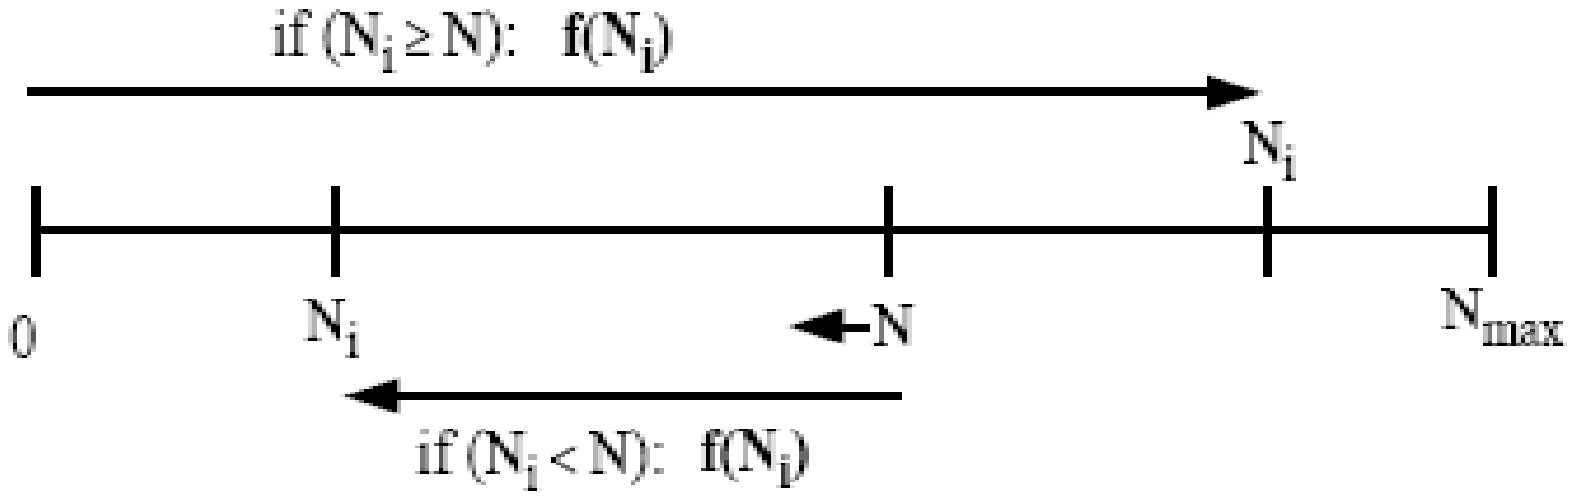

Conclusion

- Disk Scheduling – important component in the timely delivery of streams
- Admission should be done if one cares not to over subscribe