Throughput-Competitive Admission Control for Continuous Media Databases

Minos N. Garofalakis *(Univ. of Wisconsin – Madison)*

Yannis E. Ioannidis *(Univ. of Wisconsin – Madison)*
Banu Özden *(Bell Laboratories)*
Avi Silberschatz *(Bell Laboratories)*

Outline

- Introduction and Motivation
- Problem Formulation
- Contributions
  1. Algorithms and Theoretical Results
  2. Experimental Validation
- Conclusions

Introduction and Motivation

- *Continuous Media*: Real-time resource requirements for delivery

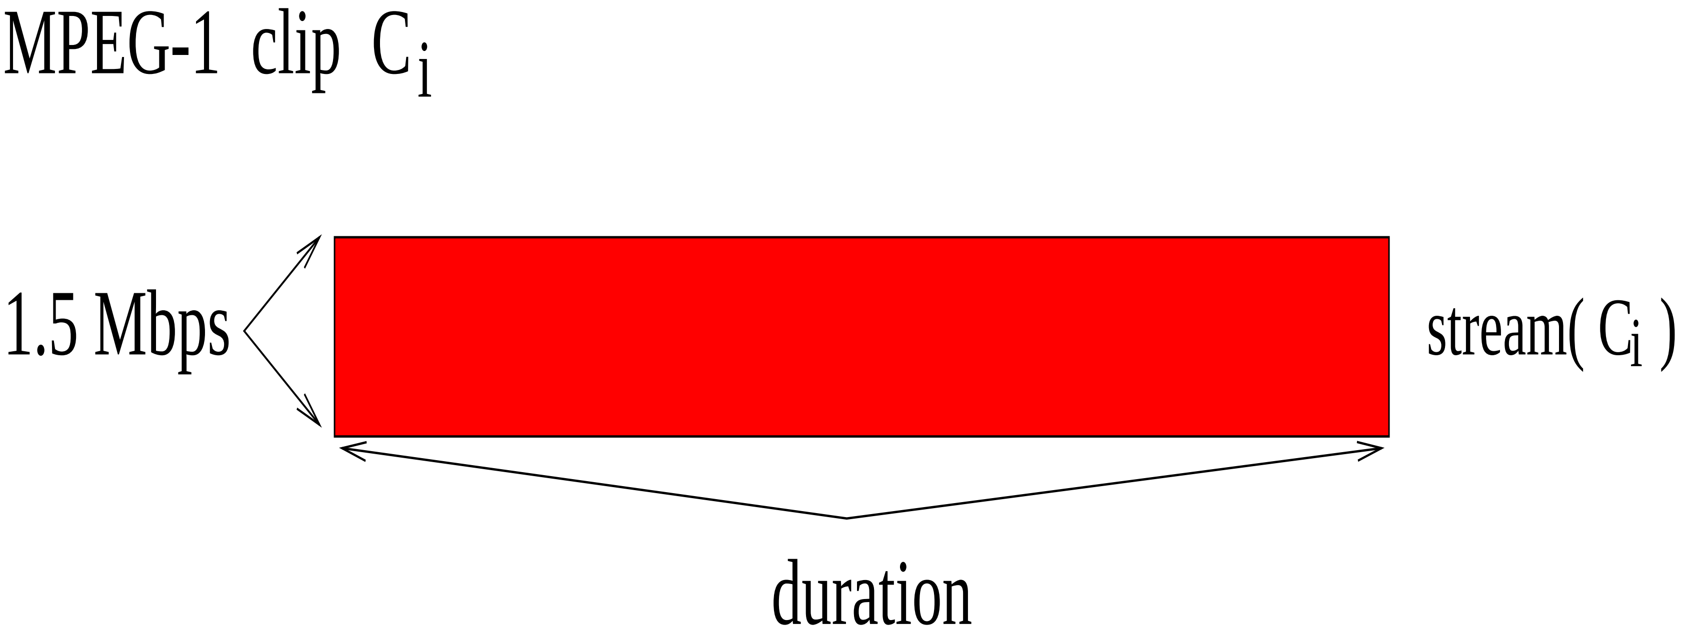

- Effective resource management for on-demand support

Introduction and Motivation (cont.)

- *Admission Control*: provide service guarantees for accepted streams

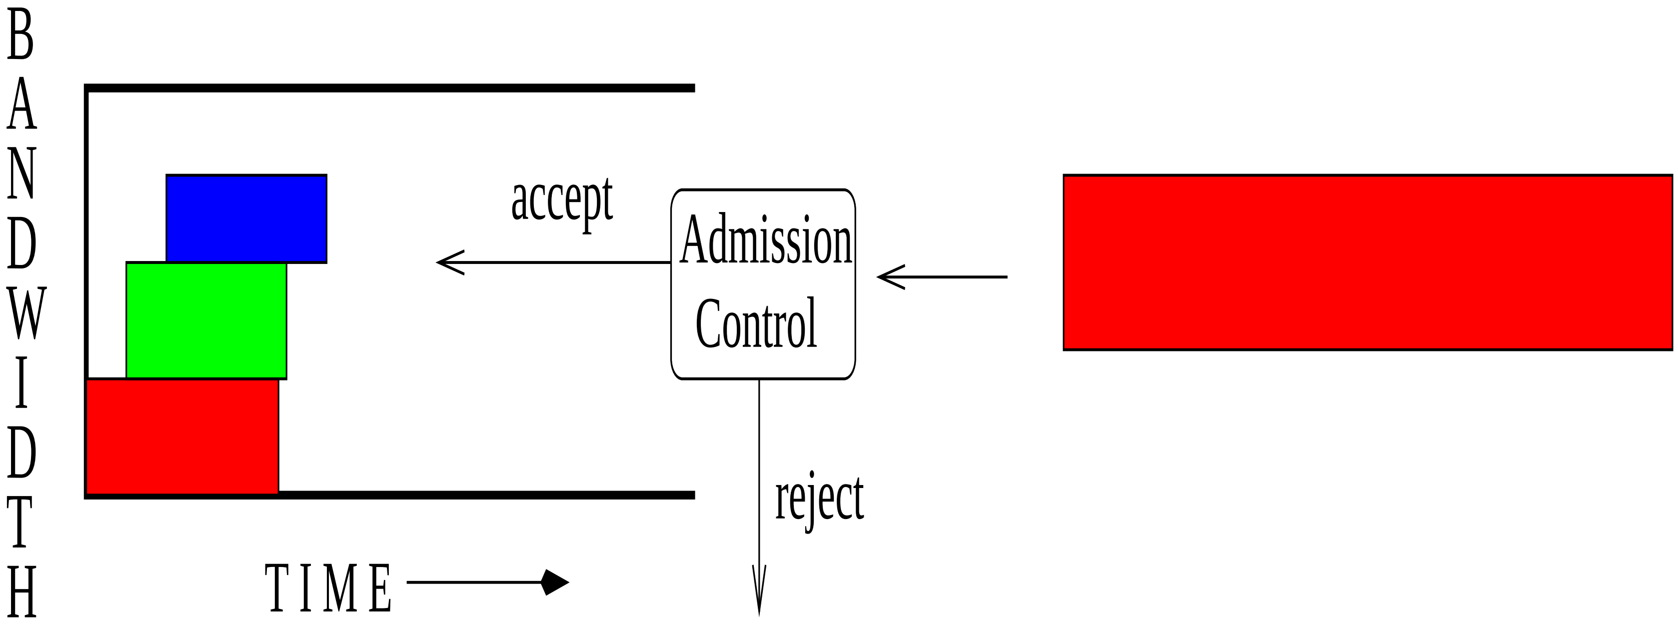

- *On-line* decision making – non-preemptive!
- Crucial for high server utilization

Problem Formulation

- Study the implications of the *on-line* nature of admission control
- Admission Control Objective: Maximize total server throughput over a sequence of requests
- Metric: *Competitive ratio*, i.e., bound on worst-case behavior over any possible sequence – very robust!

Problem Formulation (cont.)

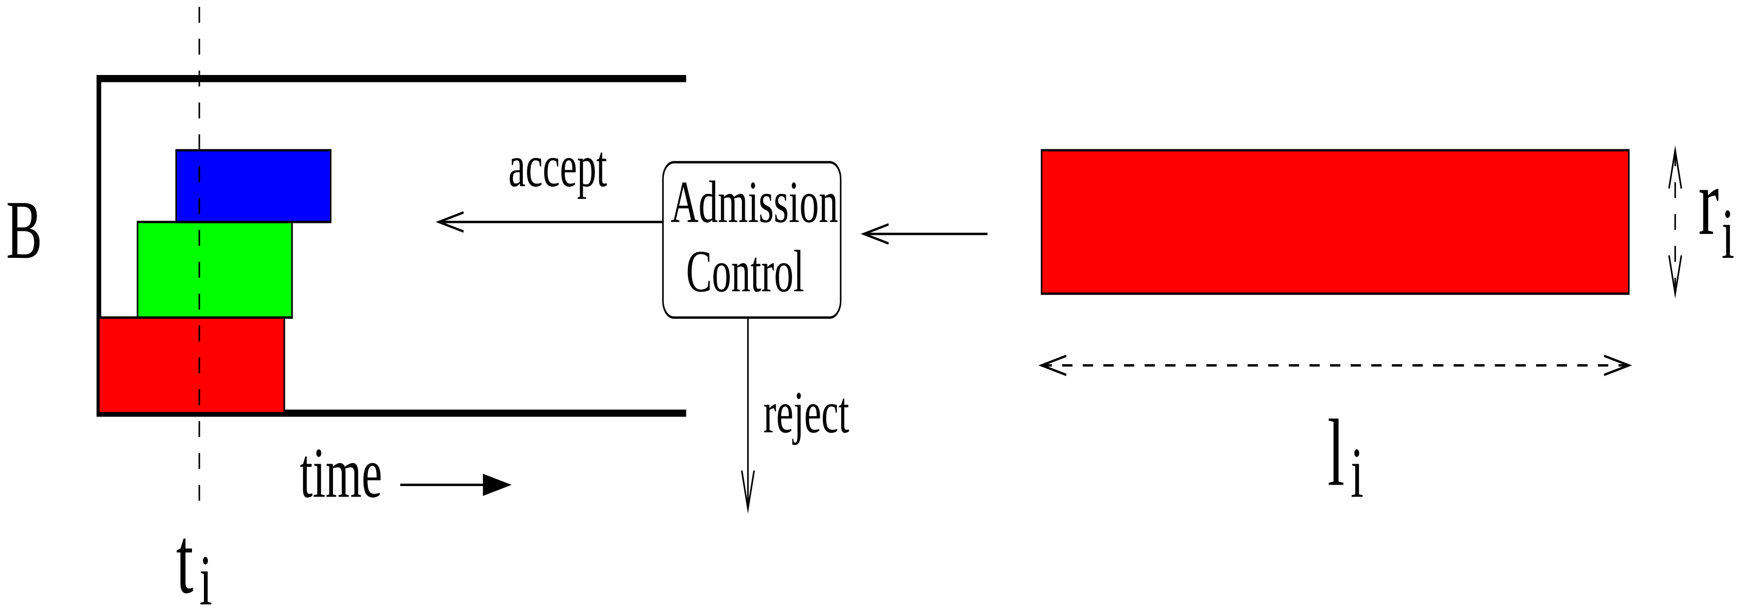

- $l_{max} = \max_i\{l_i\}$, $l_{min} = \min_i\{l_i\}$, $\Delta = \frac{l_{max}}{l_{min}}$
- $r_{max} = \max_i\{r_i\}$, $\rho = \frac{r_{max}}{B}$
- Competitive ratio of algorithm $A = \sup_{sequences} \frac{\sum_{OPT} l_i \cdot r_i}{\sum_A l_i \cdot r_i}$
- Want small (i.e., close to 1) ratios

Our Contributions

- Competitive ratio of conventional, "greedy" adm. control is $\Theta\left(\frac{\Delta}{1-\rho}\right)$
- Lower bounds of $\Omega\left(\frac{\log\Delta}{1-\rho}\right)$ on the competitive ratio of any deterministic *or* randomized strategy
- Novel strategies, based on *Bandwidth Prepartitioning* with ratios of $O(\log\Delta)$ (i.e., near-optimal) for large B
- Strategies can easily be adapted to accommodate clip popularities to improve average case
- Experimental validation

Conventional Wisdom: "Greedy" Admission Control

- *Work-Conserving strategy* ($\mathcal{WC}$): admit incoming request if $\exists$ sufficient bandwidth, reject otherwise

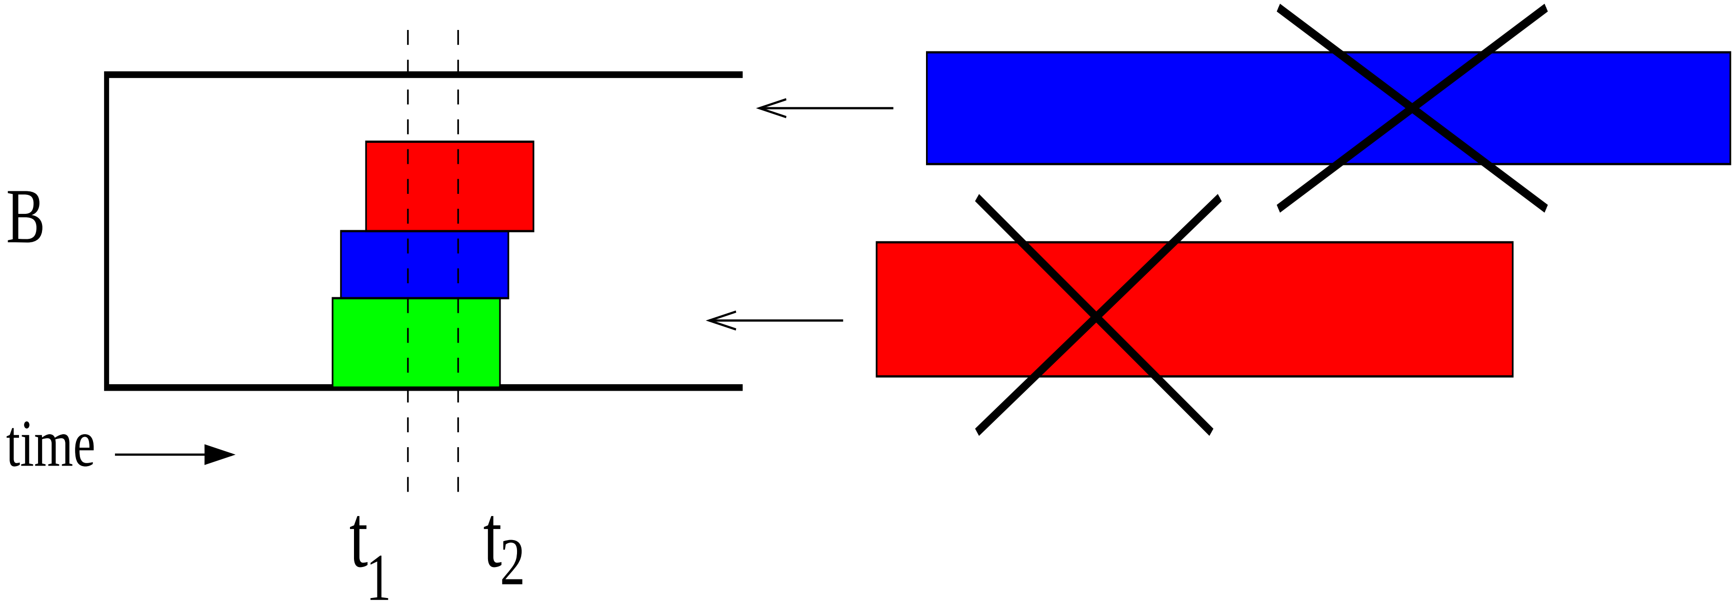

Theorem: $\mathcal{WC}$ is $\frac{1+\Delta}{1-\rho}$-competitive

Lower Bounds on Competitiveness

Theorem: Any deterministic *or* randomized on-line admission control strategy has a competitive ratio of $\Omega\left(\frac{\log \Delta}{1-\rho}\right)$

- Exponential gap w.r.t. $\mathcal{WC}$ — lots of space for improvement!

The Basic Idea: Simple Prepartitioning

- Idea: isolate requests with large length differences so that short requests cannot monopolize the server

Simple Bandwidth Prepartitioning ($\mathcal{SBP}$):

1. Divide bandwidth B into $\lceil \log \Delta \rceil$ partitions, of size $\frac{B}{\lceil \log \Delta \rceil}$ each
2. Schedule requests with lengths in $[2^{i-1} \cdot l_{min}, 2^i \cdot l_{min})$ in the i^{th} partition using $\mathcal{WC}$

(Note that $\Delta \leq 2$ *within each partition*)

$\mathcal{SBP}$ Example

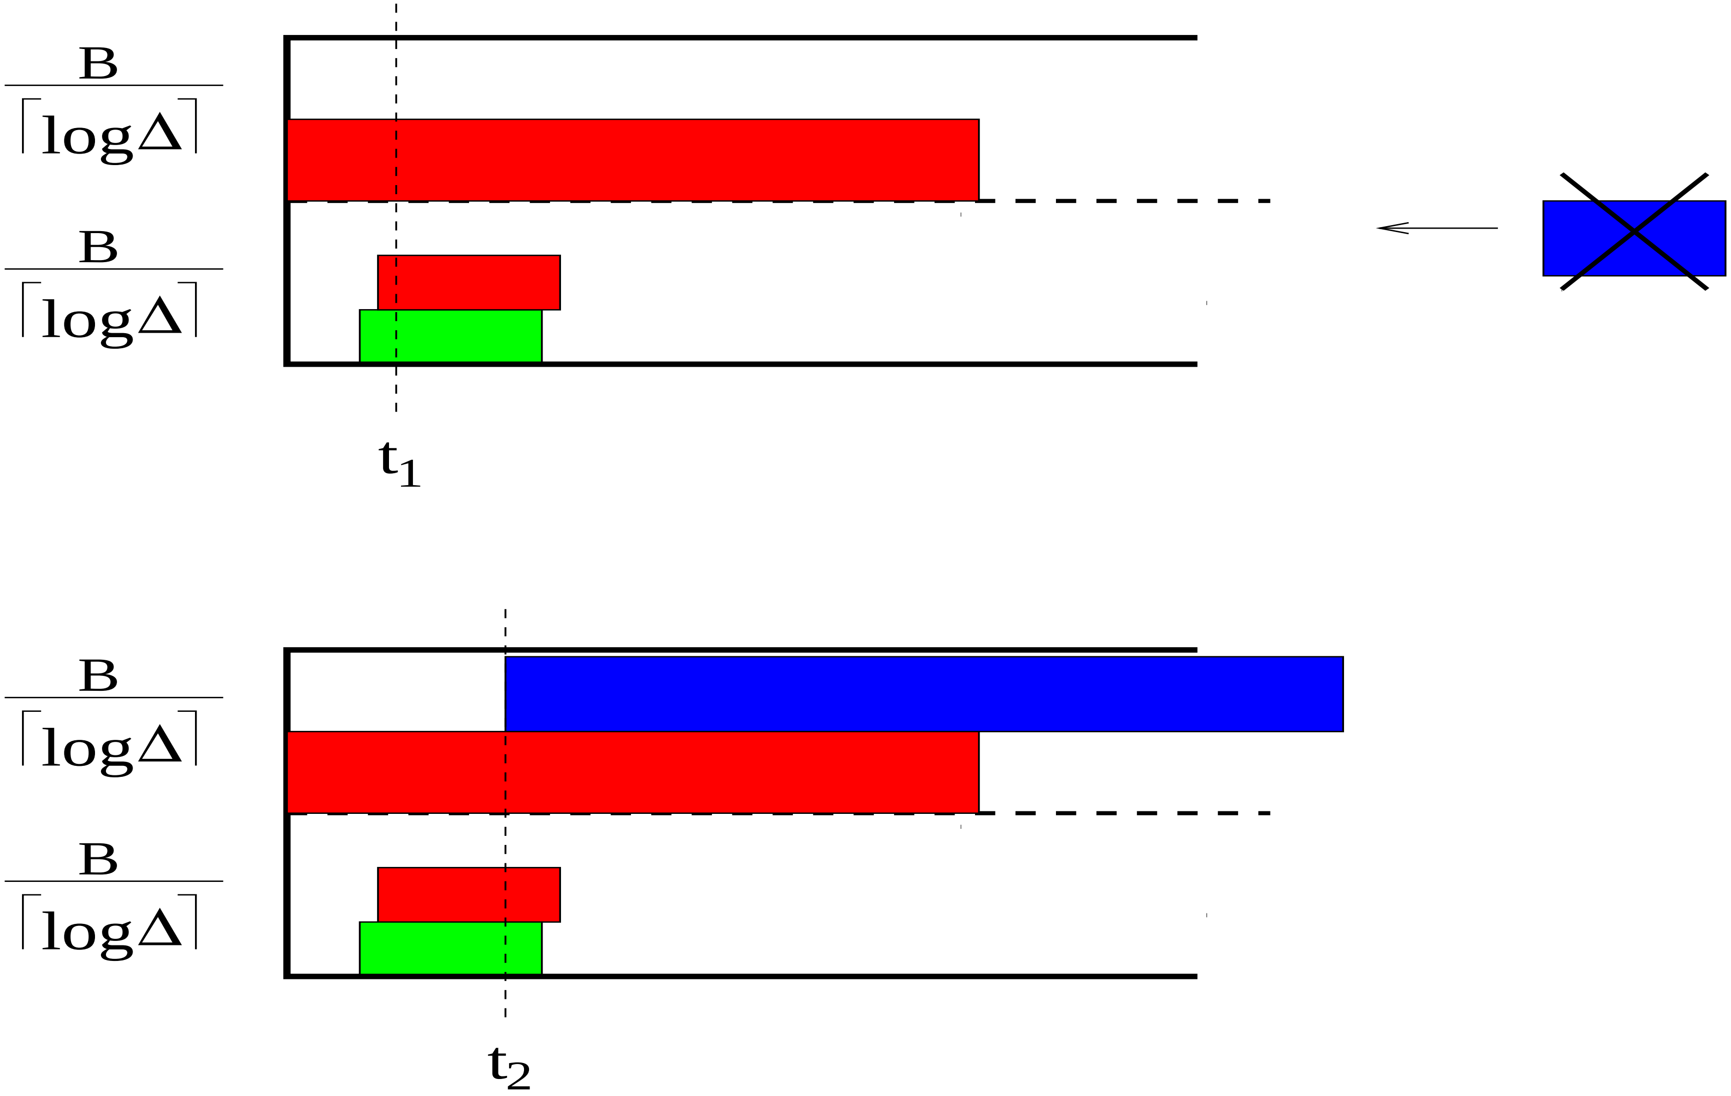

Reducing Fragmentation: Down-shift Prepartitioning

- Idea: again, prohibit short requests from monopolizing the server *but also* allow long requests to “steal” from lower partitions

Down-shift Bandwidth Prepartitioning ($\mathcal{DBP}$):

1. Divide bandwidth B into $\lceil \log \Delta \rceil$ partitions $|B_i| = \frac{B}{\lceil \log \Delta \rceil}$
2. Schedule requests with lengths in $[2^{i-1} \cdot l_{min}, 2^i \cdot l_{min})$ in $B_1 \cup \ldots \cup B_i$ using $\mathcal{WC}$

- Major benefit: reduce effects of bandwidth fragmentation due to prepartitioning

$\mathcal{DBP}$ Example

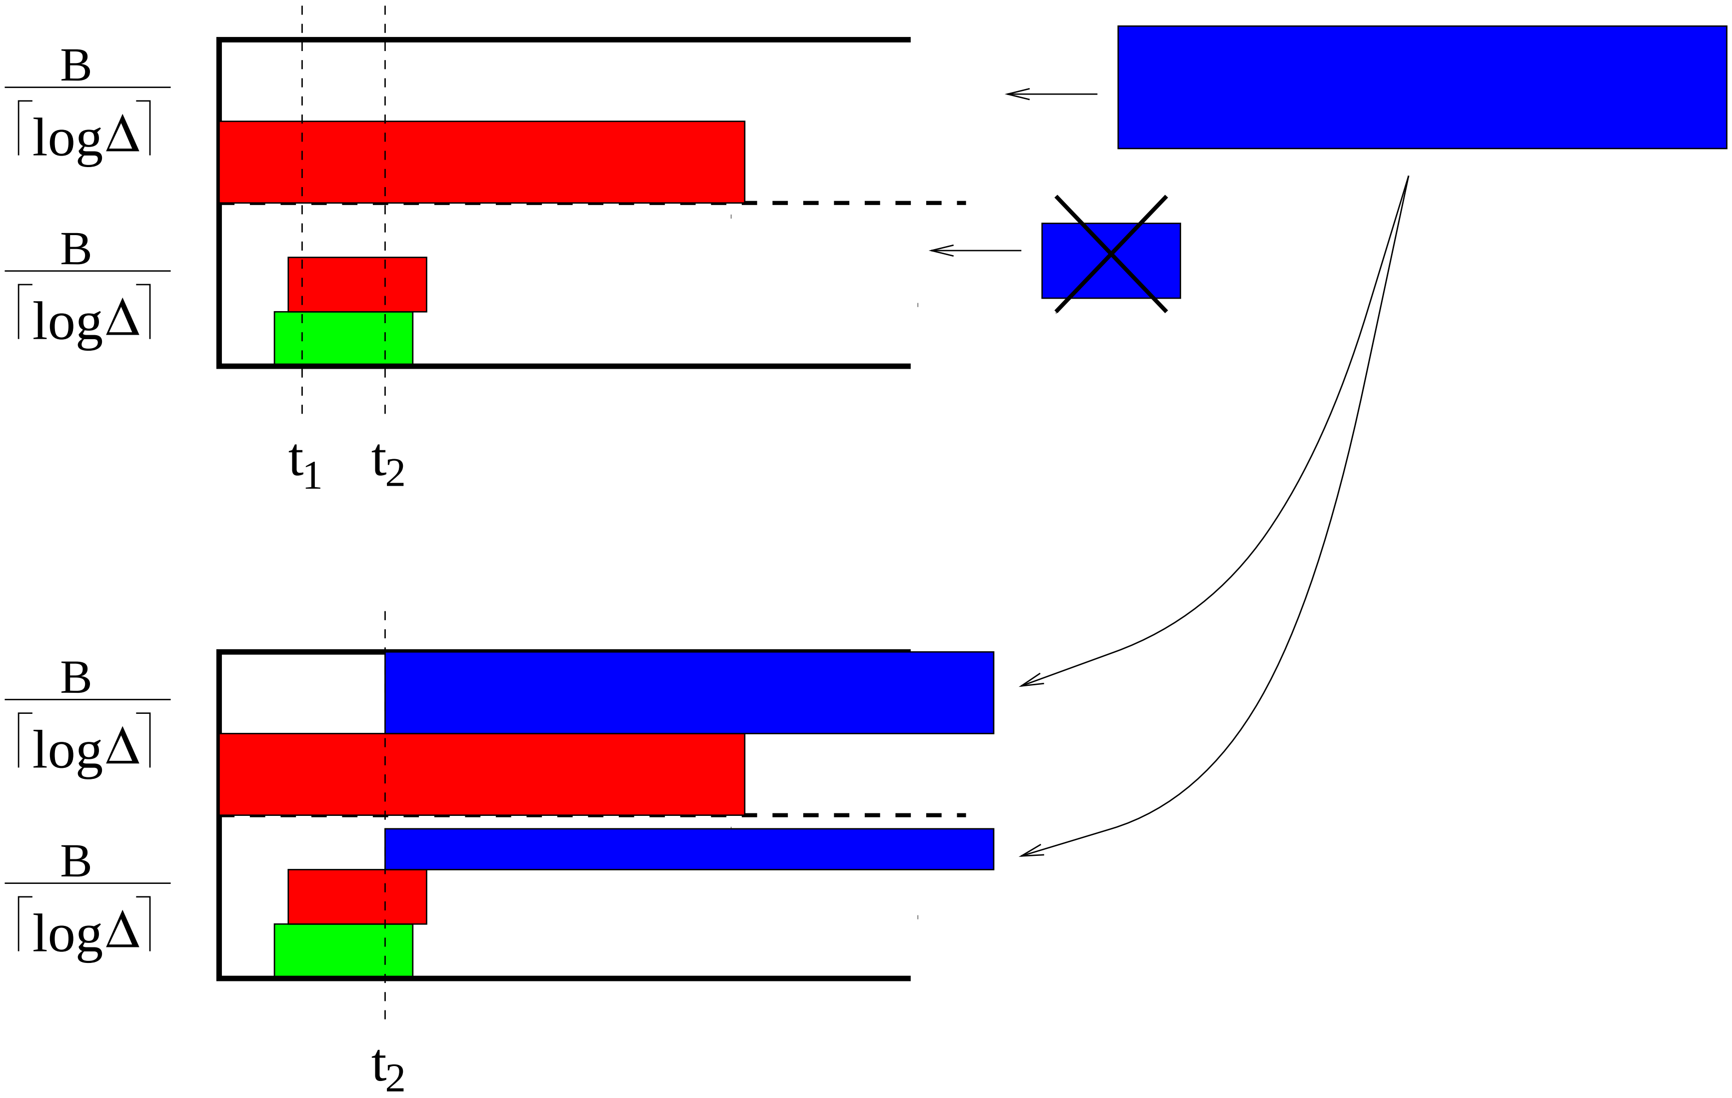

Competitiveness of Bandwidth Prepartitioning

Theorem:

(a) Assuming $\rho = \frac{r_{max}}{B} < \frac{1}{\lceil \log \Delta \rceil}$, $\mathcal{SBP}$ and $\mathcal{DBP}$ are $O\left(\frac{\log \Delta}{1 - \rho \cdot \log \Delta}\right)$ - competitive

(b) If $\rho < \frac{1}{2 \cdot \lceil \log \Delta \rceil}$, then $\mathcal{SBP}$ and $\mathcal{DBP}$ are $O\left(\log \Delta\right)$ - competitive

- Realistic assumptions. E.g., if $r_{max} = 8$ Mbps, $l_{max} = 120 \cdot l_{min}$, then $r_{max} \cdot \lceil \log \Delta \rceil = 56$ Mbps ($<< B$)
- For reasonably large B, $\mathcal{SBP}$ and $\mathcal{DBP}$ are *near-optimal*

Better Average Case: Popularity-based Prepartitioning

- Idea: employ clip *popularities* to define partition sizes (simple $\mathcal{SBP}$ or $\mathcal{DBP}$ may underutilize the server)
- p_j = cumulative popularity of clips with length l_j
- $PL_i \;=\; \{\, (p_j, l_j) \,|\, l_j \in [2^{i-1} \cdot l_{min}, 2^i \cdot l_{min}) \,\}$

Popularity-based Bandwidth Prepartitioning ($\mathcal{PBP}$):

- Same as $\mathcal{DBP}$, but define the size of the i^{th} partition as:

$$|B_i| \;=\; \frac{f(PL_i)}{\sum_i f(PL_i)} \cdot B$$

for some function f

Experimental Results

- Goal: Prepartitioning was based on competitive (worst-case) analysis — what about average-case?
- $\mathcal{WC}$ vs. $\mathcal{PBP}$ (results shown with $f(PL_i) = \sum_{PL_i} p_j \cdot l_j$)
- *Arrival Process* : Poisson , Bursty , Poisson + Short Bursts – capture short-term overloads (e.g., 6 o'clock news)
- *Popularities*: Zipfian
  Length-Popularity Correlation : Positive , Negative , Random
- *Metric* : fraction of server capacity utilized

$$\frac{\text{scheduled} \sum l_i \cdot r_i}{\text{B} \times \text{simulation time}}$$

Experimental Results (cont.)

- Poisson Arrivals

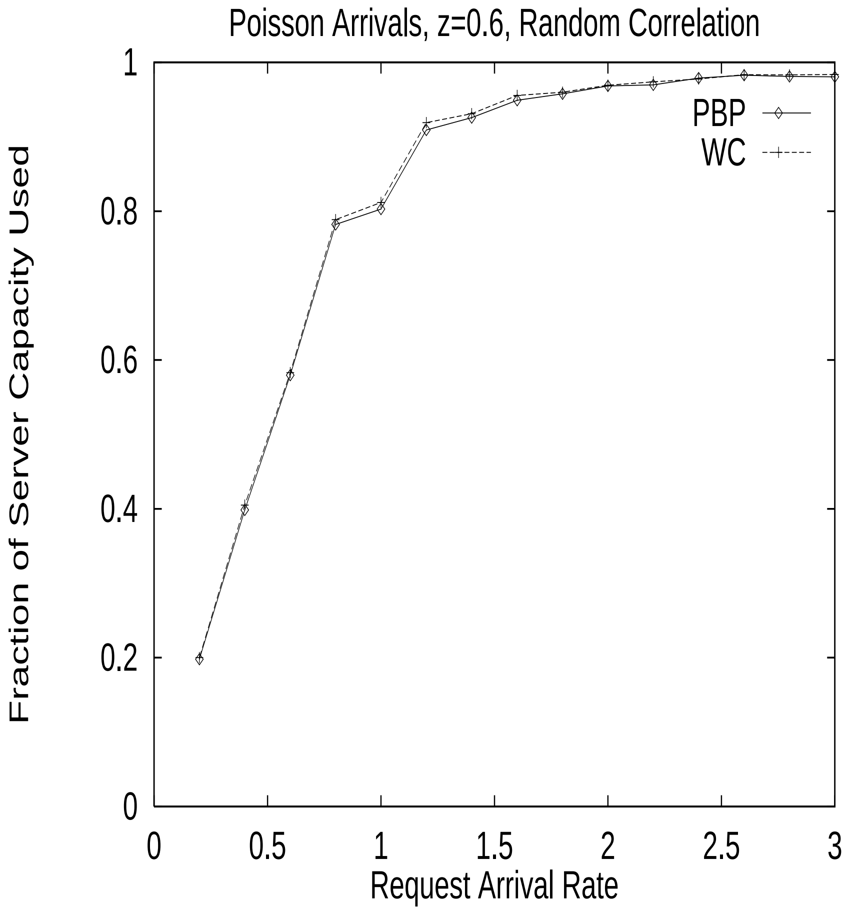

- Bursty Arrivals, Random Correlation

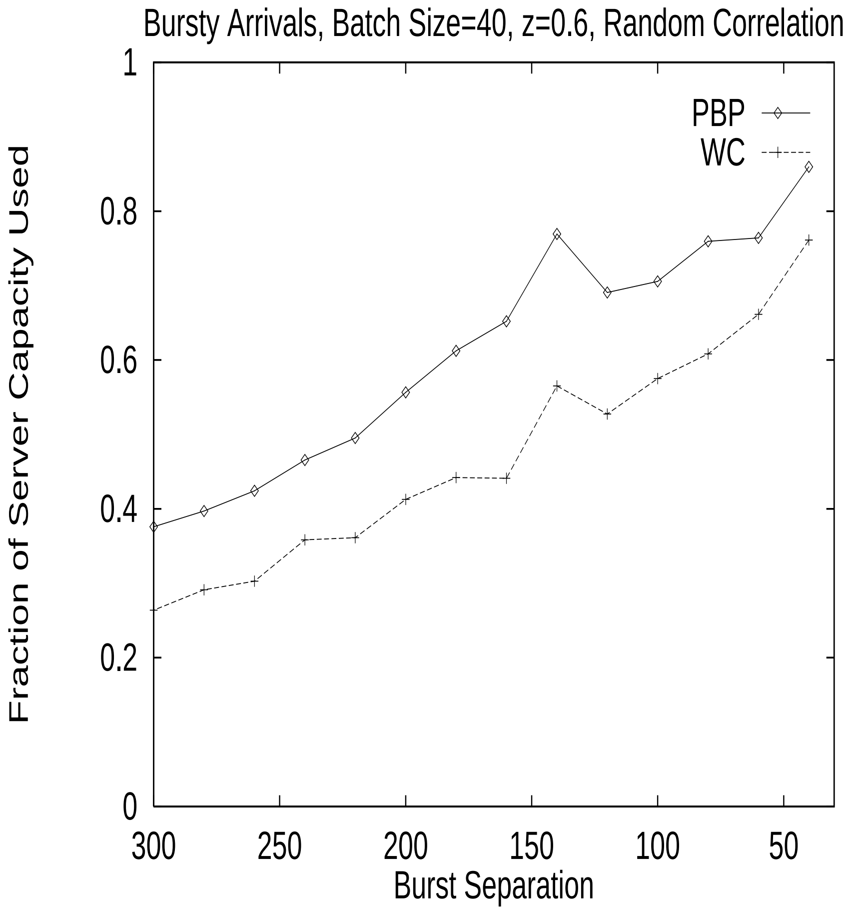

Experimental Results (cont.)

- Bursty Arrivals, Negative Correlation
- Poisson+Short Bursts

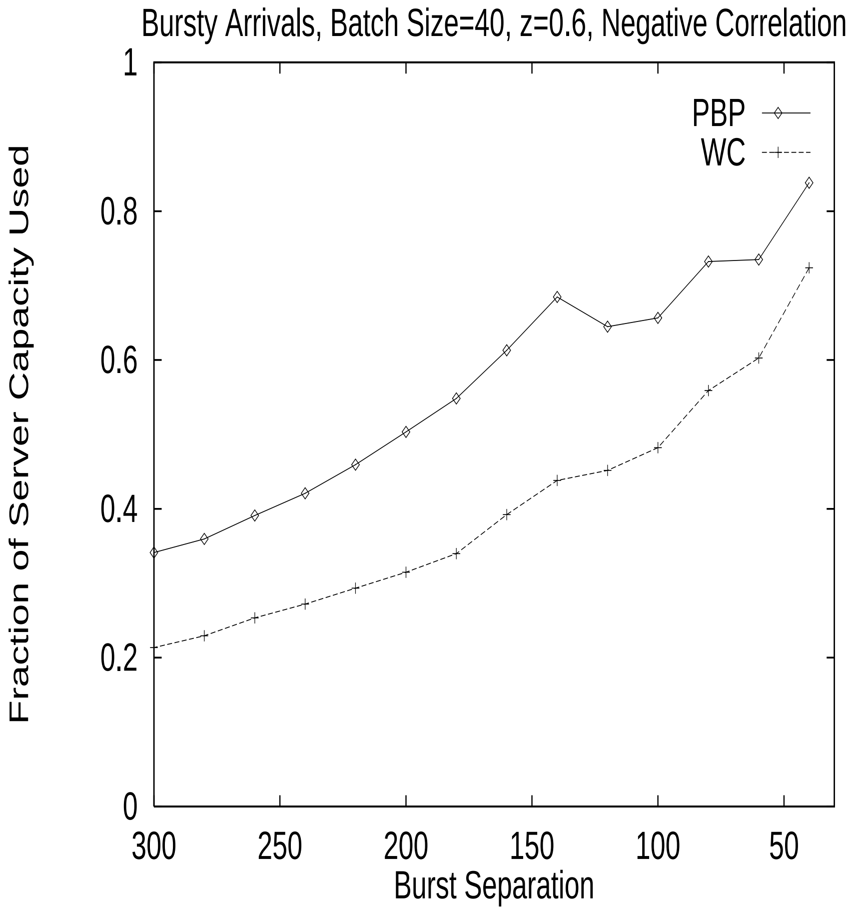

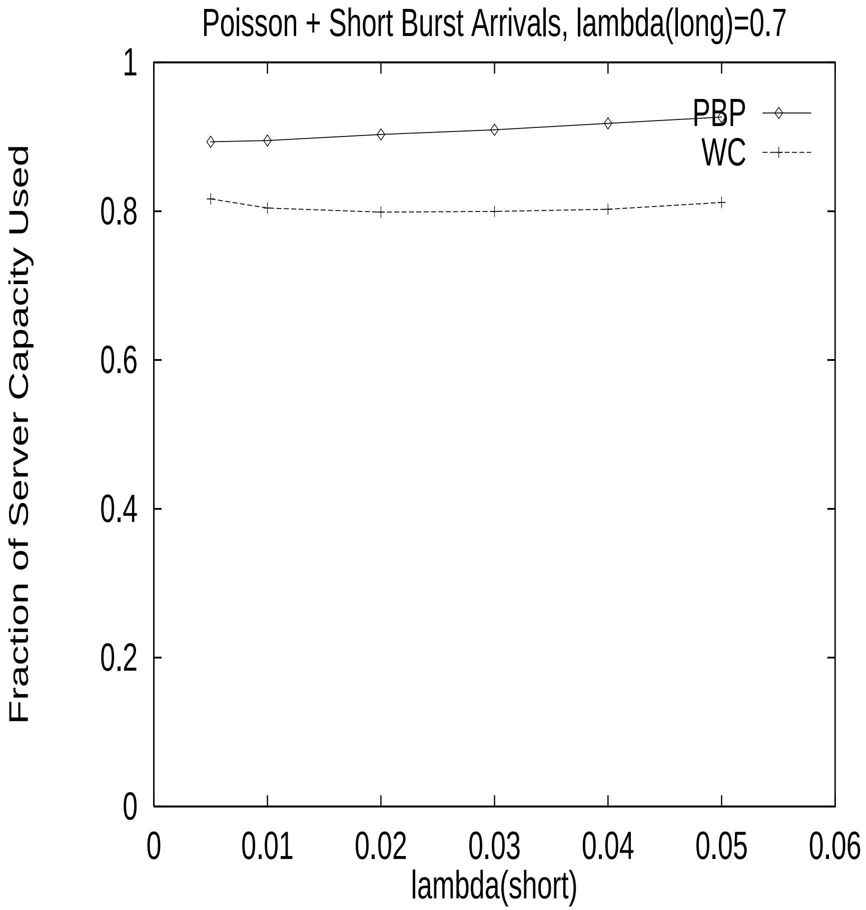

Conclusions and Future Work

- Competitive analysis for admission control in Continuous Media DBs:

  1. Conventional $\mathcal{WC}$ has competitive ratio linear in $\Delta = \frac{l_{max}}{l_{min}}$
  2. Lower bounds of $\Omega(\log \Delta)$ for any deterministic or randomized strategy
  3. Prepartitioning strategies – near-optimal for sufficiently large bandwidth
  4. Exploit knowledge of popularities for good average-case performance
  5. Experimental validation

- Incorporate other resources (e.g., memory)
- On-line load balancing in distributed continuous media servers

Competitive Ratio: Formal Definition

- $\overline{\sigma}$: request sequence
- $V_A(\overline{\sigma}) \sum_{S_A} l_i \cdot r_i$: total "benefit" of A over $\overline{\sigma}$

$$\kappa(A) = \begin{cases} \sup_{A^*,\overline{\sigma}} \frac{V_{A^*}(\overline{\sigma})}{V_A(\overline{\sigma})} & , \quad \text{if } A \text{ is deterministic} \\ \\ \sup_{A^*,\overline{\sigma}} \frac{V_{A^*}(\overline{\sigma})}{E_A[V_A(\overline{\sigma})]} & , \quad \text{if } A \text{ is randomized} \end{cases}$$

Experimental Results: More Details

- Poisson: arrival rate λ
- Bursty: short bursts of request batches arriving at displacements of *burst separation*
- Poisson+Bursts: poisson arrival of long requests (λ_{long}) with bursts of short requests (λ_{short})

System Parameter	**Value**
Server Bandwidth Capacity	100 channels / 100–250 Mbps
Request Lengths	5 minutes – 150 minutes
Request Rates (Variable Bandwidth Case)	500 Kbps - 8 Mbps
Zipfian Popularity Skew (z)	0.0 - 2.0

Related Work

- *On-line Algorithms:* On-line Interval Scheduling , On-line circuit routing
- *Shared Video-on-Demand* [Aggarwal et al., ESA'95]