School of Computing Science
Simon Fraser University, Canada

Analysis of Peer-Assisted Video-on-Demand Systems with Scalable Video Streams

Mohamed Hefeeda

(Joint work with Kianoosh Mokhtarian)

22 February 2010

SFU

Motivations

- **Wide deployment of P2P streaming systems**
  - **PPLive, UUSee, SopCast, CoolStreaming, ….**
- **Video streams: 300 kbps to 1.5+ Mbps**
- **Typically more than upload bandwidth of most peers ➔ seed servers needed to help ➔ peer-assisted streaming systems**

Motivations

- Receivers are quite heterogeneous
  - network, processing power, screen resolution, …

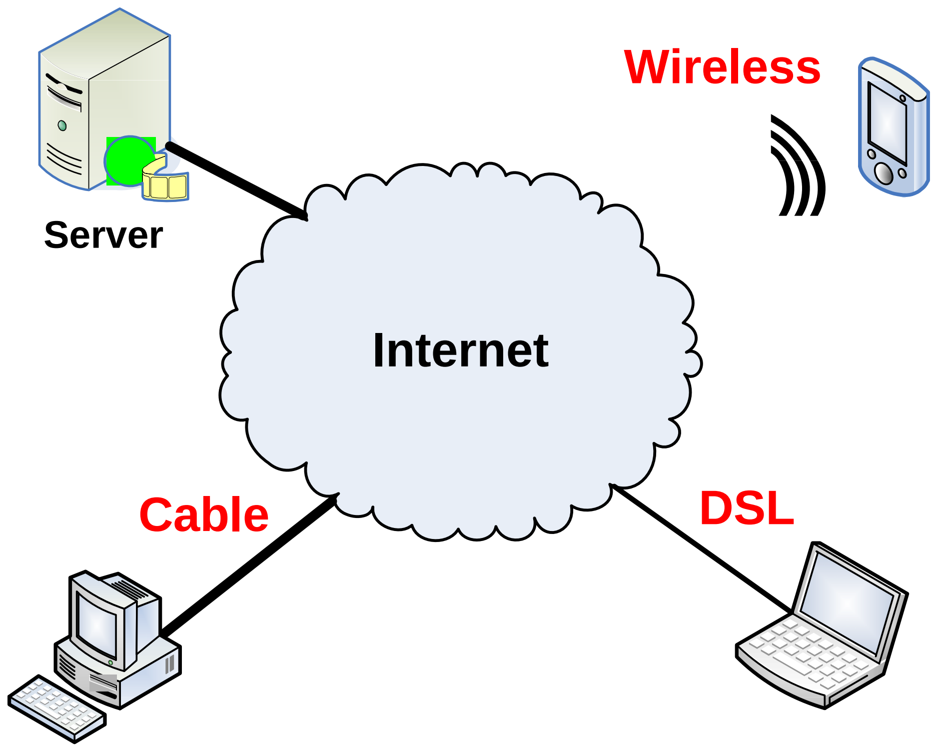

Supporting Receiver Heterogeneity

- **Encode low bitrate videos**
  - **Supports more receivers**
  - **Low video quality for everyone**
- **Distribute multiple versions (simulcasting)**
  - **Need to encode video several times**
  - **Switching among versions (wait for next I-frame)**
  - **Reduced connectivity (P2P): different stream versions ➔ different P2P networks**

Supporting Receiver Heterogeneity

- Multiple Description Coding (MDC)
  - Video quality ~ number of descriptions received
  - Considerable bitrate overhead
  - Computationally complex
- Scalable Video Coding (SVC)
  - Encode and distribute one video stream
  - Extract and decode various substreams
  - Lower overhead, simpler than MDC
  - Our focus in this paper

Our Contributions

- Analytical model for peer-assisted streaming systems with SVC
  - General model; details worked out for typical systems
  - Validated using simulations
- Provides insights and answers questions such as
  - What is the expected throughput or video quality?
  - How much seeding capacity is needed to increase quality to certain level?
  - How many peers can the system support in flash crowd scenarios?

Previous Works

- Capacity growth of P2P system over time [Tu 05]
  - Finding the time to turn off servers
  - Similar analysis based on fluid model in [Kumar 07]
  - Only nonscalable video streams
- Streaming capacity/startup delay and piece selection strategies [Zhou 07, Parvez 08]
  - Random matching of peers
- Analysis based on constructing specific overlay topology [Small 06, Liu 08]
  - Cannot support most of today's systems

System Model

- **Trackers**
  - Receive periodic update reports from peers
- **Peers**
  - Join by contacting one of the trackers
  - Try to serve lower layers first
    - As many layers as they can
  - Download in a streaming form, at the video bitrate
- **Seed servers**
  - Serve requests when not enough capacity in peers
  - Finite resources

System Model

- **Video file divided into S small segments**
  - Time unit = segment length = T_{seg} (e.g., 1 minute)
  - Peers in P_s (peers currently watching segment s) can serve to peers in $P_{1,\ldots,s-1}$
- **Serving requests by seed servers**
  - Random
  - Contribution based [Mokhtarian 09]

System Model

- **Inputs (characteristics of P2P system)**
  - Joint distribution of upload and download bandwidths (random variables D and U)
    - $Pr(x_1 \le \mathrm{D} \le x_2 , y_1 \le \mathrm{U} \le y_2)$
  - The expected peer seeding duration T_{seed}
  - Capacity of seed servers C, number of video layers L, bitrate of each layer $r_l(t)$ at each segment t
  - Distribution of peer arrivals (random variable N)
    - $Pr(\mathrm{N} = n)$: probability of n arrivals in one time unit (T_{seg})
  - Peer failure/departure rate $\gamma\ (t)$ at segment t

System Model

- **Outputs**
  - System throughput
  - Average video quality served
  - Max number of supported peers (flash crowd)

General Analytical Model: Overview

- **High-level procedure**
  - ***Step 1*: Estimate the capacity that peers in P_s can serve to others (i.e., $P_{1,\ldots,s-1}$) as a function of the capacity served to P_s**
  - ***Step 2*: Employ this function for $s = S, S-1, \ldots, 1$ to analyze the distribution of the complete video**

General Analytical Model: Step 1

- **Step 1: Distribution of one video segment**
  - $f(e, l, r_l, n)$: **expected value of output capacity for layer** l **of a segment when serving it with capacity** e **to** n **peers**
    - $e \geq r_l \times n$: **trivial → $f(.)$ in this case denoted by $f_{all}(.)$**
    - $e < r_l \times n$: **only $[e/r_l]$ of the n peers can be served**
      - **Selected randomly → $f_{rnd}(.)$**
      - **Selected based on contribution → $f_{opt}(.)$**
  - **Auxiliary rnd variable U_l: upload bw of peers for layer l**
    - **Distr of U_l estimated using U and peers upload behavior**

General Analytical Model: Step 1

- **Step 1: Distribution of one video segment**
  - $f(e, l, r_l, n)$: **expected value of the output capacity**

$$f_{all}(e, l, r_l, n) = n \times E[\ \mathbf{U}_l \mid \mathbf{D} \geq R_l\ ]$$

$$f_{rnd}(e, l, r_l, n) = k \times E[\mathbf{U}_l \mid \mathbf{D} \geq R_l]$$

  - k: **num peers that capacity e can serve with layer l**
  - R_l: **bitrate of the first l layers**
  - $E[\ \mathrm{U}_l \mid \mathbf{D} \geq R_l\ ]$: **expected upload bw (for layer l) of a peer given that it can download layers 1 to l**

$$E[\mathbf{U}_l \mid \mathbf{D} \geq R_l] = \frac{\int_0^\infty Pr(\mathbf{U}_l = x, \mathbf{D} \geq R_l)\ x\ dx}{Pr(\mathbf{D} \geq R_l)}$$

General Analytical Model: Step 1

- **Step 1: Distribution of one video segment**
  - $f_{opt}(e, l, r_l, n)$: **expected output capacity of the first k peers out of a set of n**
  - **Assume a sorted list of peers according to serving capacity for layer l. Let $k+1$-st peers have capacity x**
    - **The expected capacity of the first k peers:**

$$E[k \times \mathbf{U}_l \mid \mathbf{U}_l > x, \mathbf{D} \geq R_l]$$

$$\Rightarrow f_{opt}(e, l, r_l, n) = \int_0^{\infty} E[k \times \mathbf{U}_l \mid \mathbf{U}_l > x, \mathbf{D} \geq R_l] \times$$

$$\Pr(k \text{ peers have capacity} > x, \text{ one has } x, n - k \text{ have} < x)\ dx$$

General Analytical Model: Step 2

- **Step 2: distribution of video segments over time**
  - **Peers dynamics (at a given segment)**
    - **Using $f_{opt\,|\,rnd\,|\,all}(.)$ obtained so far, the expected output capacity of a set of peers at a segment is:**

$$\sum_{n=0}^{k} \Pr(n \text{ peers demand } l) \times f_{all}(e, l, r_l, n) \times \Pr(\text{a peer stays}) +$$

$$\sum_{n=k+1}^{\infty} \Pr(n \text{ peers demand } l) \times f_{opt|rnd}(e, l, r_l, n) \times \Pr(\text{a peer stays})$$

  - **Pr(**there are ***n*** peers demanding layer ***l*** at a given seg**): calc'ed using distr of N & D (arrivals and bandwidths)**
  - **Pr(**a peer at a given seg stays until next seg**): calculated using peer failure rate**

General Analytical Model: Step 2

- **Step 2: distribution of video segments over time**
  - We know how to calculate the expected output of peers at a segment (prev slide): $\hat{f}(e, l, r_l, \mathbf{N}, \epsilon)$
  - e: input capacitiy, l: layer, r_l: layer rate, $\mathbf{N}$: random variable denoting arrivals, ε: probability of departure
  - At time instance i, throughput of peers at segment t is:

$$x_{t,l}^{i} = \hat{f}\Big(\underbrace{x_{t,l}^{i-1} \times \frac{T + T_{seed}}{T} + C_{t,l}}_{\mathrm{I}} \; , \; l \; , \; r_l(t) \; , \; \mathbf{\Lambda}_t \; , \; \gamma(t) \Big)$$

  - Term I: input capacity to seg, $r_l(t)$: rate of layer l at seg t, Λ_t (rnd variable): num peers at t, $\gamma(t)$: departure rate at t

General Analytical Model: Step 2

- **Step 2: distribution of video segments over time**

$$x_{t,l}^{i} = \hat{f}\Big(\underbrace{x_{t,l}^{i-1} \times \frac{T + T_{seed}}{T} + C_{t,l}}_{\mathrm{I}} \; , \; l \; , \; r_l(t) \; , \; \mathbf{\Lambda}_t \; , \; \gamma(t) \Big)$$

- **$x^{i}_{t,l}$ converge over time; $x_{t,l}$: the converged value**
- **Expected num peers receiving layer l of segment t in the long term: $m_{t,l} = x_{t,l} / r_l(t)$**

→ **Steady-state throughput:** $\sum_{t=1}^{T} \sum_{l=1}^{L} m_{t,l} \times r_l(t)$

→ **Average video quality:** $\frac{1}{\sum_{t=1}^{T} \mathbf{\Lambda}_t} \sum_{t=1}^{T} \big(\mathbf{\Lambda}_t \sum_{l=1}^{L} m_{t,l} \times q_l(t) \big)$

- **Λ_t: because different num peers in different segs**

General Analytical Model: Step 2

- Step 2: distribution of video segments over time
  - Max peers that system can support in a flash crowd?
  - Recall: the expected num peers receiving layer l of segment t in the long term: $m_{t,l} = x_{t,l} / r_l(t)$
  - Max of number of supportable peer = number of minimal substreams (base layers) that can be served using the expected capacity of the system = [total capacity – capacity needed for minimally serving current peers] / bitrate of base layer of first few segs =

$$\frac{1}{r_1(1)}\left(C + \sum_{t=1}^{T}\sum_{l=1}^{L} m_{t,l} \times r_l(t) - \sum_{t=1}^{T} \Lambda_t \times r_1(t) \right)$$

Analysis of Example System: Procedure

- **Employing the general analytical model to analyze a sample P2P streaming system**
  - **Video streams: L layers, each at rate r kbps**
  - **Download bandwidths ~ $U(0, M)$**
  - **Upload bandwidth = α × download bandwidth**
  - **Peer arrivals: Poisson distribution with arrival rate λ**
  - **Peer departures: given probabilities γ_t at each seg t**

  → Auxiliary random variables, intermediary functions $f_{opt\,|\,rnd\,|\,all}(.)$, and finally, $m_{t,l}$ values calculated

Analysis of an Example System: Results

- **Distribution of one layer of one video segment**

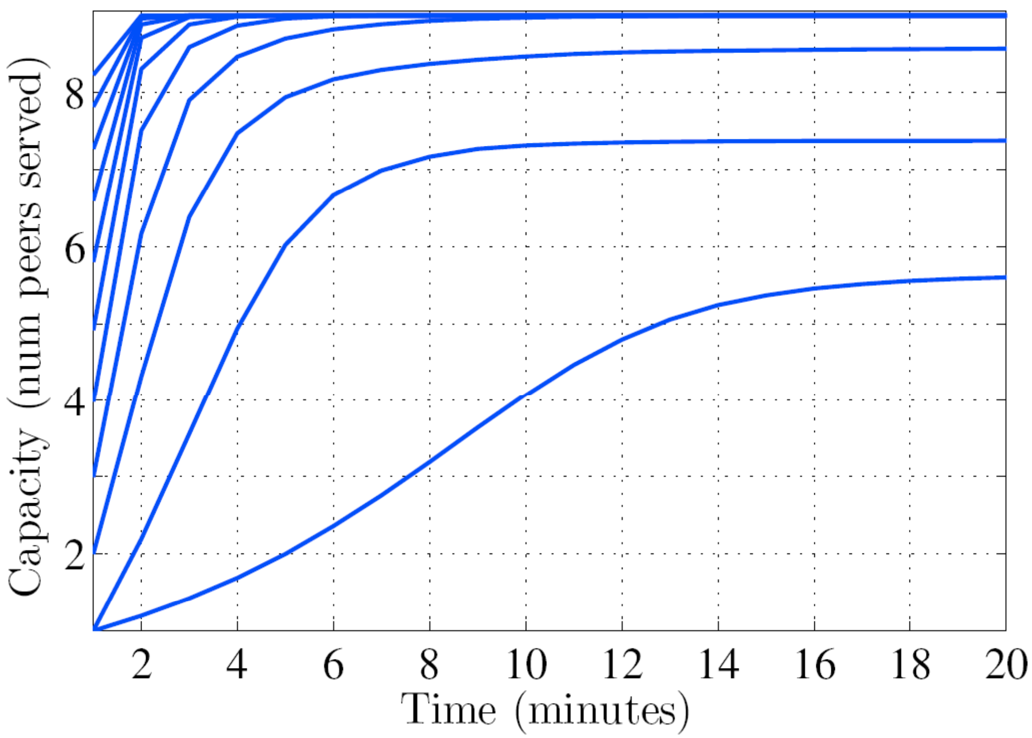

Num peers receiving layer 5. Lowest curve: seed server capacity of 1 layer rate.

- 1-hour video, 1-min segments, 10 layers at rate 200 kbps each, max download bw: 10 Mbps, upload bw: 20% of download, arrival rate: 10 peers per minute, 25% of peers eventually leave

Analysis of an Example System: Results

- **Distribution of the complete video**

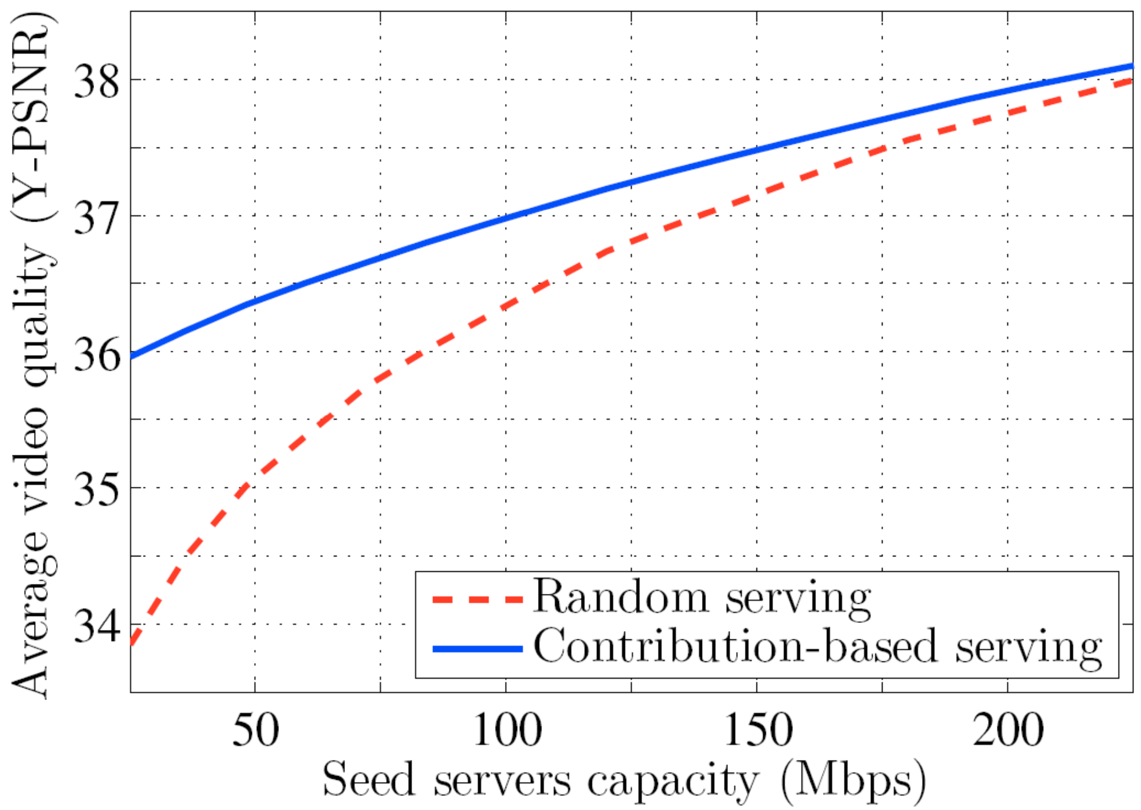

Average video quality delivered to peers for different seeding capacities

Validation of the Analysis

- **Simulation setup**
  - Peers and video characteristics
    - On average 500—600 peers in the network
  - Event-driven simulation, 10-second time step
  - Peers disobey some of our model assumptions and have other dynamics; to simulate realistic settings

Validation of the Analysis: Results

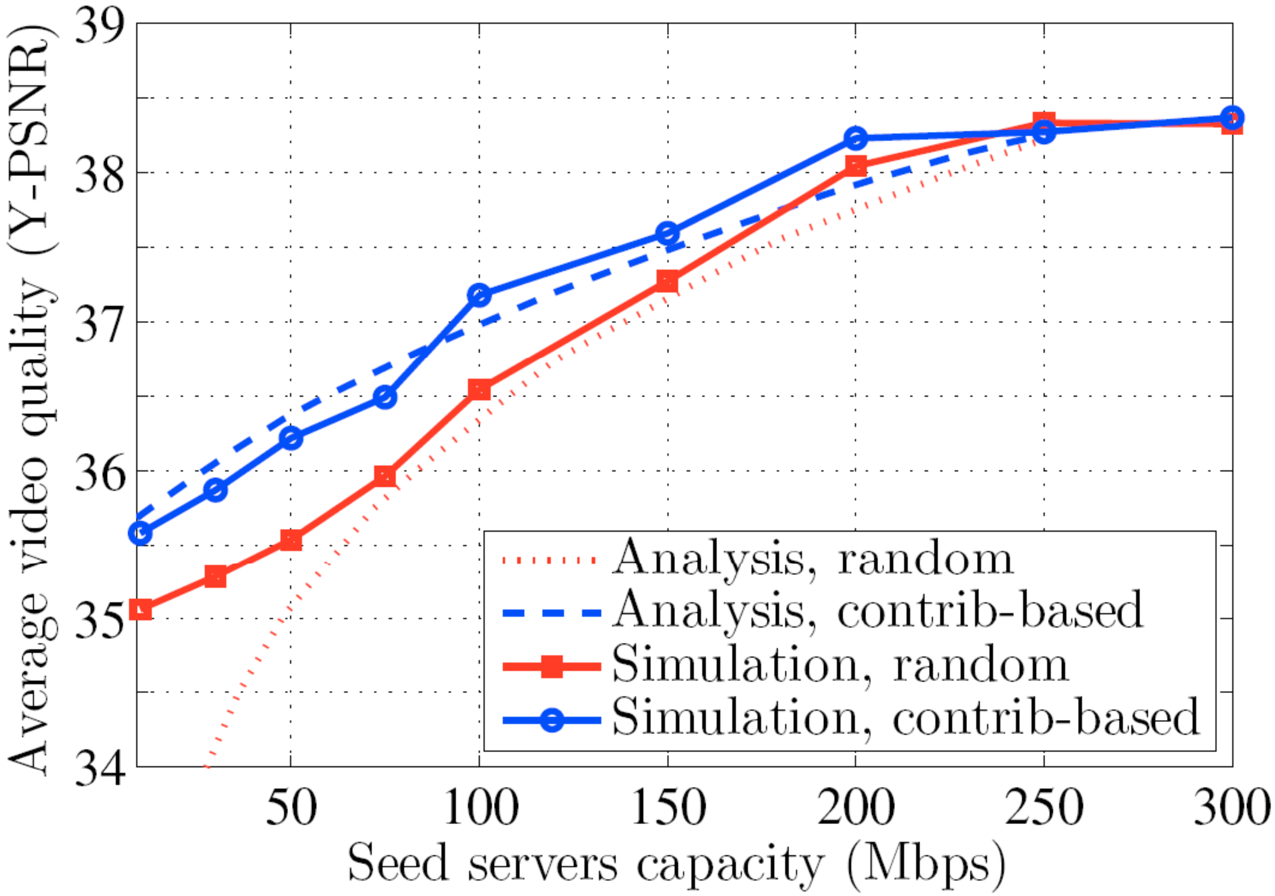

Theoretical results are close to simulation results

Analyzing Seed Server Allocation Methods

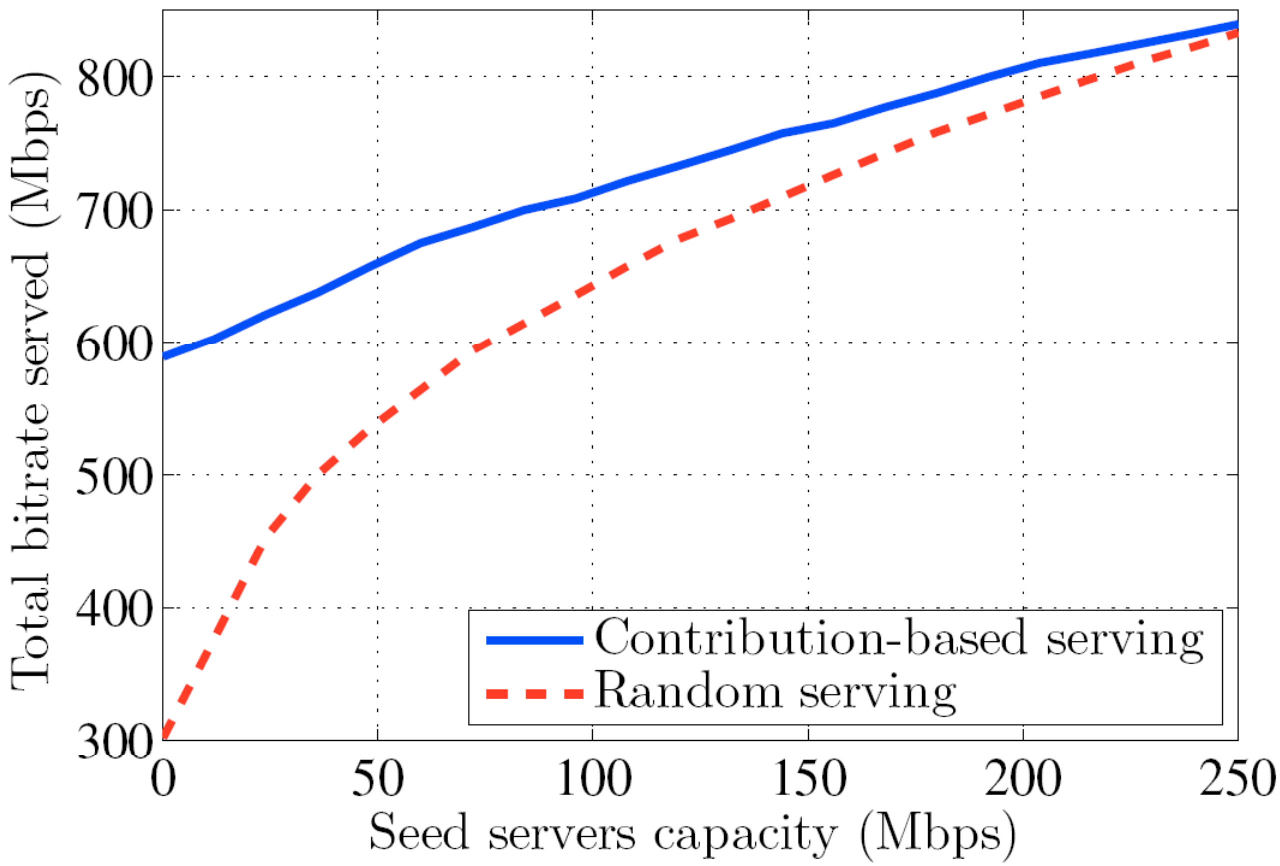

Expected throughput of the system in the long term

Analyzing Flash Crowd Scenarios

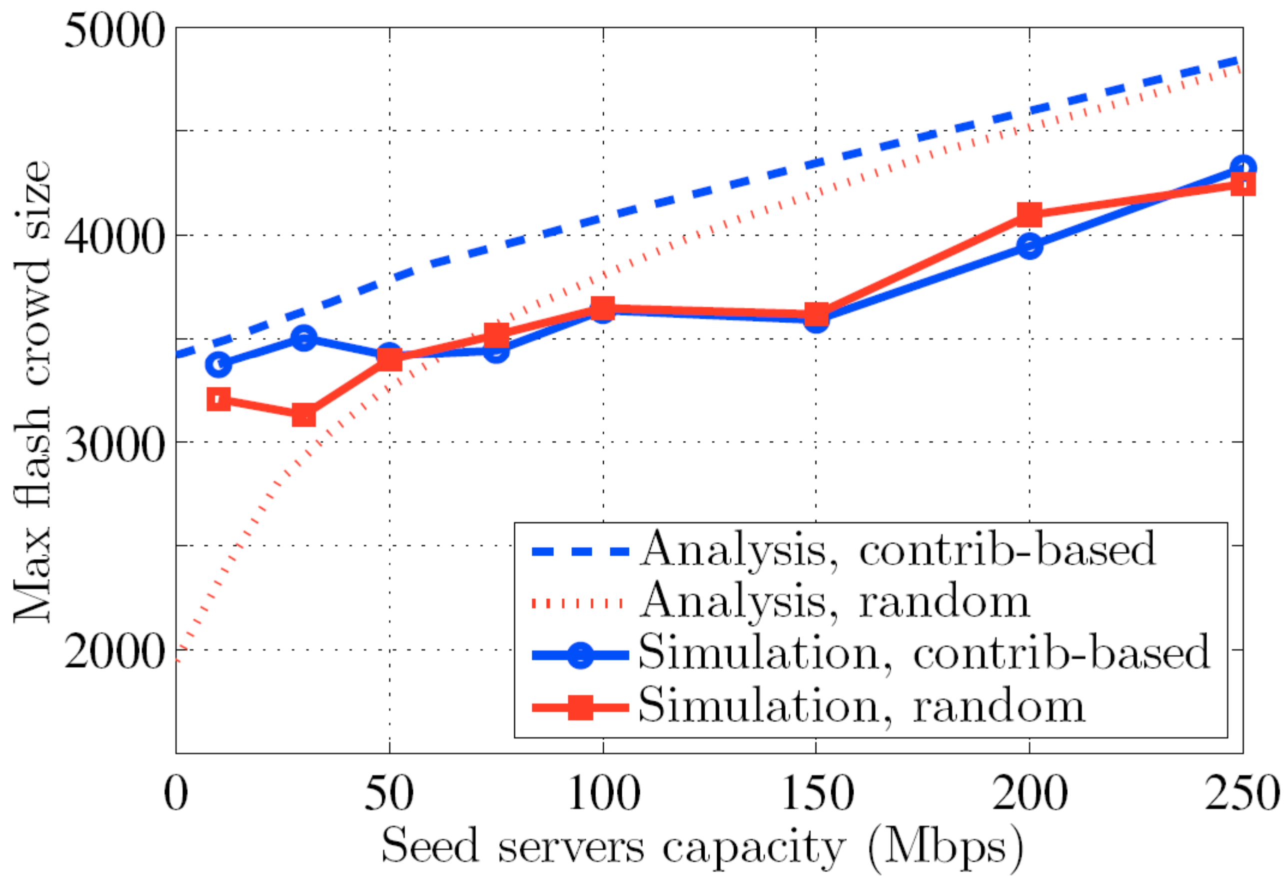

Max number of peers that the system can serve (with minimal substream) in a flash crowd

Conclusions

- **Theoretical analysis of P2P/peer-assisted streaming systems with scalable videos**
  - **The long-term throughput and delivered video quality**
    - **The cost-benefit tradeoff for provisioning a higher server capacity to have a desired throughput/quality**
  - **Max number of peers that the system can support in a flash crowd**
  - **Validated though simulations; results confirm the accuracy**

Thank You!

Questions??

- **More info at:**

  http://nsl.cs.sfu.ca/

References

- [Small 06]: T. Small, B. Liang, and B. Li. Scaling laws and tradeoffs in peer-to-peer live multimedia streaming. In *Proc. of ACM Multimedia Conference*, pages 539--548, Santa Barbara, CA, October 2006.
- [Liu 08]: S. Liu, R. Zhang-Shen, W. Jiang, J. Rexford, and M. Chiang. Performance bounds for peer-assisted live streaming. In *Proc. of ACM Conference on Measurement and Modeling of Computer Systems (SIGMETRICS'08)*, pages 313--324, Annapolis, MD, June 2008.
- [Tu 05]: Y. Tu, J. Sun, M. Hefeeda, Y. Xia, and S. Prabhakar. An analytical study of peer-to-peer media streaming systems. *ACM Transactions on Multimedia Computing, Communications, and Applications*, 1(4):354--376, November 2005.
- [Kumar 07]: R. Kumar, L. Yong, and K. Ross. Stochastic fluid theory for P2P streaming systems. In *Proc. of IEEE INFOCOM'07*, pages 919--927, Anchorage, AK, May 2007.
- [Zhou 07]: Y. Zhou, D. Chiu, and J. Lui. A simple model for analyzing P2P streaming. In *Proc. of IEEE International Conference on Network Protocols (ICNP'07)*, pages 226--235, Beijing, China, October 2007.
- [Parvez 08]: N. Parvez, C. Williamson, A. Mahanti, and N. Carlsson. Analysis of Bittorrent-like protocols for on-demand stored media streaming. In *Proc. of ACM Conference on Measurement and Modeling of Computer Systems (SIGMETRICS'08)*, pages 301--312, Annapolis, MD, June 2008.