Overlay Monitoring and Repair in Swarm-based Peer-to-Peer Streaming

Nazanin Magharei, Reza Rejaie
University of Oregon
http://mirage.cs.uoregon.edu

Introduction

- Swarm-based (Mesh-based) P2P Streaming (SPS) of live content is very popular, e.g. IPTV
- In SPS, peers form a random overlay mesh over which they incorporate swarming content delivery
- Connectivity of the Overlay plays a key role on the performance of content delivery in SPS
  - Randomness in connectivity ensures diversity of content
- In practice, a group of peers may exhibit a stronger (biased) internal connectivity and form a cluster due to :
  - ISP Localization, regional bootstrap nodes, network-aware connections or flash crowd
- Does clustering/localization adversely affect the performance of SPS mechanisms?

In this paper

- Understanding the effect of overlay clustering (or localization) on the performance of SPS mechanisms for delivery of live content
- Propose a distributed Overlay Monitoring and Repair (OMR) mechanism that maintains proper connectivity of the overlay

Talk Outline

- SPS: Background
- Importance of Overlay Connectivity
- Effect of Overlay Clustering
- OMR Design
- Evaluation
- Conclusion

SPS: Background

- The SPS mechanism has two basic components:
  - **Overlay Construction:** Peers form a randomly connected mesh where each peer has several parents and children
    - Number of parents and children of each peer is proportional to its incoming and outgoing bandwidth
  - **Content Delivery:** Periodic reporting + pull requesting (swarming) through a packet scheduling scheme

Organized View

- To understand the effect of overlay connectivity on content delivery
- Group peers into levels based on their shortest distance from source
- Peers one hop away from source are level 1, two hops away are level 2, ..
- The number of levels is :

$$Max\_Depth <= Log_{\frac{OutDeg}{N}} Deg_{src}$$

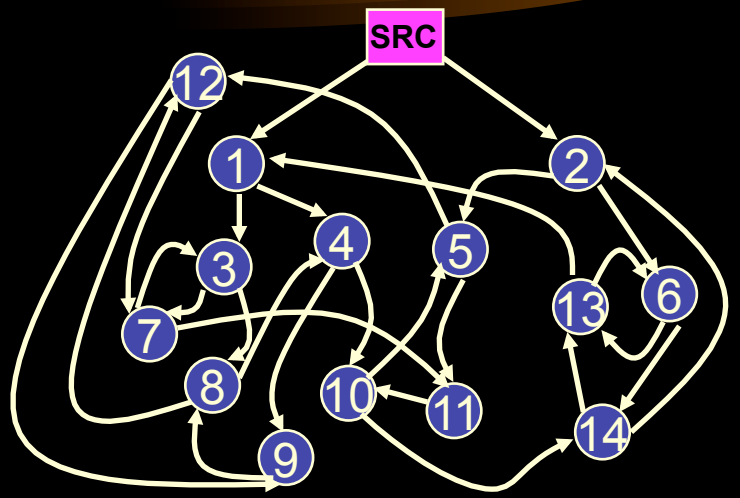

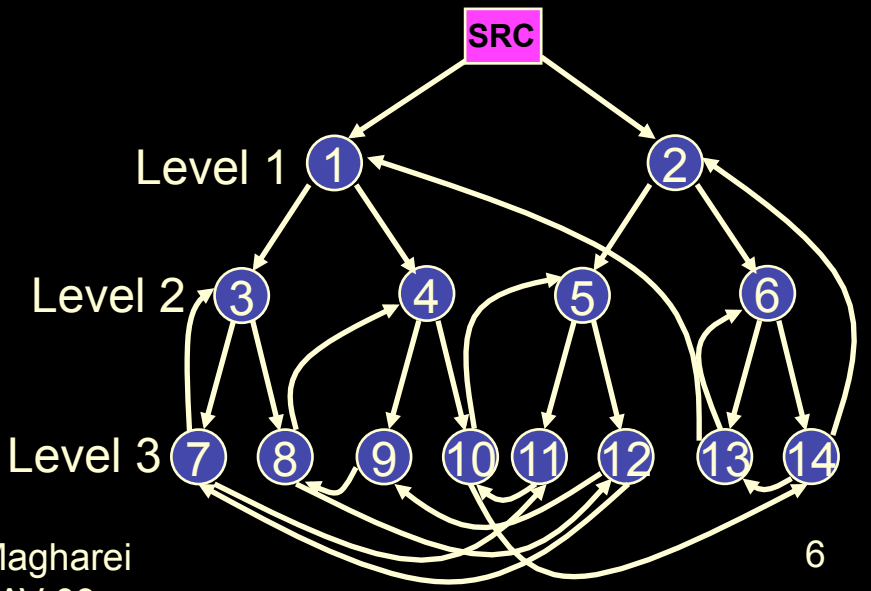

Organized View, cont.

- In a well performing packet scheduling scheme, delivery of packets has two phases:

1. *Diffusion phase*
   - Pulling packets from parents in the higher level
   - Collection of peers receive a packet during its diffusion phase form a sub-tree called diffusion sub-tree
   - Collection of packets delivered to a level 1 peer are called a substr
2. *Swarming phase*
   - Pulling packets from parents in the same or lower levels

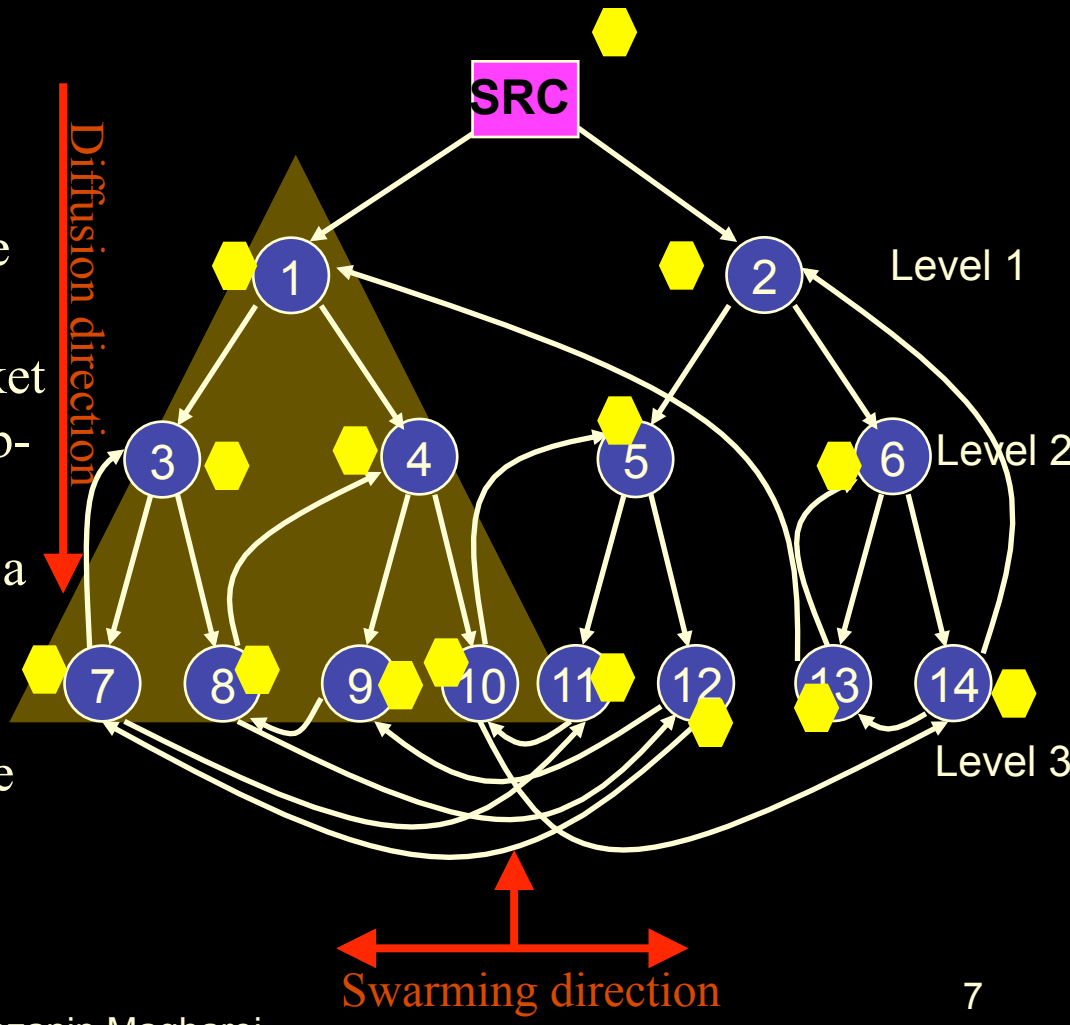

Impact of Overlay Connectivity

- Diffusion and swarming phases for delivery of each substr enable us to derive buffer requirement
- Maximum intervals for diffusion of each substr is *Max_Depth*
- The number of swarming intervals for a substr to a peer depends on the relative location of its swarming parents
- In a random overlay, swarming parents may not be located at the mutually exclusive diffusion sub-trees

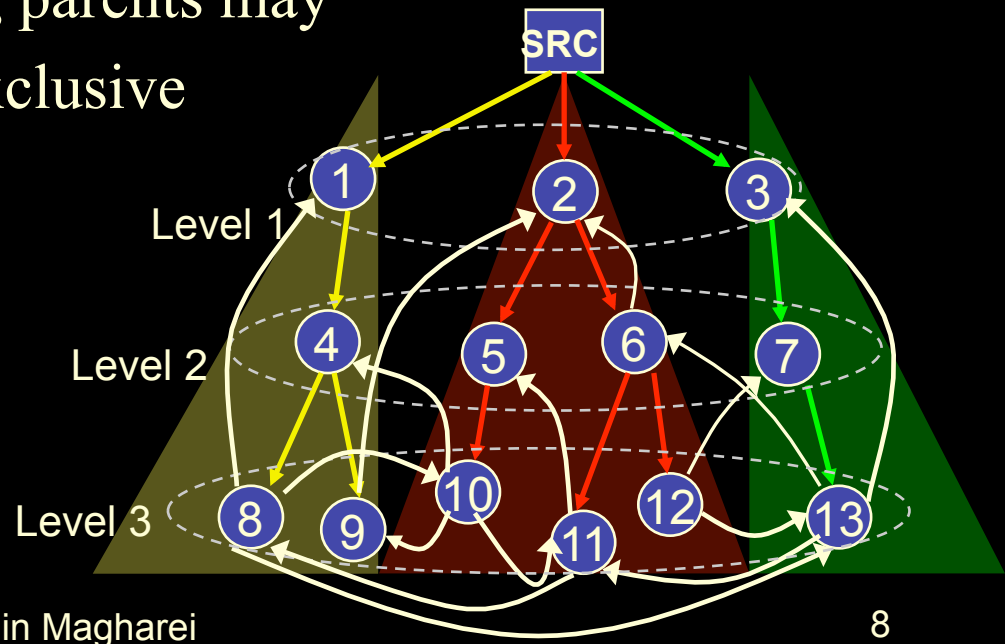

Impact of Overlay Connectivity, Cont.

- Swarming connections can be divided into 4 categories based on relative location of connected peers:
  - C_{ld}, C_{id}, C_{ls}, C_{is}
- The swarming connections between peers in different sub-trees (C_{ld} and C_{id}) are more effective for content delivery
  - They can provide a substr in one interval

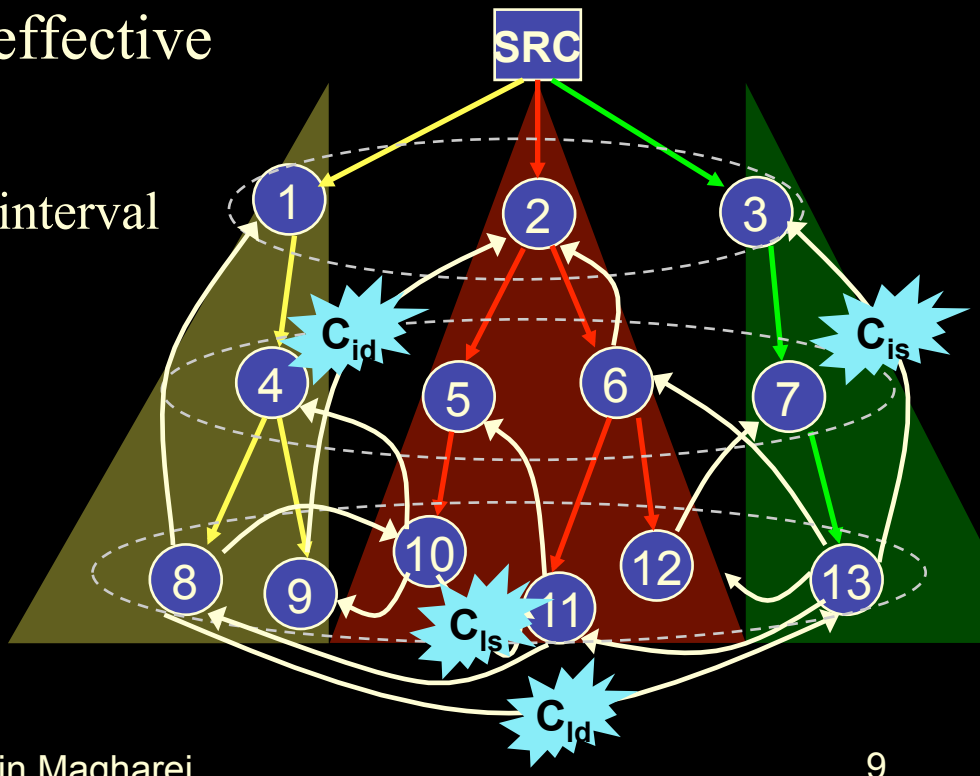

Effect of Overlay Clustering

- What if the connectivity is not random?
- A Group of peers may connect to each other and form clusters due to many reasons (e.g. ISP localization)
- Formation of clusters causes a poor connectivity between different clusters
  - Limits the flow of content among them
  - ➢ Adversely affect the performance of content delivery

Effect of Overlay Clustering, cont.

- If all peers in a cluster become part of a diffusion sub-tree, then, swarming connections
  - within the sub-tree (C_{ls} or C_{is}) increases
  - to other sub-trees (C_{ld} or C_{id}) decreases
- Clustering that decreases the number of connections between sub-trees can
  - increase the number of required swarming intervals and/or
  - make some sub-trees unreachable

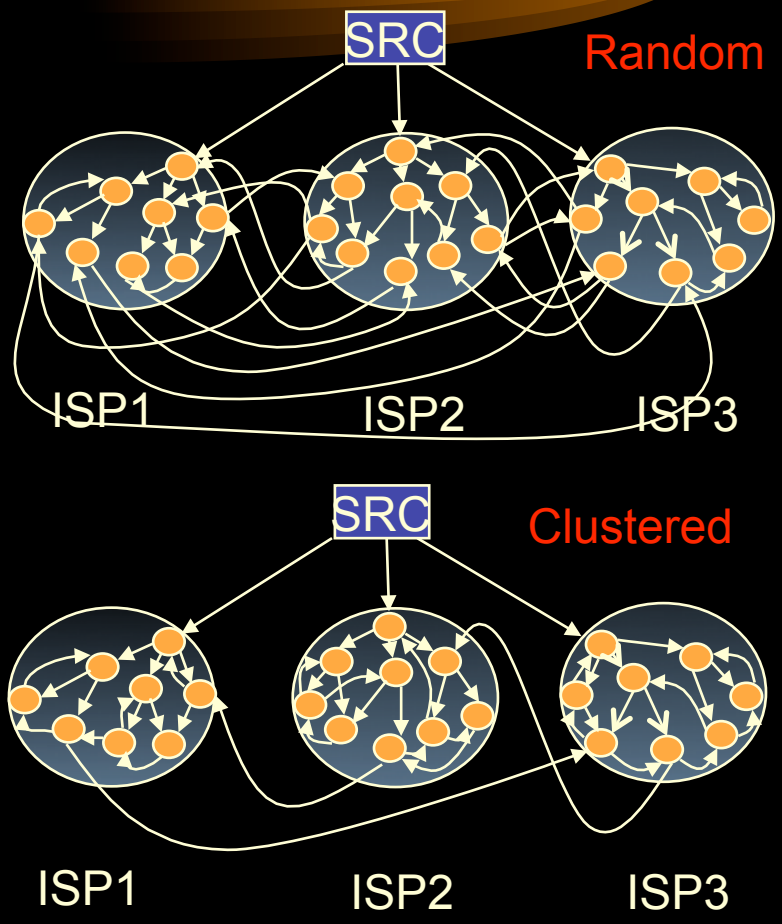

Effect of Overlay Clustering, cont.

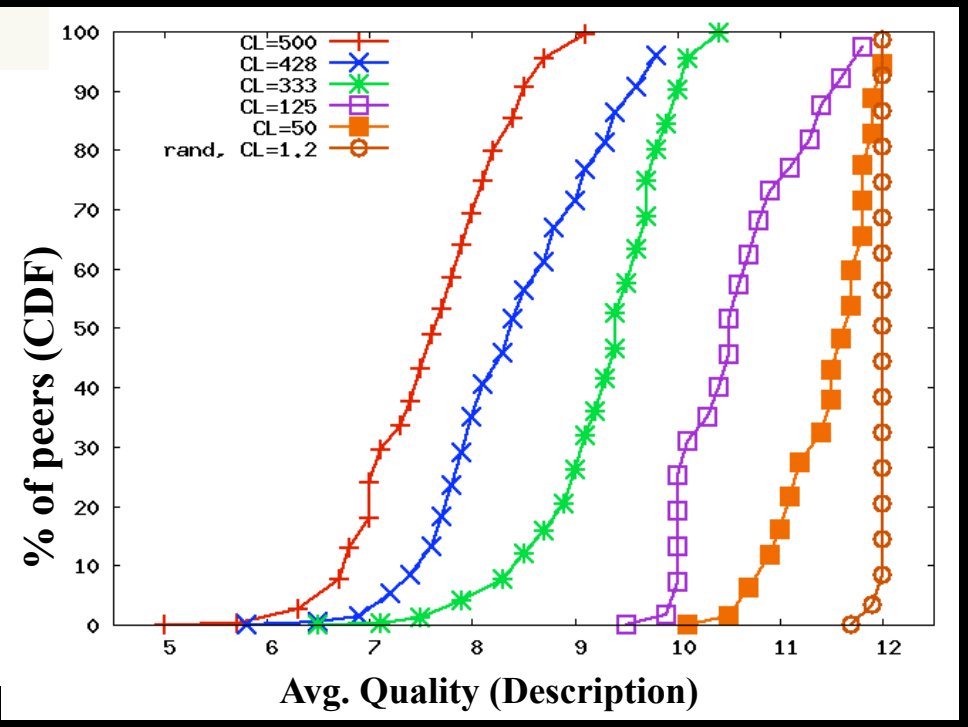

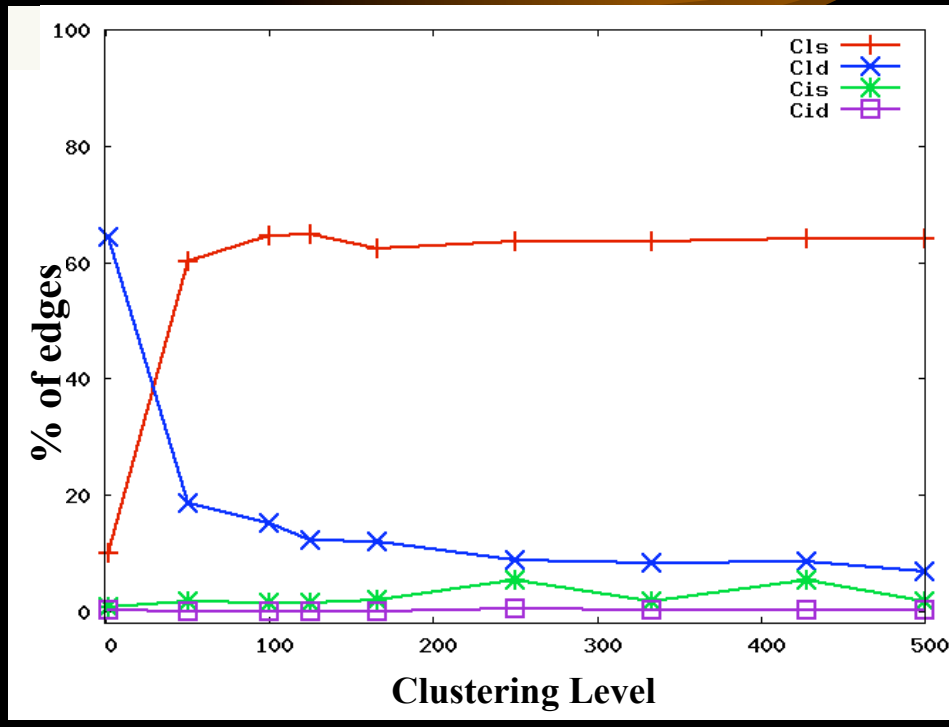

- 5000 homogeneous peers, # of clusters (ISPS) = 10, Content is MDC encoded with 12 descriptions

$$CL = \frac{\sum_{i=0}^{i=n} InDeg_i}{Ext\_incom}$$

Overlay Monitoring & Repair (OMR)

- How can we detect the effect of overlay clustering on the content delivery?
- How can such an overlay be minimally rewired in a scalable fashion to significantly improve performance of content delivery?
- OMR: A QoS mechanism that maintains the connectivity of the overlay such that SPS operates properly
- Two components
  - Detection: identifies any problem with the overlay connectivity
  - Reaction & Repair: rewires a minimum number of connections by reacting peers

Detection

- How to detect the effect of overlay clustering?
- Intuition: Poor content availability is a good indicator of poor connectivity when a good packet scheduling is used
  - Each peer *P* monitors its received substrs and the available substrs among its swarming parents
  - When peer *P* is missing at least one substr for a given period of time, it periodically invokes the reaction and subsequently repair algorithm

Detection, cont.

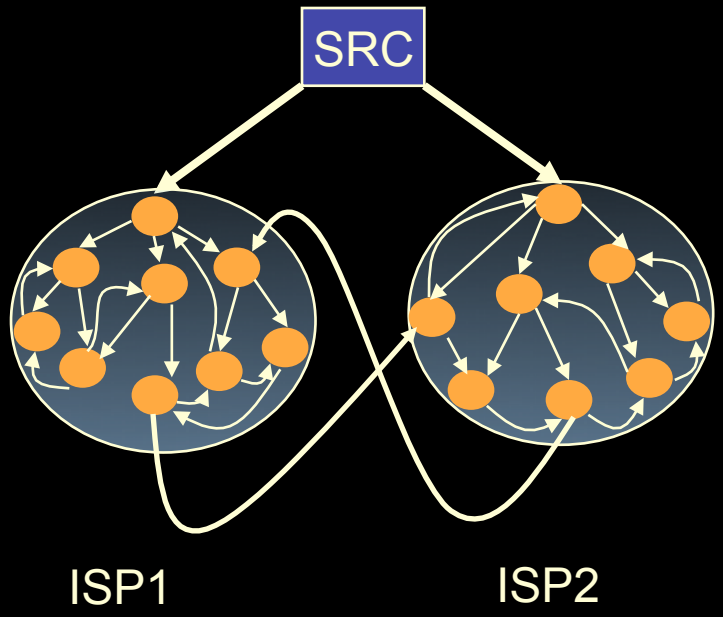

Reaction

- How many peers and which ones should react?
- Intuition: Probabilistically increasing the effective swarming connections improves the performance in a scalable fashion
  - How many? Probabilistic reaction controls the number of reacting peers
  - Who? Biased towards peers able to achieve the most improvement due to their location in the overlay

Reaction, cont.

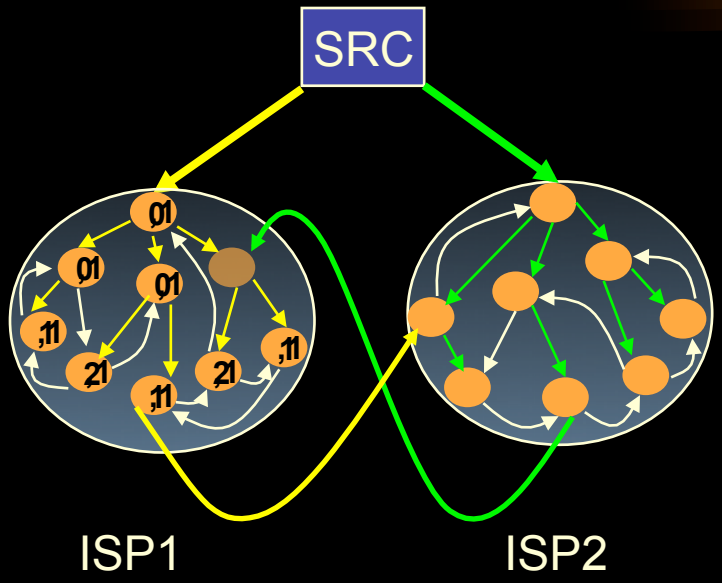

- Reaction probability for a missing substr *r* depends on
  - Contribution Factor (CF): % of swarming children missing *r*
  - Parent Factor (PF): % of swarming parents in the same sub-tree, missing *r*
  - $CF^{\alpha} \times \frac{1}{PF^{\beta}}$

Repair

- Where is the repair position?
- Intuition: The higher levels in the missing sub-tree leads to a smaller required buffer size
  - Peers with a higher contribution are placed at a higher position in the missing sub-tree *r*
  - The probability for new position (*nd*) is computed by:

$$1-(1-CF)^{(nd-1)}$$

Repair, cont.

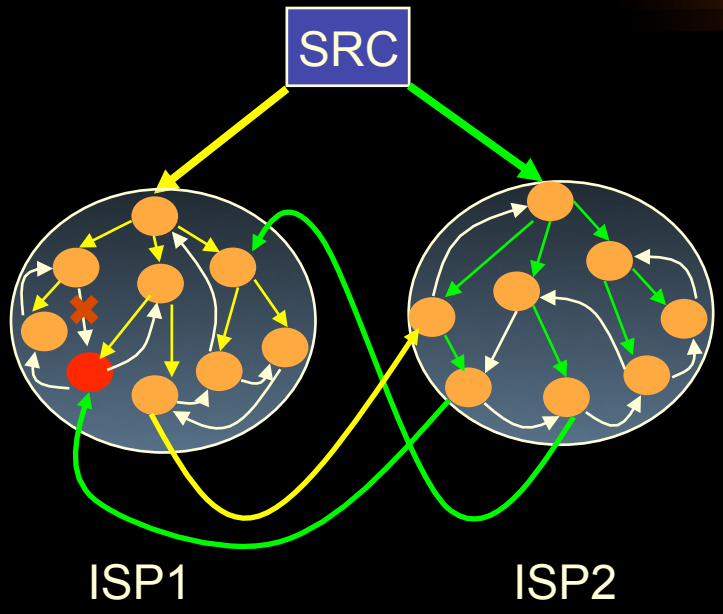

- Reacting peer traverses sub-tree *r* in the random depth-first search fashion till reaches level *nd* (e.g. 3)
- Reacting peer swaps one of its parents (preferably swarm) with the new parent

Performance Evaluation

- Preliminary evaluation of OMR using event-driven simulator
- The simulations are performed in multiple rounds, each round includes two steps:
  - P2P content delivery over the existing overlay
  - Running the OMR mechanism at each peer
- Resource constraint setting:
  - Source BW = Peer BW= Stream Rate
- Homogeneous and symmetric peer BW
- Peer population is 5000
- Number of clusters (ISPs) = 10
- The initial "localized overlay" has the highest level of clustering
  - External incoming BW to each cluster = Stream BW
- $\alpha=1$, $\beta=2$

How Does OMR Perform?

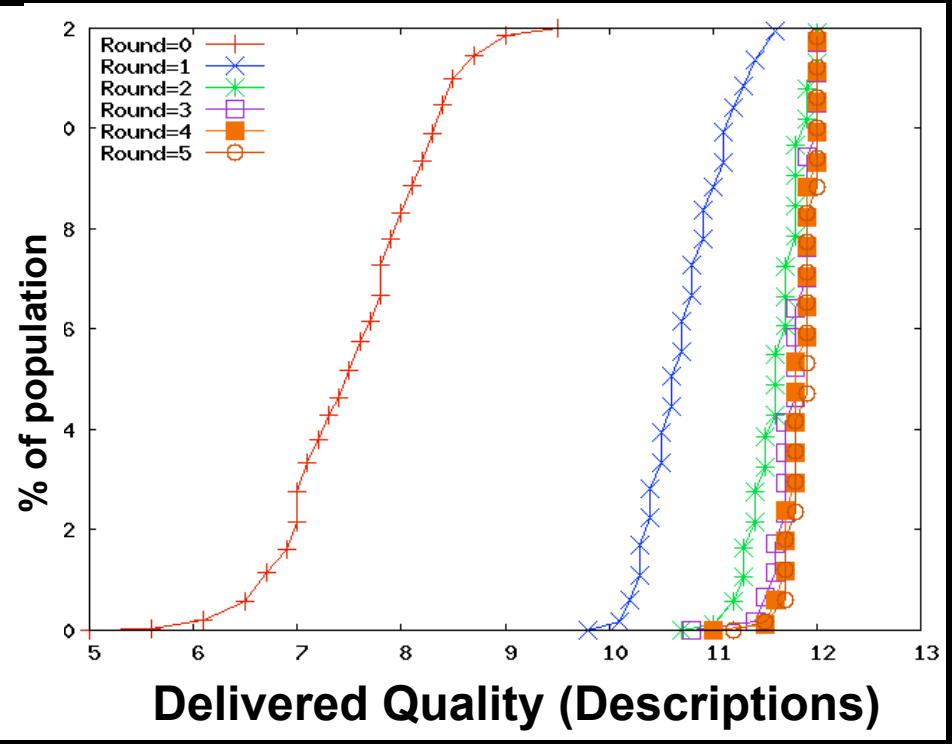

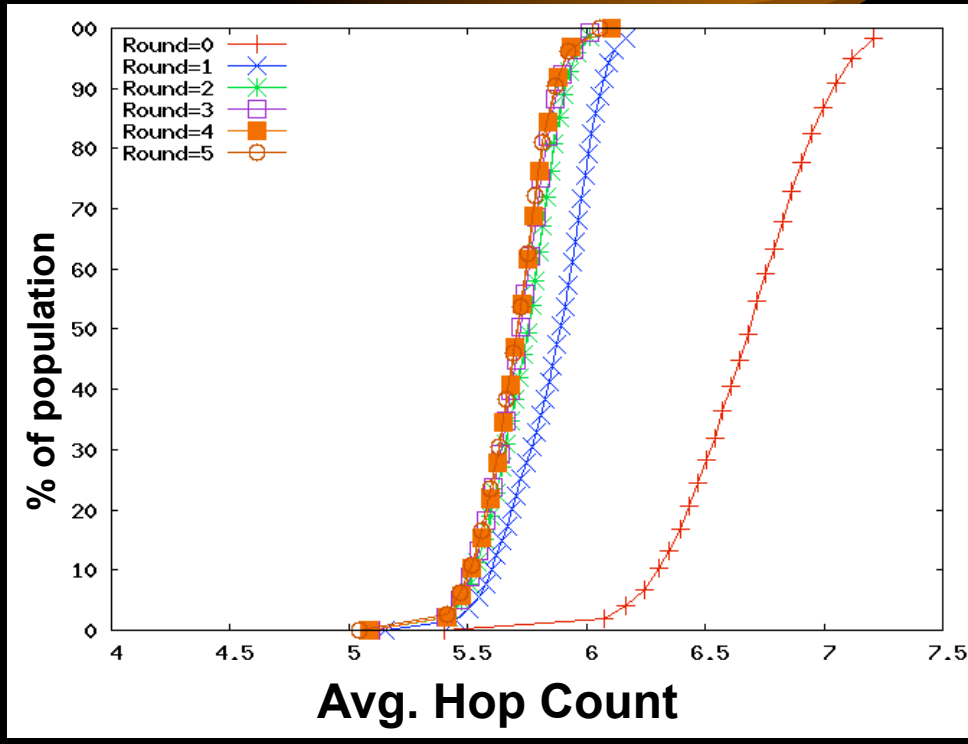

- *Significant improvement in delivered quality mostly during the first two rounds*
- Tighten the overlay effectively and reducing the buffer size

How Does OMR Perform? Cont.

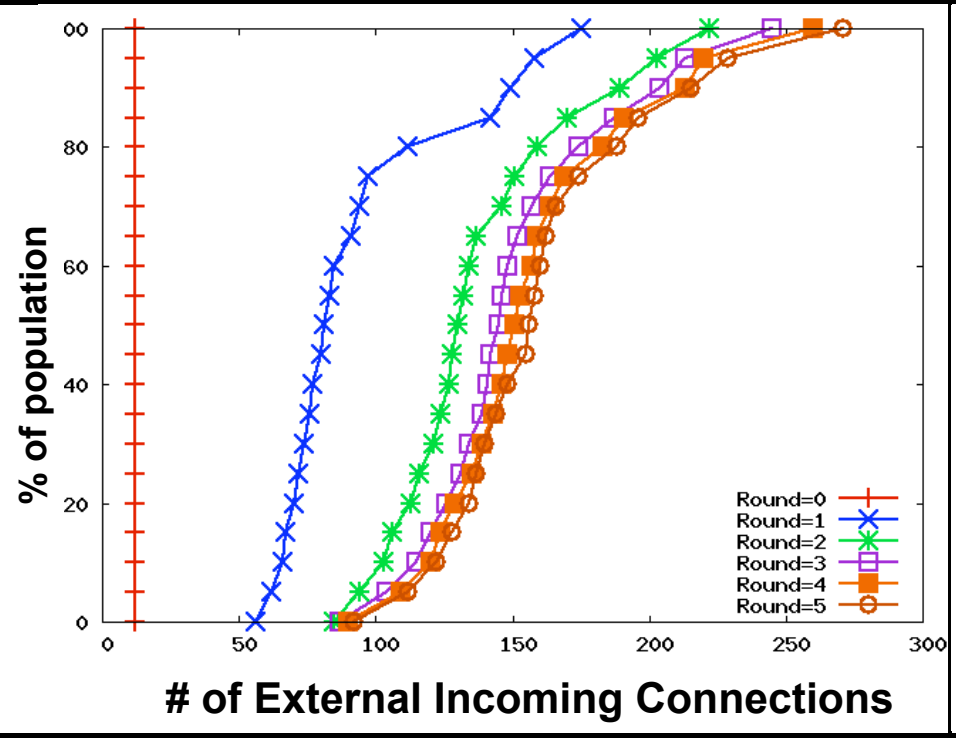

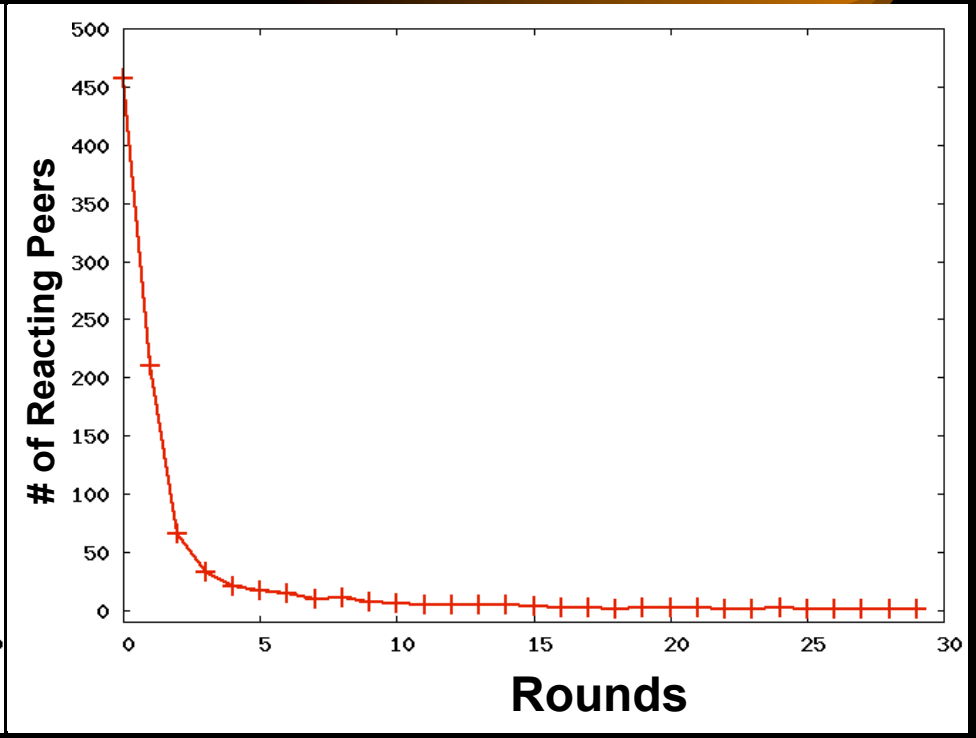

- *Connectivity between clusters significantly increases*
- About 9% of peers react in the first round and this number greatly decreases in the subsequent rounds

Scalability

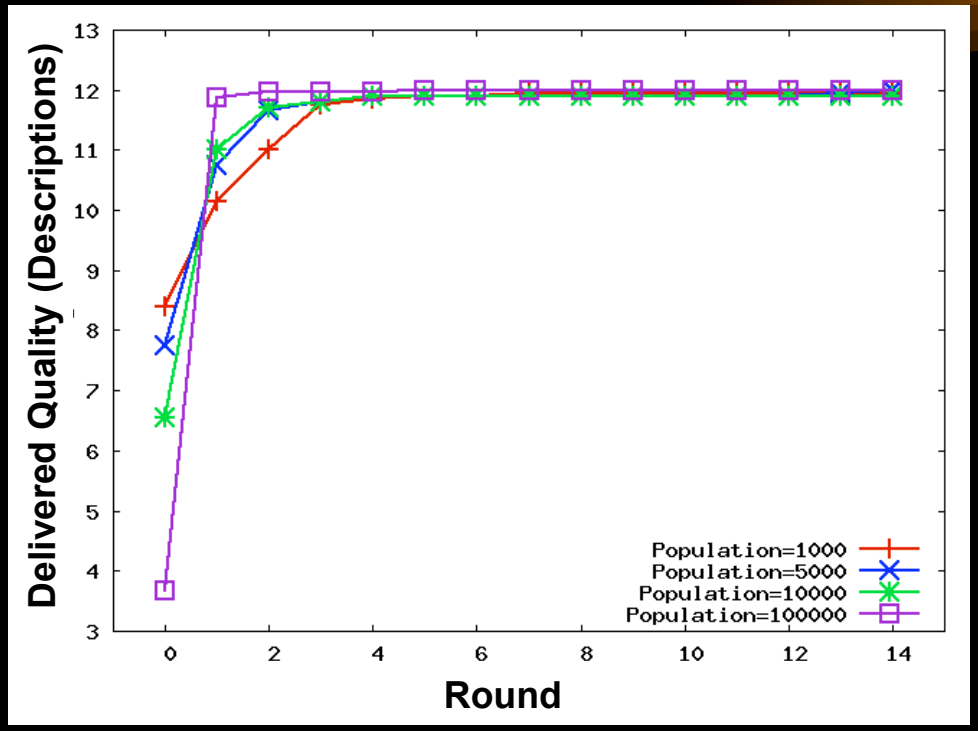

- Regardless of the population, quality improves during the initial rounds
- *Larger overlays initially have larger clustering level → lower quality*
- Time to reach maximum quality is lower for larger overlay as more peers initially reacts to the lower initial performance

BW Heterogeneity

- Content is MDC encoded
- High bandwidth peers can receive the full stream quality
- Low bandwidth peers can receive 50% of the quality
- Two scenarios of
  - SC1 80% are high bandwidth
  - SC2 20% are high bandwidth
- Low bandwidth peers require a smaller number of substrs and are less sensitive to unavailability of a particular substr.
  - ➢ We primarily focus on the delivered quality to high bandwidth peers

BW Heterogeneity, cont.

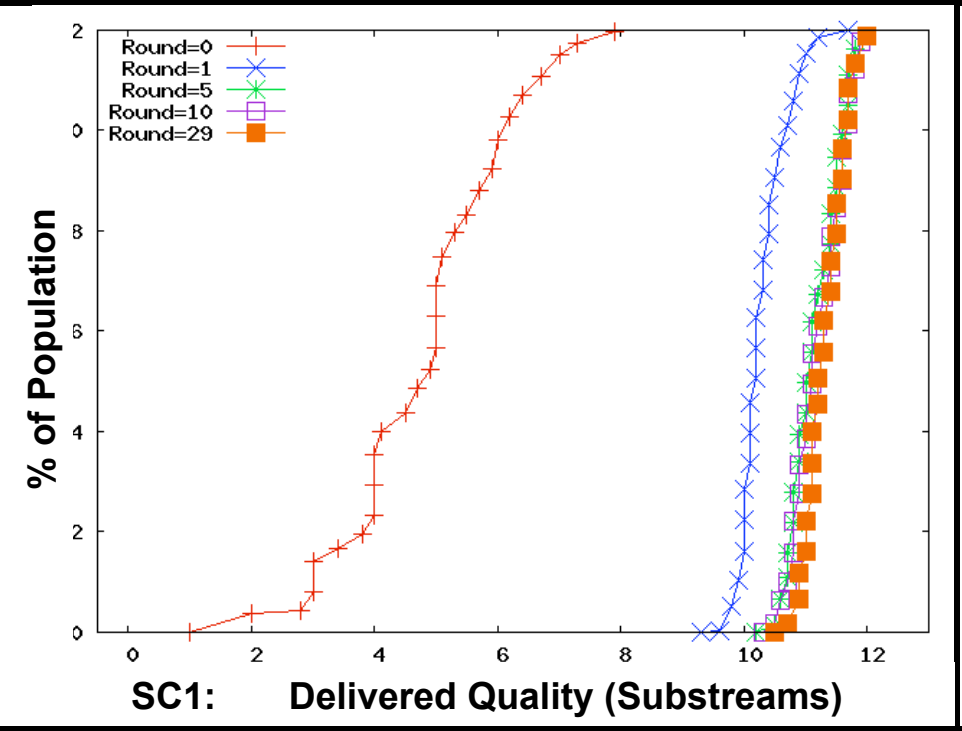

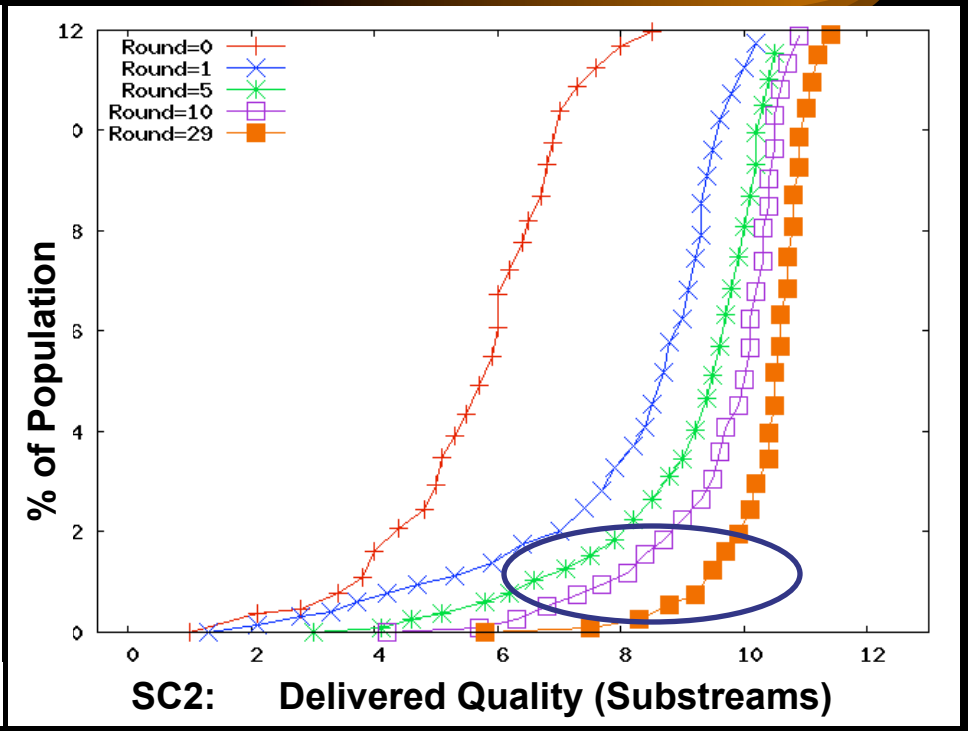

- *OMR significantly improves the quality despite heterogeneity of BW*
- Unlike SC1, the delivered quality remains relatively skewed in SC2
  - The large % of low BW peers in SC2, increases the likelihood that a high BW peer *P* has several low BW parents or children. → decreases the value of *CF* and *PF* in the reaction algorithm

Conclusions

- Clustering significantly reduces the performance of live Swarm-based P2P streaming
- OMR detects poor connectivity in the overlay and reacts by rewiring minimum overlay connections
- Our preliminary simulation results
  - Demonstrate OMR is able to improve performance
  - Reveal a few more issues that requires further investigation
- Ongoing Work:
  - The effect of churn and peer bandwidth heterogeneity,
  - The effect of α and β on controlling the tradeoff between responsiveness and aggressiveness
- For more information visit *http://mirage.cs.uoregon.edu*