On-Demand Media Distribution Services for the Masses

Ioanis Nikolaidis
yannis@cs.ualberta.ca
Computing Science Department
University of Alberta

Outline

- Motivation, Assumptions and Models
- Previous Work
- A Family of Greedy Schedules (GEBB)
- Loss-less VBR Broadcast (LLBE)
- Prefix Caching and Multicasting
- Conclusions

Motivation

- The initial motivation: Video-on-Demand.
  - Delivery process must ensure no buffer underflow (starvation) of the playout process.
- The problem.
  - maintain large selection and
  - virtually (near) on-Demand response.
    - Simple solution: staggered broadcast.
- Simple = Over-Engineered
  - Large bandwidth requirements for clients.
  - Expensive server equipment.

Assumptions

- Video encoding schemes produce Variable Bit Rate (VBR) data streams.
  - (Near-)Constant Bit Rate (CBR) still common.
- Storage capacity at the client set-top-box is essentially limitless.
- A large portion of the distribution cost is moved to the “edge” of the network.
- Bandwidth at the access point has increased dramatically (still, no FTTH).
- Disk I/O and de-compression are potential bottlenecks.

The Business Model

- Compete against commute times.
- Seasonal but predictable user traffic.
  - Time-of-the-day differences.
- Ephemeral but predictable selections.
  - Top-hits most frequently requested.
- Problems:
  - sufficient selection range,
  - per-video objectives,
  - pricing, etc.

The Broadcast Distribution Model

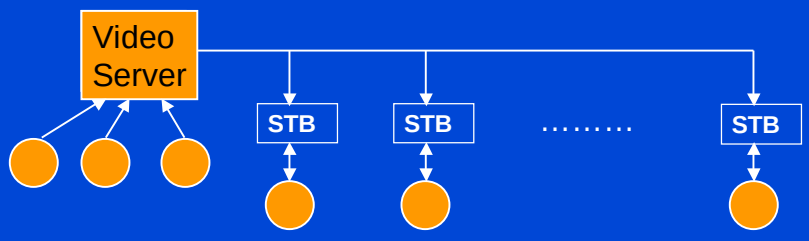

The Internet Distribution Model

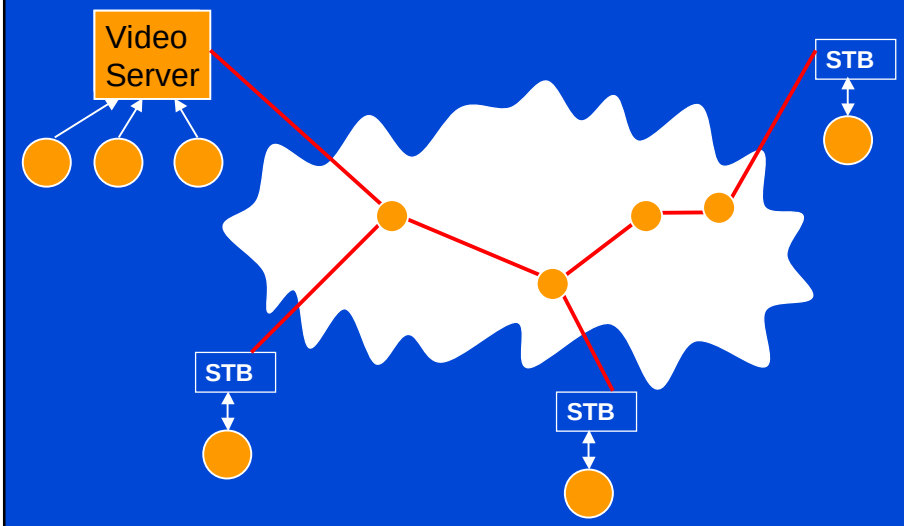

Efficiency

- Bandwidth
  - Server
  - Client
- Storage
  - Client

Playout Latency

- Interval between a client “tunes-in” until uninterrupted playout can begin.
  - random (but bounded), or,
  - deterministic.

Assumptions (cont.)

- (Either) Dedicated broadcast channel(s)
  - e.g., satellite or cable distribution.
- (Or) Multicast + RSVP support
  - to bound delay jitter in best effort (Internet) net.
- CBR Video Encoding
  - eventually relaxed to VBR.
- No VCR-like functionality while receiving.
  - You can always store the entire video.

What is a channel?

- A logical entity, an allocated fraction of a link's bandwidth.
- Implementation:
  - Time-Division Multiplexing,
  - Frequency-Division Multiplexing,
  - Both,
  - Weighted Fair Queueing, etc.
- (Small) bounded jitter needs to be absorbed.

The Traffic

- User requests for videos.
- Zipf distribution for requested items:
  - Typically 10-20 "hot"-set videos.
  - Hot videos can account for 80% of requests.
  - Rarely requested outliars.
- Objective:
  - Efficient distribution of most requested videos.
  - Use unicast or staggered multicast/broadcast for rarely requested videos.

Outline

- Motivation, Assumptions and Models
- Previous Work
- A Family of Greedy Schedules (GEBB)
- Loss-less VBR Broadcast (LLBE)
- Prefix Caching and Multicasting
- Conclusions

Demand-Based Batched Multicast

- Collect requests over successive non-overlapping time intervals.
- If one or more requests for item A, attempt to admit a replica of A.
- Admission may fail (rejection blocking).
- Inefficient for popular videos.

The Spectrum of Choices

- Continuous broadcast for hot-set.
- Hybrids for lukewarm-set.
- On-demand multicast for cold-set.

We will focus on the hot-set.

Pyramid Broadcasting (PB)

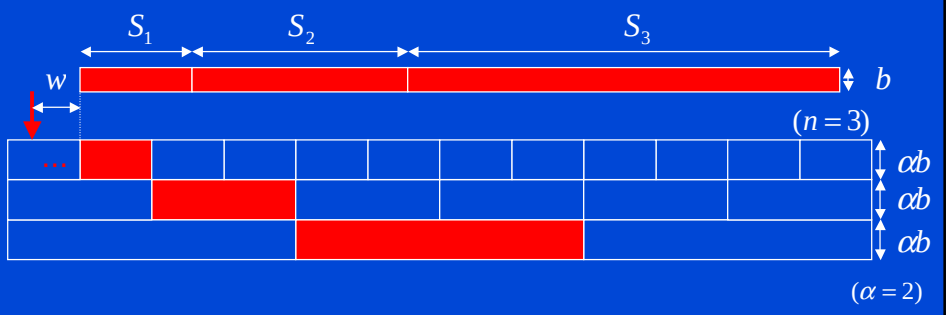

$$S_i = \begin{cases} \dfrac{S(\alpha-1)}{(\alpha^n-1)} & i=1 \\ \alpha S_{i-1} & i \geq 2 \end{cases}$$

$$w \leq S_1 / \alpha$$

$$\sum_{i=1}^{n} b_i = n\alpha b$$

Skyscraper Broadcasting (SB)

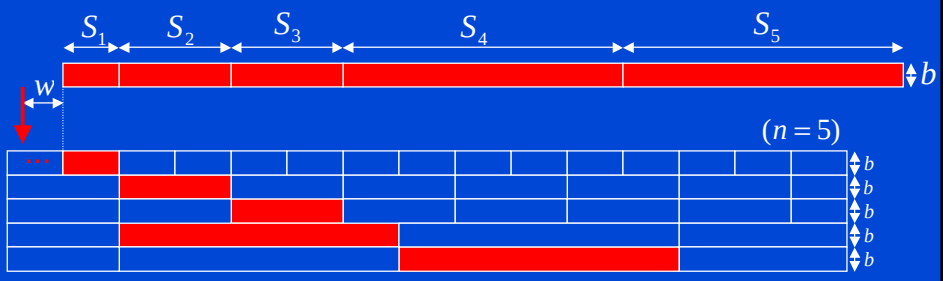

$$S_i = f(i)S_1$$

$$f(i) = \{1,2,2,5,5,12,12,25,25,\ldots\}$$

$$w \le S_1$$

$$\sum_{i=1}^{n} b_i = nb$$

Client-Centric Approach (CCA)

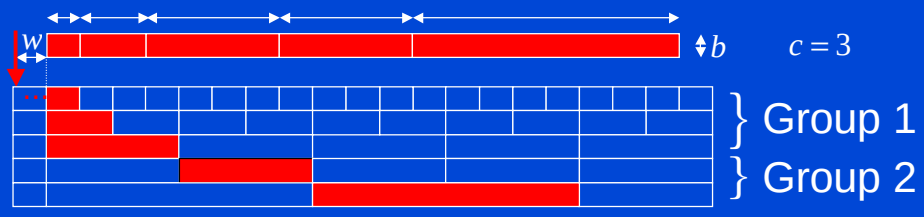

$$f(i) = 2^{i-\left\lceil \frac{i}{c} \right\rceil} = \begin{cases} 1 & i = 1 \\ 2f(i-1) & i \bmod (c+1) \neq 0 \\ f(i) & i \bmod (c+1) = 0 \end{cases}$$

$$w \le S_1$$

$$\sum_{i=1}^{n} b_i = nb$$

Harmonic Broadcasting

- The initial proposal was flawed.
- Fixed by Paris, Carter and Long.
- Cautious Harmonic Broadcasting.

Quasi-Harmonic Broadcasting

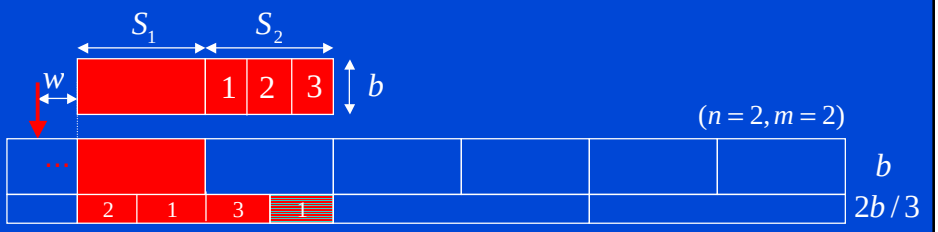

$$S_i = S/n \qquad w \le S_1$$

$$b_i = \frac{mb}{im-1} \qquad \sum_{i=1}^{n} b_i = bH(n) + O(m^{-1})$$

Relaxing the Timing Constraint

- Do not wait for beginning of a segment.
- Start buffering as soon as you tune-in.

A *Greedy* Client Download Strategy

Poly-Harmonic Broadcasting (PHB)

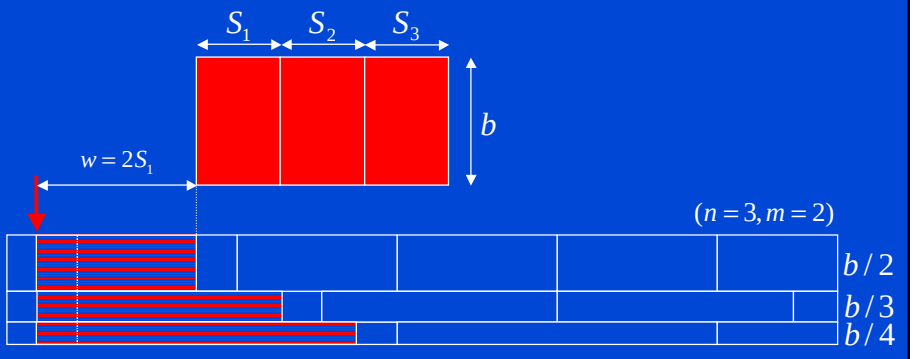

$$S_i = S/n \qquad w = mS_1$$

$$b_i = \frac{b}{m+i-1} \qquad \sum_{i=1}^{n} b_i = b\left(H(n+m-1) - H(m-1)\right)$$

Outline

- Motivation, Assumptions and Models
- Previous Work
- A Family of Greedy Schedules (GEBB)
- Loss-less VBR Broadcast (LLBE)
- Prefix Caching and Multicasting
- Conclusions

Motivating Question

- Poly-Harmonic Broadcasting (PHB) exhibits the best bandwidth efficiency of all VoD Broadcast Protocols.
  - Can we do better than PHB?
  - What is the lower bandwidth bound?

Problem: all known schemes are ad-hoc.

Three Common Components

- Per-channel/per-segment bandwidth.
- Duration of segments.
- Timing (continuity) constraints.

The Unnecessary Assumption

- Why start a segment reception only at the beginning of the segment?
- Segment position/timing markers:
  - for error recovery,
  - due to their nature (frame units),
  - for multiplexing.

Criterion

- From tune-in until its consumption, a segment must be transmitted exactly once in its entirety.

A Greedy Broadcast Schedule

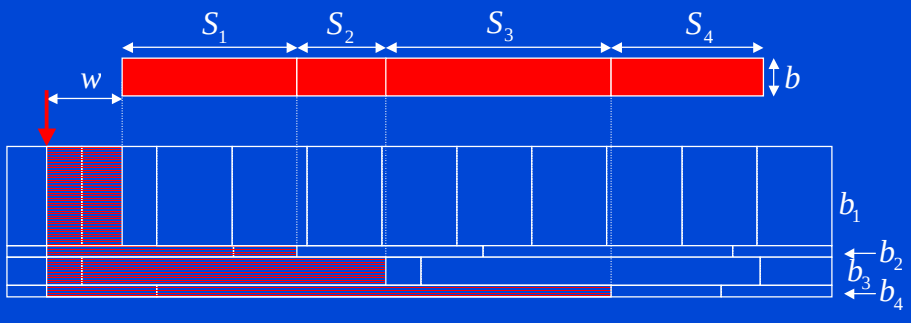

minimize $\sum_{i=1}^{n} b_i$ subject to: $b_i\left(w+\sum_{j=1}^{i-1} S_j\right) = bS_i$

$$\sum_{j=1}^{i-1} S_j = S$$

The Solution

$$\text{minimize} \sum_{i=1}^{n} b_i \equiv \text{minimize} \sum_{i=1}^{n} (1+b_i)$$

$$\text{s.t.} \quad \prod_{i=1}^{n} (1+b_i) = \frac{S}{w} + 1$$

Lagrange Multiplier Method

$$b_i = b^* = \left(\sqrt[n]{\frac{S}{w} + 1} - 1 \right) b$$

Greedy Equal-Bandwidth Broadcasting (GEBB)

$$b_i = b^* = \left(\sqrt[n]{\frac{S}{w} + 1} - 1 \right) b$$

$$S_i = w \left(\sqrt[n]{\frac{S}{w} + 1} - 1 \right) \left(\sqrt[n]{\frac{S}{w} + 1} \right)^{i-1}$$

GEBB Properties (Bandwidth)

$$\sum_{i=1}^{n} b_i = nb\left(\sqrt[n]{\frac{S}{w}+1}-1\right)$$

$$\lim_{n\to\infty}\sum_{i=1}^{n} b_i = b\ln\left(\frac{S}{w}+1\right)$$

GEBB Properties (Storage)

$$\left(\frac{S}{e}+\frac{w}{e}\right)b$$

GEBB vs. PB

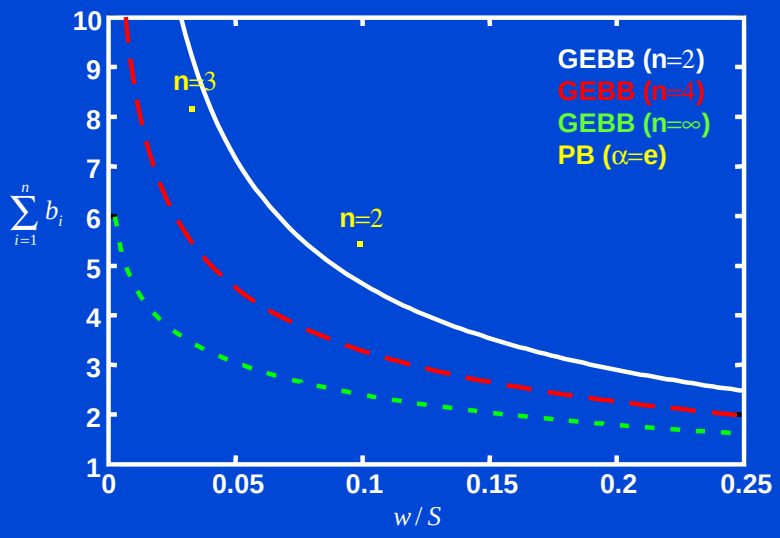

GEBB vs. SB

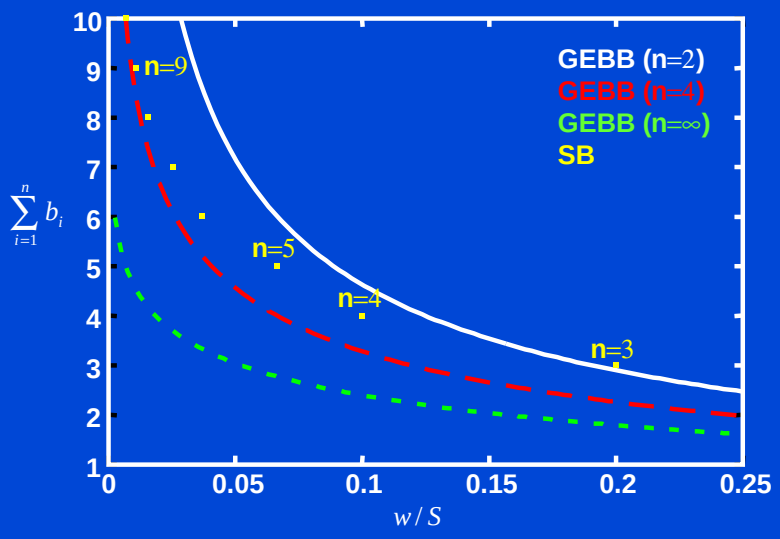

GEBB vs. QHB

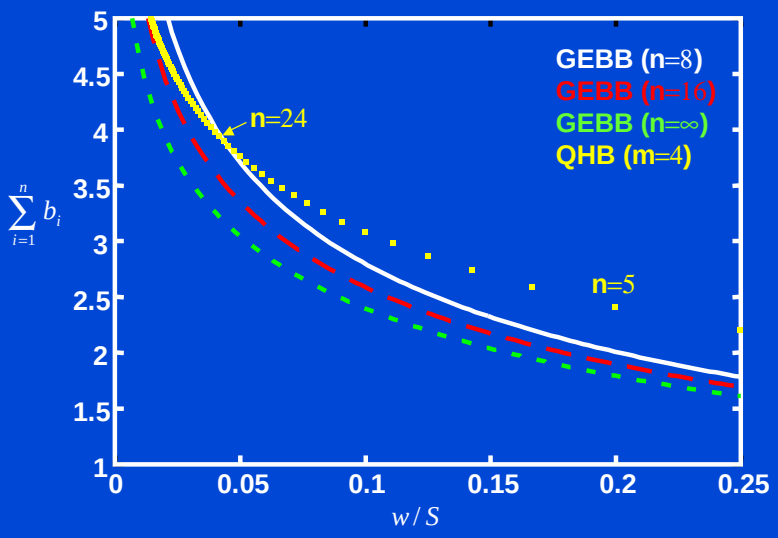

GEBB vs. PHB

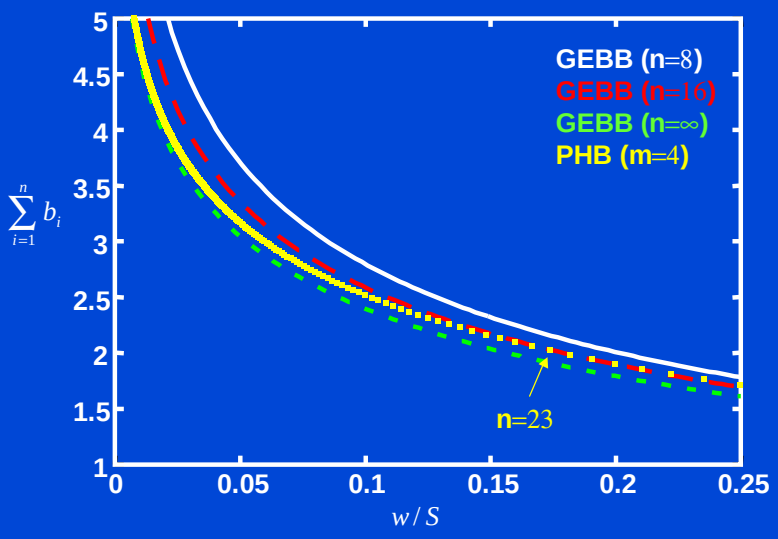

Outline

- Motivation, Assumptions and Models
- Previous Work
- A Family of Greedy Schedules (GEBB)
- Loss-less VBR Broadcast (LLBE)
- Prefix Caching and Multicasting
- Conclusions

From CBR to VBR

- Save (around 50%) in average bandwidth.
- Broadcast schemes defined for CBR.
- Two ways for CBR to VBR transition:
  - trivial: peak bandwidth allocation,
  - indirect:
    - apply CBR-based scheme, then
    - apply technique to minimize data loss.

VBR Support

- Pre-fetching(-sending) e.g. JSQ.
- Smoothing.
- Buffered multiplexing.
- Others?

Note: all of the above add a jitter component.

A Greedy Schedule for VBR

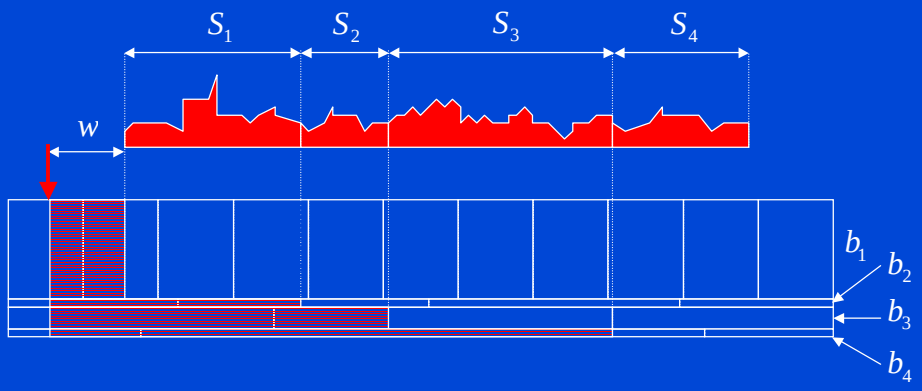

minimize $\sum_{i=1}^{n} b_i$ subject to: $b_i\left(w+\sum_{j=1}^{i-1} S_j\right)=\sum_{k=1+\sum_{j=1}^{i-1} S_j}^{\sum_{j=1}^{i} S_j} f_k$

The Solution

Can be transformed to an n-edge single source shortest path problem. Each edge corresponds to a (potential) channel.

$$n_i = \sum_{j=1}^{i} S_j$$

$$c_{n_{i-1},n_i} = \frac{\sum_{k=1+n_{i-1}}^{n_i} f_k}{(w+n_{i-1})} (= b_i)$$

$$n_0 = 0$$

$$n_K = N$$

$$c_{2,7} = \frac{\sum_{k=3}^{7} f_k}{(w+2)}$$

The Solution (cont.)

An n-step Bellman-Ford algorithm for single source shortest path.

$$Time : O(nN^2)$$

$$Space : O(N^2)$$

But,

$$N \approx 40000 - 160000$$

The Solution (cont.)

$$Space : O(nN)$$

No cost matrix, instead costs are calculated upon demand in constant time using an array of pre-calculated prefix sums.

$$v_i = \sum_{j=1}^{n_i} f_j$$

$$c_{n_{i-1},n_i} = \frac{v_i - v_{i-1}}{(w + n_{i-1})}$$

Bandwidth vs. Segments

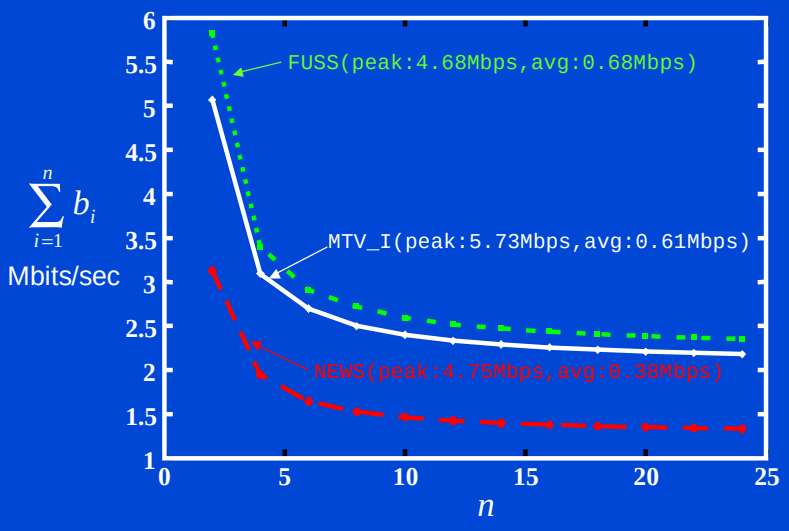

Bandwidth vs. Latency

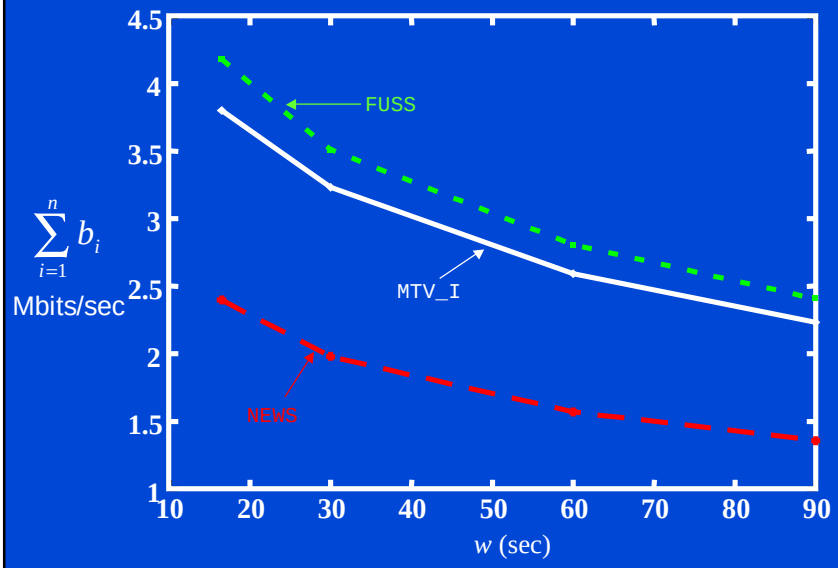

Per-Segment Bandwidth

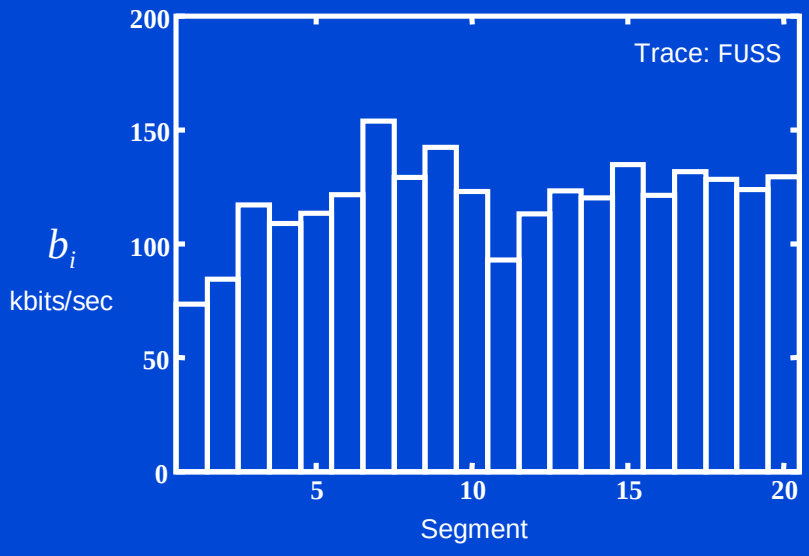

Per-Segment Size

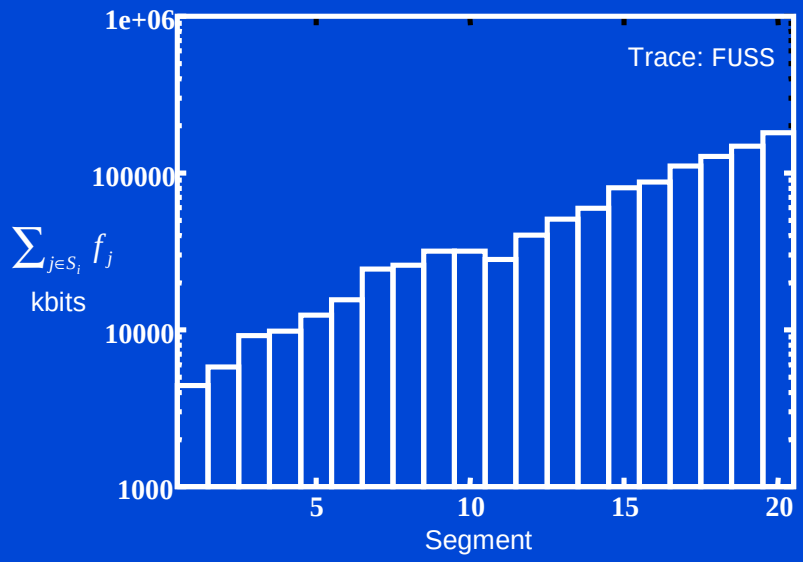

vs. Lossy Schemes

Scheme	Loss Prob.	Bandwidth
LLBE	0.000	33.587
VBR-B	0.153	86.958
TAF	0.104	60.722

$$n = 7, c = 7, N = 40000, F = 25\,fps, w \le 16.5\,\text{sec}.$$

Outline

- Motivation, Assumptions and Models
- Previous Work
- A Family of Greedy Schedules (GEBB)
- Loss-less VBR Broadcast (LLBE)
- Prefix Caching and Multicasting
- Conclusions

The Internet Distribution Model (2)

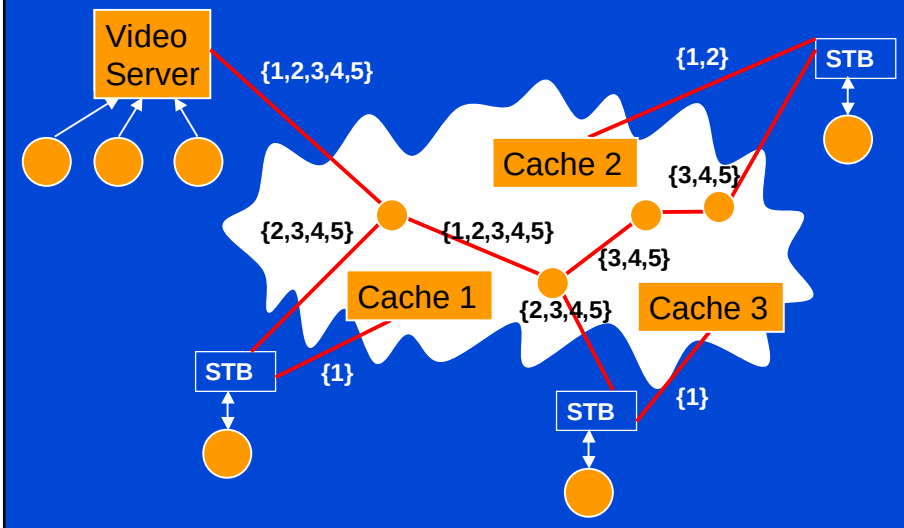

Prefix Caching

- A two-tier system:
  - Servers and Local Prefix Caches.
  - Servers:
    - Multicast continuously all the segments.
    - How far a segment multicast is forwarded depends on
      - (a) declared interest of receivers,
      - (b) load conditions of the network.
    - To match the available bandwidth on a link, forward only as many of the *latter* segments as possible.

Prefix Caching (cont.)

- *Local* Prefix Caches:
  - Provide only the first few segments.
  - Their purpose is to “buy time” for the latter segments to be retrieved via the multicast from the servers, using as few latter segment multicast flows as possible.

Prefix Caching (cont.)

- The tradeoffs:
  - The more the fraction of segments in local caches,
    - the higher the cost of redundantly maintaining (possibly unpopular) content,
    - the more likely the segments are delivered using unicast,
    - the more likely the cache causes congestion of the local part of the network.
  - The less the fraction of videos in local caches,
    - the less likely that all segments can traverse from server to client (due to the network load),
    - the more likely the servers congest the wide area part of the network. (Higher Blocking)

Cached Content vs. Multicast Bandwidth

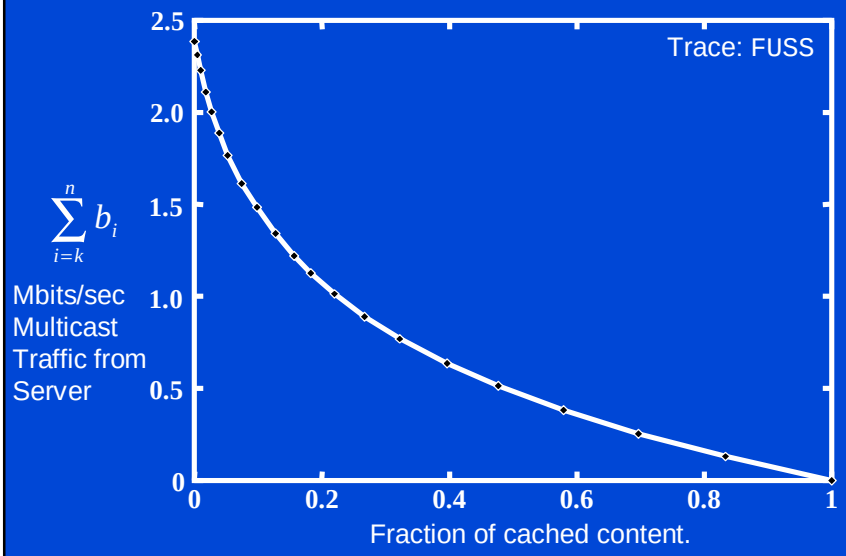

Outline

- Motivation, Assumptions and Models
- Previous Work
- A Family of Greedy Schedules (GEBB)
- Loss-less VBR Broadcast (LLBE)
- Prefix Caching and Multicasting
- Conclusions

Conclusions

- The Periodic Schedule Construction for given playout latency which minimizes the necessary server bandwidth is solved for both CBR and VBR video content.
- Open problems:
  - incorporate (per-)client bandwidth constraints,
  - incorporate FF/REW (interactive) operations,
  - determine the prefix cache location and content.

References

[1] Trace Adaptive Fragmentation for Periodic Broadcast of VBR Video
Li, Nikolaidis *NOSSDAV 1999*

[2] A Traffic Envelope and Transmission Schedule Computation Scheme for VoD Systems
Li, Nikolaidis *ISCC 1999*

[3] On the Design of Efficient Video-on-Demand Broadcast Schedules
Hu, Nikolaidis, Van Beek *MASCOTS 1999*

[4] An Inherently Loss-Less and Bandwidth Efficient Periodic Broadcast Scheme for VBR Video
Li, Nikolaidis *UofA, CS, TR99-04*

[5] An Inherently Loss-Less and Bandwidth Efficient Periodic Broadcast Scheme for VBR Video
Nikolaidis, Li, Hu *SIGMETRICS 2000*

[6] SCB: StairCase Broadcast for Media-on-Demand Systems
Li, Nikolaidis *MoMuC 2000*

[7] Towards Scalable VoD Distribution on the Internet
Nikolaidis *IPCCC 2001*