Architecting Bandwidth-Aware Peer-to-Peer Overlays

Daniel Ellsworth and Reza Rejaie - University of Oregon

1. Introduction

Problem: Non-interactive video streaming in P2P Networks.

Design Goals:

- Maximum delivered BW to each peer, *BW-aware overlay*.
- Scalability (of data and control traffic) with group size.
- Ability to accommodate churn and heterogeneity.

Existing Solutions:

- Adopted the idea of "application-level multicast" to form a single (or multiple) delay-optimized, source-rooted tree(s)
- Video is encoded in layered (or multiple-description) format where each layer is delivered through a separate tree.
- Delivered BW over a tree structure is inherently limited by minimum outgoing BW among upstream peers in the tree,

➢*Existing solutions are unable to construct bandwidth-aware overlays, thus unable to maximize delivered quality to each peer.*

2. Design Issues

- Unstructured (Gnutella-like) P2P networks are preferred because they are more robust to churn than DHTs.
- Overlay connections must be *congestion controlled.*
- Each peer should connect to the overlay at multiple points (parents) to maximize its incoming bandwidth.
- A mechanism for streaming video from multiple congestion controlled parents is required

Main Components for each peer:

1) Peer Discovery (PD): How to find other peers?

2) Parent Selection (PS): How to select "good" parents?

- Good parents have minimum delay and maximum BW. Note that these two criteria might be in conflict.
- Scalable estimation of pair-wise delay is feasible (e.g. using GNP), but active measurement of pair-wise bandwidth is expensive and does not scale!

3. Proposed Approach

- *The Idea: Each peer selfishly and independently searches for a subset of parents to maximize its own BW (i.e. quality).*
- The overlay construction problem is formulated as a collection of local optimizations instead of a single global optimization.
- Competing peers reach an "equilibrium" in which the resulting overlay is efficient and the delivered BW to each peer is maximized.
- We also developed a receiver-driven streaming scheme from multiple congestion-controlled parents to address delivery.

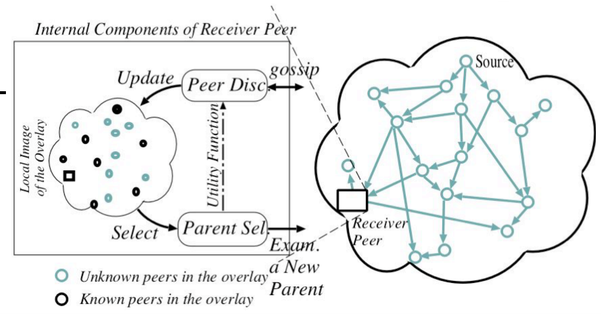

The search for good parents is conducted in two phases:

1) Gossip-based Peer Discovery (PD): identifies *potentially good parents* based on their static properties (i.e. their utility).

- Relative utility of a parent for a peer depends on its relative distance and its access link BW.
- Potentially good parents are maintained in a local ***image***.
- This approach heuristically reduces the scope of search from the entire group to a small number of potentially good parents.

2) Selfish Parent Selection (PS): progressively examines peers in the local image to identify the best subset of parents.

- Actual BW of a parent is *passively measured during delivery.*

4. Gossip-based Peer Discovery

- Each peer periodically selects a random target peer from its local image to request a gossip message (i.e. pull gossiping).
- Given a well-known utility function (i.e. joint-ranking), each peer can determine the relative utility of any two peers.
- The target peer provides information about parents from its local image that have maximum utility for the requesting peer.
- Information in a gossip message is integrated into the local image in order to improve some aspect of the image quality; 1) freshness or 2) overall utility.
- Our initial results show that each peer can find *N* best parents (with highest utility) within a few rounds independent of group size.

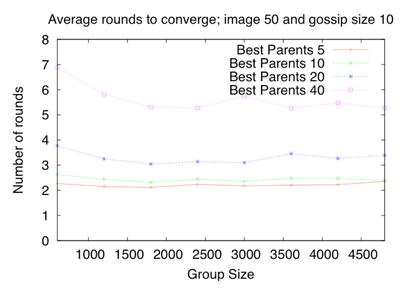

5. Selfish Parent Selection

- *Key Question: How does **uncoordinated** and selfish parent selection by individual peers collectively affect the **shape** and **stability** of the resulting unstructured overlay?*

1) Leverage available BW from each parent as an *implicit signal* to detect any relevant change in the overlay (e.g. peer departure, arrival) or network (e.g. shared congestion).
2) Fundamental tradeoff between overlay stability and responsiveness to a change.
3) To damp oscillation, reaction of affected peers to a signal are diversified, hysteresis and binning strategies are incorporated.

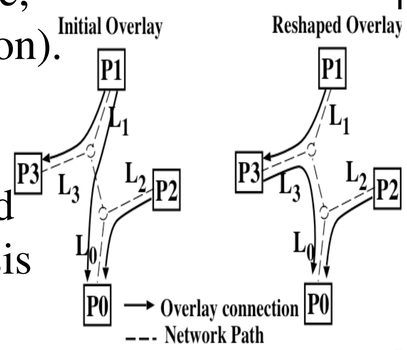

5. Conclusion and Future Work

- We sketched a scalable bandwidth-aware overlay construction scheme as a two-level search.
- A new idea is to cast the problem as a collection of uncoordinated local optimizations rather than a single global optimization.
- Detailed evaluations of both gossip-based peer discovery and parent selection are ongoing, and implementation will follow.
- Examination of gossip as a generic resource discovery tool.
- For further details visit ***http://mirage.cs.uoregon.edu***