The Feasibility of Supporting Large-Scale Live Streaming Applications with Dynamic Application End-Points

SIGCOMM' 04

Presented by 江文德

Outline

- Introduction
- Live Streaming Workload
- Key Requirements for Feasibility
  - Are There Enough Resource?
  - Is There Any Stability?
  - Can Efficient Overlays be Constructed?
- Summary
- References

Introduction

- Commercial content delivery network such as *Akamai Technologies* [1] have developed and deployed large-scale dedicated infrastructure to deliver both live streams and video-on-demand.
- In contrast to infrastructure, application end-point overlay multicast has recently received attention.

Introduction

- The lack of dependence on infrastructure support makes application end-point architectures easy to deploy and economically viable.
- Question:
  - Whether or not such architectures are feasible at large scales of 1000 to 10000 nodes.

Live Streaming Workload

- Data collection:
  - The logs are collected from the thousands of streaming servers belonging to Akamai Technologies.
  - The logs collected over a 3-month period from Oct. 2003 to Jan. 2004.
  - The traffic consists of three of the most popular streaming media formats, QuickTime, Real, and Windows Media.

Live Streaming Workload

- Summary Statistics

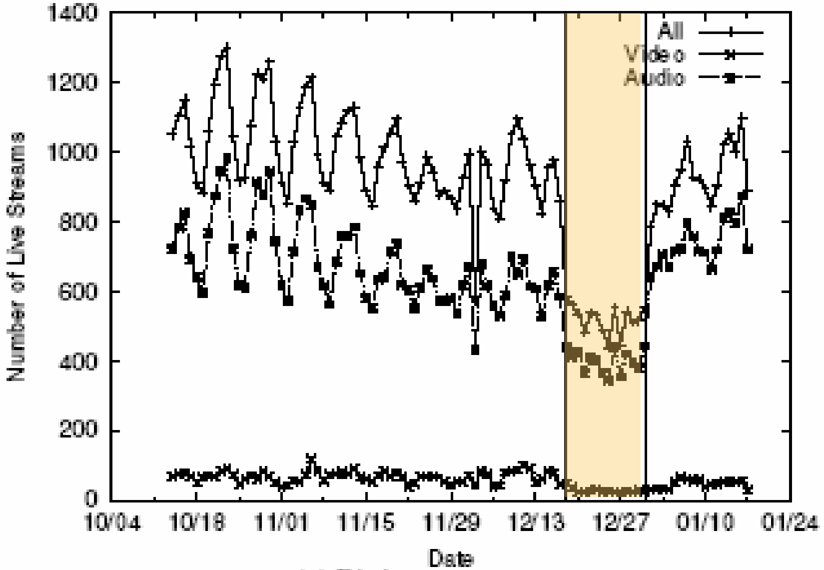

(a) Distinct streams.

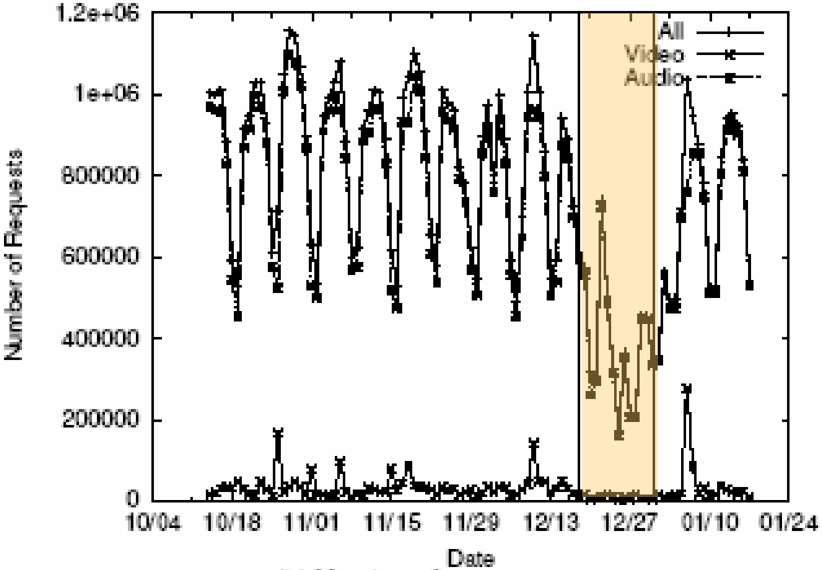

(b) Number of requests.

Live Streaming Workload

- Summary Statistics
  - There were a total of 660 large scale streams, of which 55 were video streams and 605 were audio streams.

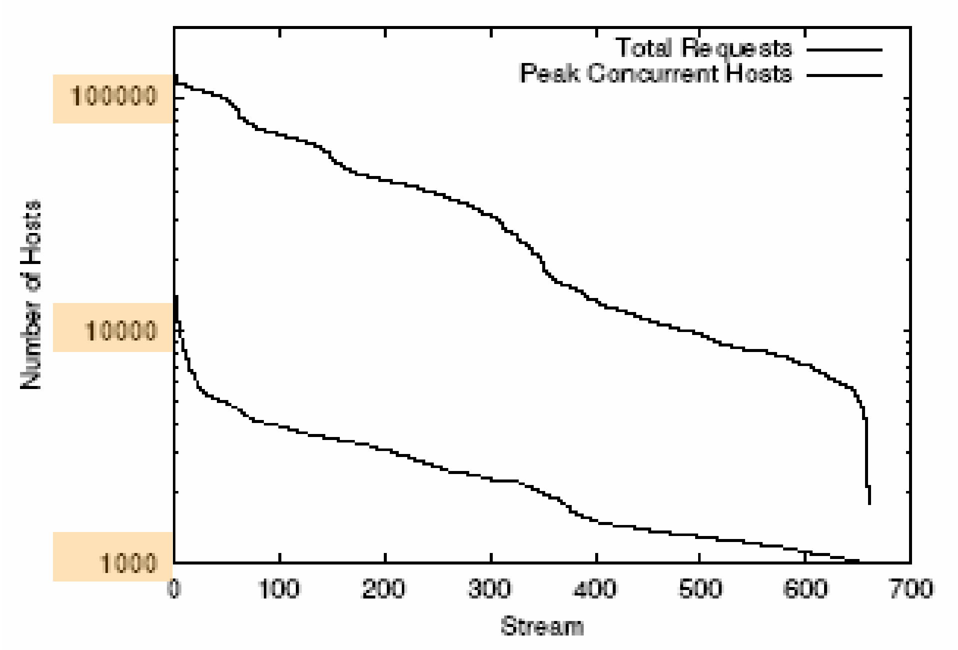

Key Requirements for Feasibility

- Are there enough resource to construct an overlay?
- Can a stable and connected overlay be maintained in the presence of group dynamics?
- Can an efficient overlay be constructed?

Are There Enough Resource?

- Outgoing Bandwidth Estimation.
  - Step1:
    - User voluntarily go to the web site (broadbandreports.com) to test their connection speeds.
    - For the 72% of hosts, we assigned their bandwidth values to the ones reported by broadbandreports.com.
  - Step2:
    - Aggregate the remaining IP addresses into /24 prefix blocks and conducted packet-pair measurements to measure the bottleneck bandwidth to several hosts in each block.
    - Use the measure results form the prefix blocks to assign bandwidth estimations to an additional 7.6% of the IP addresses.

Are There Enough Resource?

- Outgoing Bandwidth Estimation.
  - Step3:
    - Translate access technology into raw bandwidth.

Access technology	Packet-pair measurement	Outgoing bandwidth estimate
Dial-up modems	0 kbps $\leq$ BW $<$ 100 kbps	30 kbps
DSL, ISDN, Wireless	100 kbps $\leq$ BW $<$ 600 kbps	100 kbps
Cable modems	600 kbps $\leq$ BW $<$ 1 Mbps	250 kbps
Edu, Others	BW $\geq$ 1 Mbps	BW

    - Assign estimates to additional 7.1% of the IP addresses.
  - Step4:
    - Use the host's DNS name to infer its access technology.
    - Get estimates for an additional 2.2%+1.2% of the IP addresses.

Are There Enough Resource?

- Degree:
  - To simplify the presentation, we normalize the bandwidth value by the encoding bit rate.
  - E.g. outgoing link bandwidth = 300kbps, encoding bit rate = 250kbps ➔ [300/250] = 1 degree.

Type	Degree-bound	Number of hosts
Free-riders	0	58646 (49.3%)
Contributors	1	22264 (18.7%)
Contributors	2	10033 (8.4%)
Contributors	3-19	6128 (5.2%)
Contributors	20	8115 (6.8%)
Unknown	-	13735 (11.6%)
Total	-	118921 (100%)

Are There Enough Resource?

- Resource Index [2] :
  - The ratio of the supply of bandwidth to the demand for bandwidth in the system for a particular encoding bit rate.

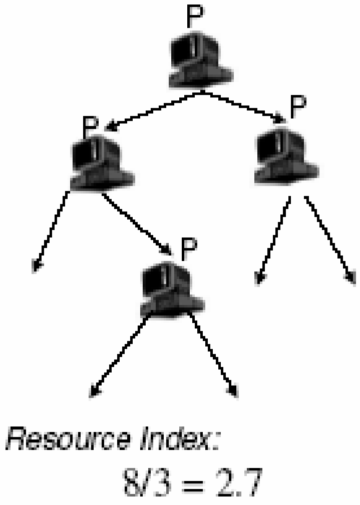

Are There Enough Resource?

- A Resource Index of 1 indicates that the system is fully saturated, and a ratio less than 1 indicates that not all the participating hosts in the broadcast can receive the full encoding rate.
- As the Resource Index gets higher, the environment becomes more feasible.

Are There Enough Resource?

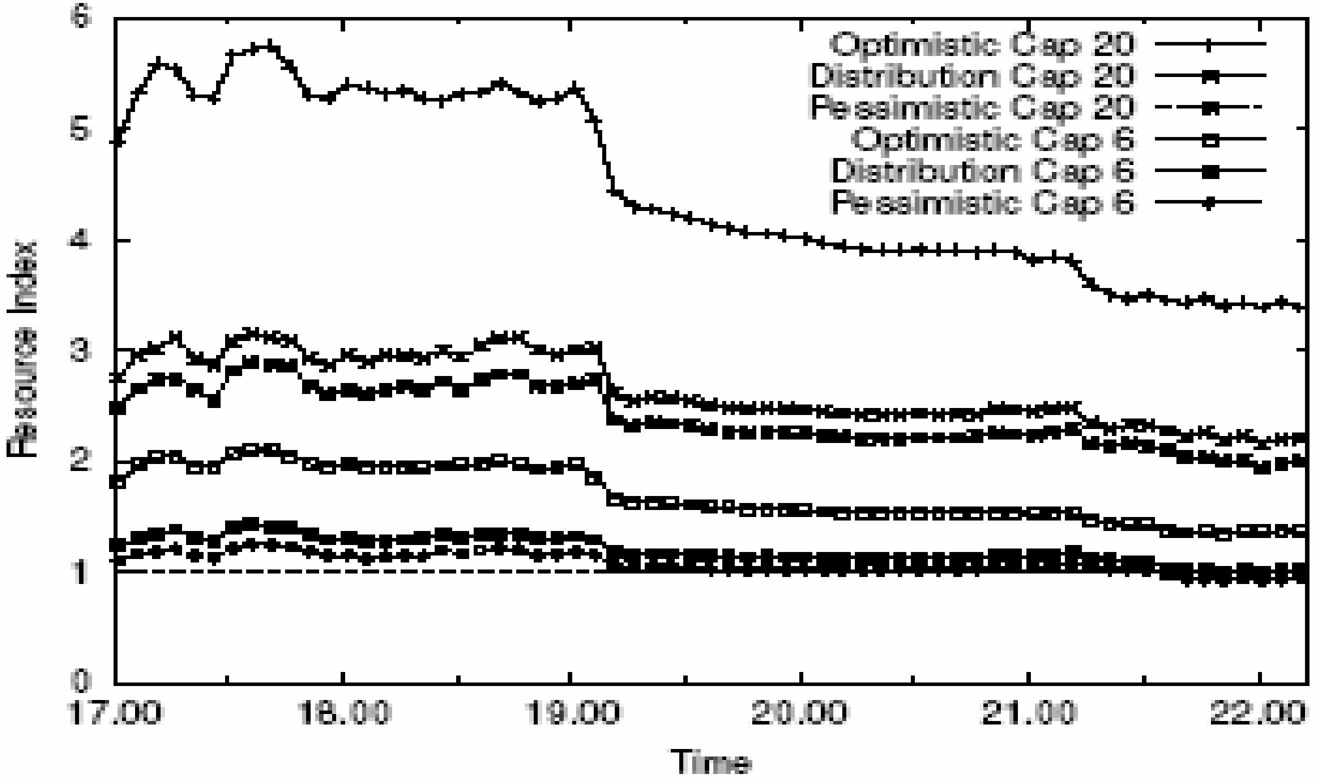

Are There Enough Resource?

- Trace Replay: Single-Tree Protocol
  - 3 assignment algorithms for the 10% of IP Address without bandwidth estimates.
  - *The pessimistic estimate* assumes that all unknowns are free riders.
  - *The distribution algorithm* assign a random value drawn form the same distribution as the known resources.
  - *The optimistic estimate* assumes that all unknowns can contribute up to the maximum resource allocation.

Are There Enough Resource?

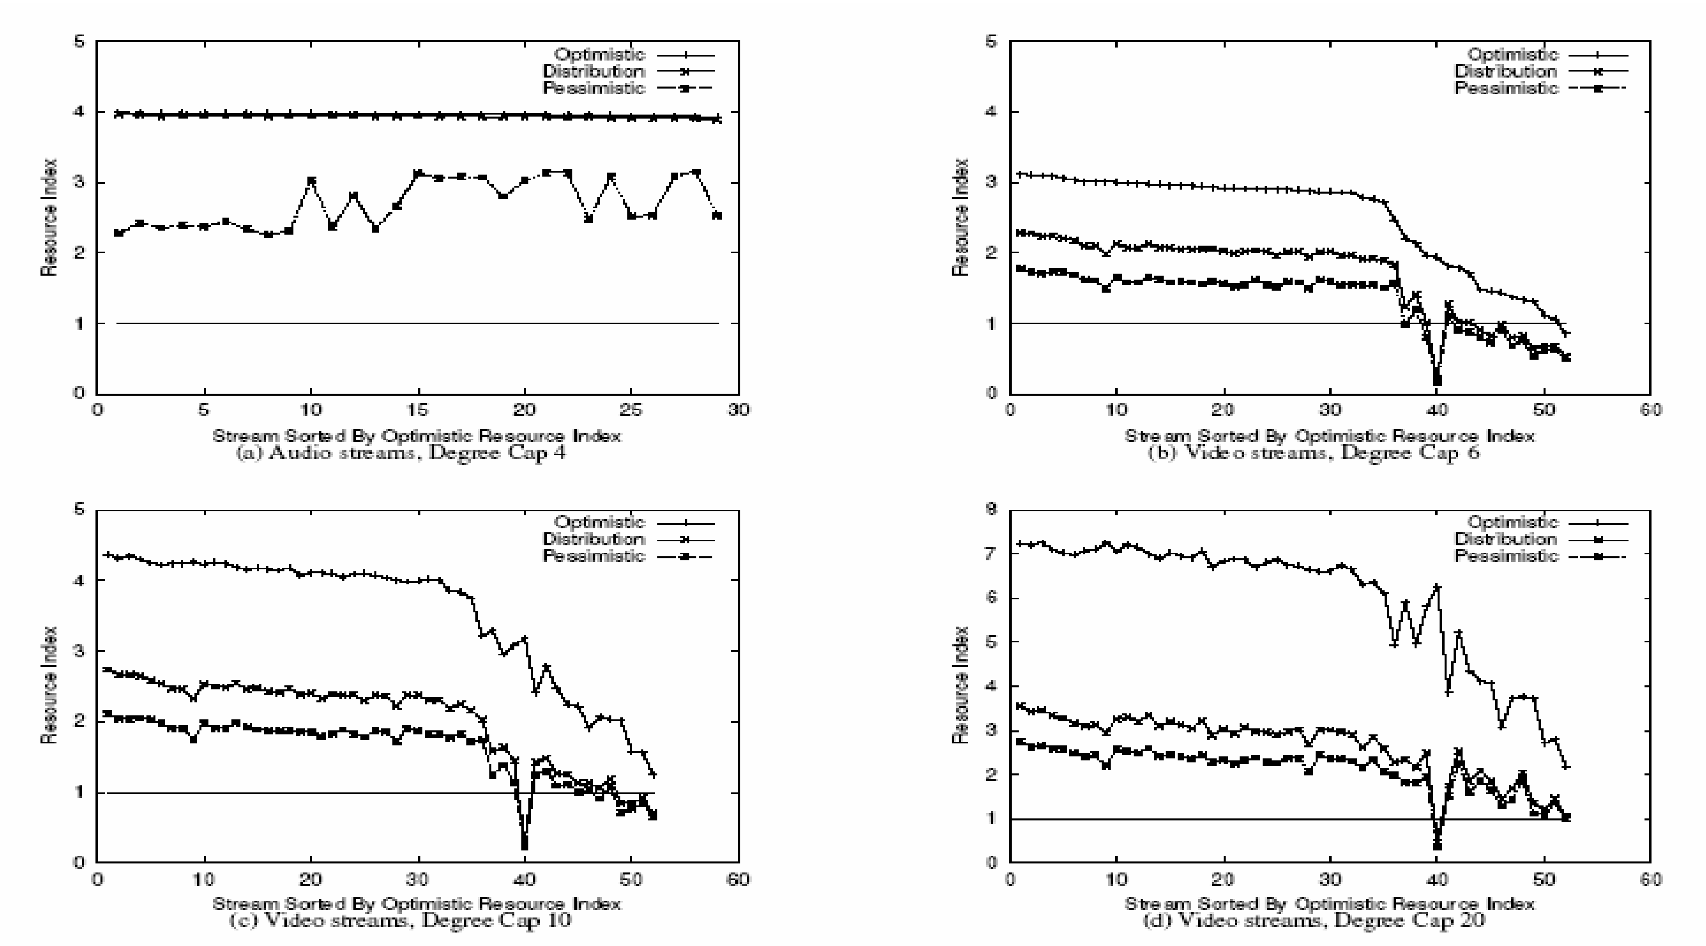

(a) Audio streams, Degree Cap 4

(b) Video streams, Degree Cap 6

(c) Video streams, Degree Cap 10

(d) Video streams, Degree Cap 20

Are There Enough Resource?

- Trace Replay: Multiple-Tree

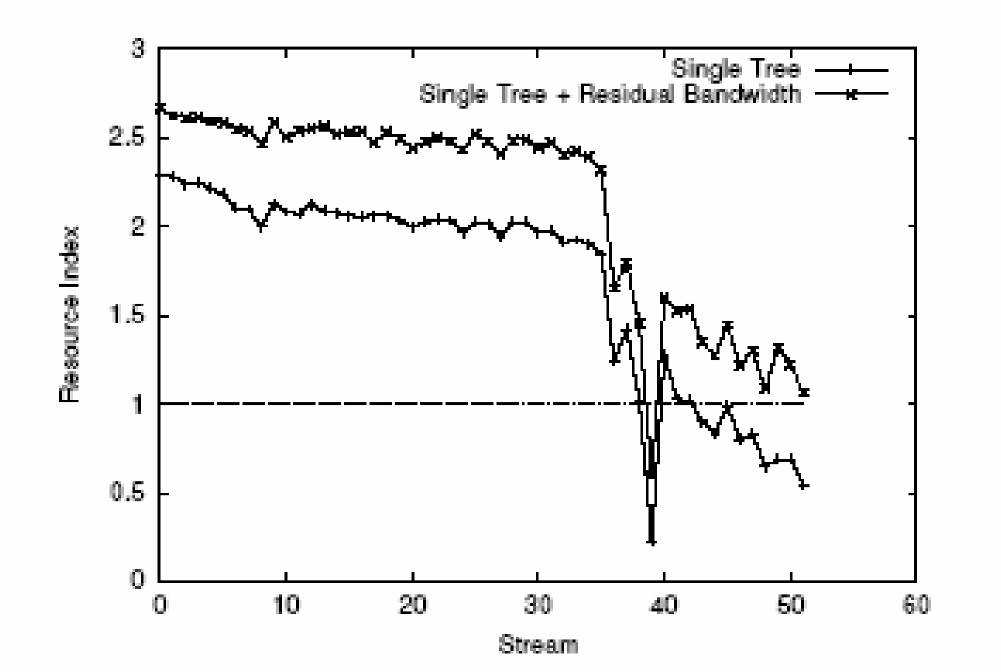

Is There Any Stability?

- Extreme Group Dynamics

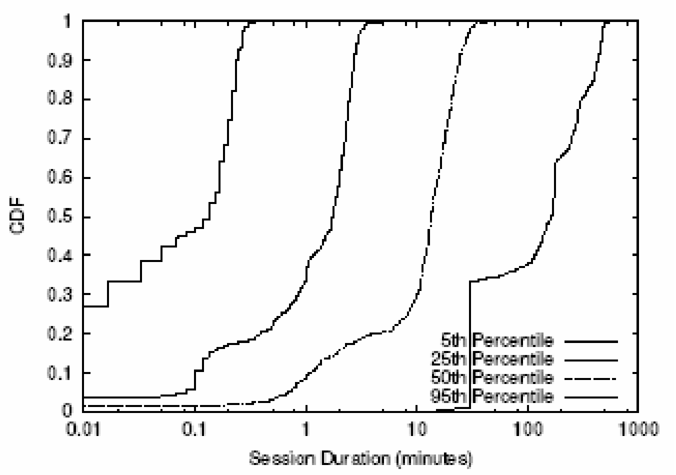

- Stability Metrics
  - Mean interval between ancestor change for each incarnation.
  - Number of descendants of a departing incarnation.

Is There Any Stability?

- Overlay Protocol
  - Simulate the effect of group dynamics on the overlay protocol using a trace-driven event-based simulator.
  - The simulator takes the group dynamics trace from the real event and the degree assignments based on the techniques in the previous section, and simulates the overlay tree at each instant in time.
  - Host Join, Host Leave, Parent Selection.

Is There Any Stability?

- Parent Selection
  - *Oracle:*
    - A host chooses the parent who will stay in the system longer than itself.
    - This algorithm requires future knowledge and cannot be implemented in practice.
  - *Longest-first:*
    - This algorithm attempts to predict the future and guess which nodes are stable by using the heuristic that if a host has stayed in the system for a long time, it will continue to stay for a long time.

Is There Any Stability?

- Parent Selection
  - *Minimum depth:*
    - A host chooses the parent with the minimum depth.
  - *Random:*
    - A host chooses a random parent.

Is There Any Stability?

- Results: Single-Tree Protocol

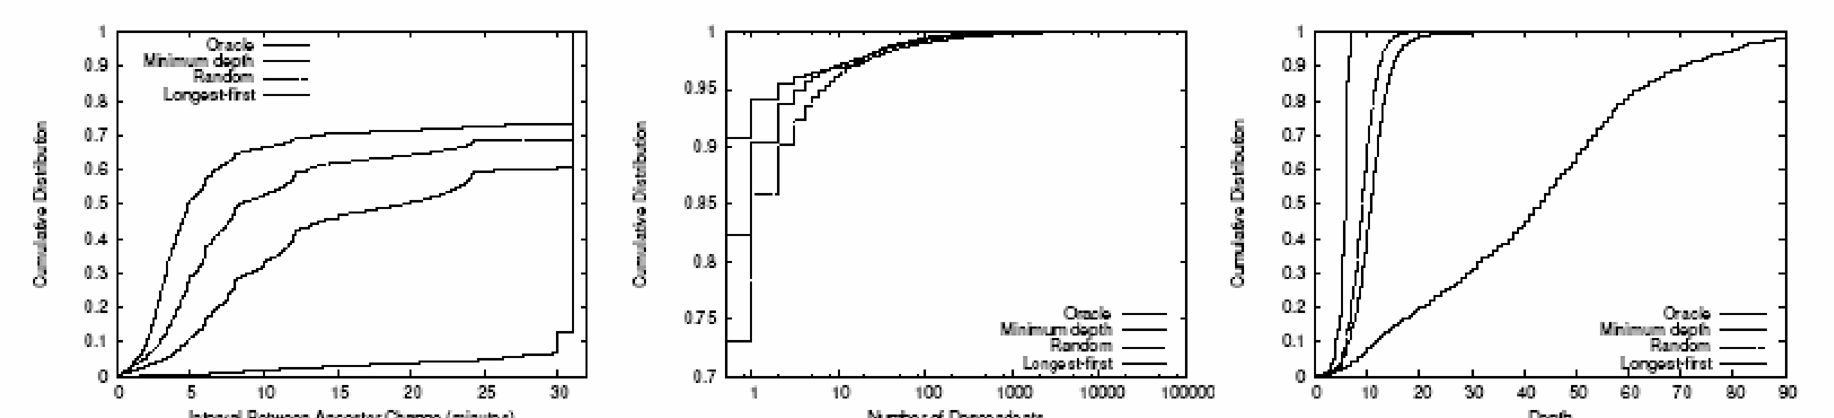

(a) CDF of interval between ancestor change. (b) Number of descendants of a departing host. (c) Tree depth.

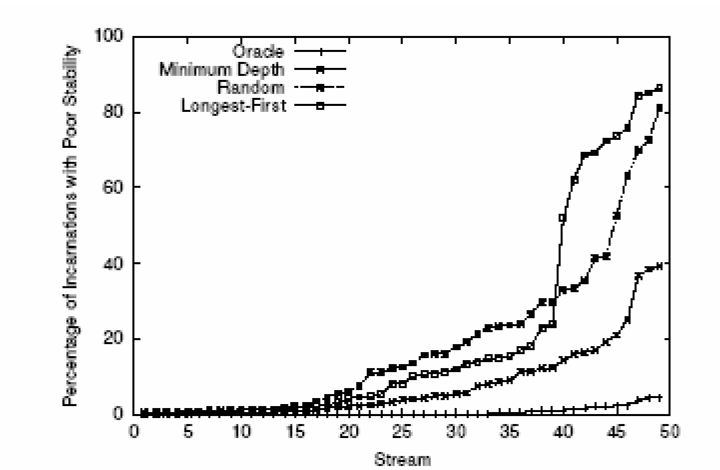

Is There Any Stability?

- Results: Multiple-Tree Protocol

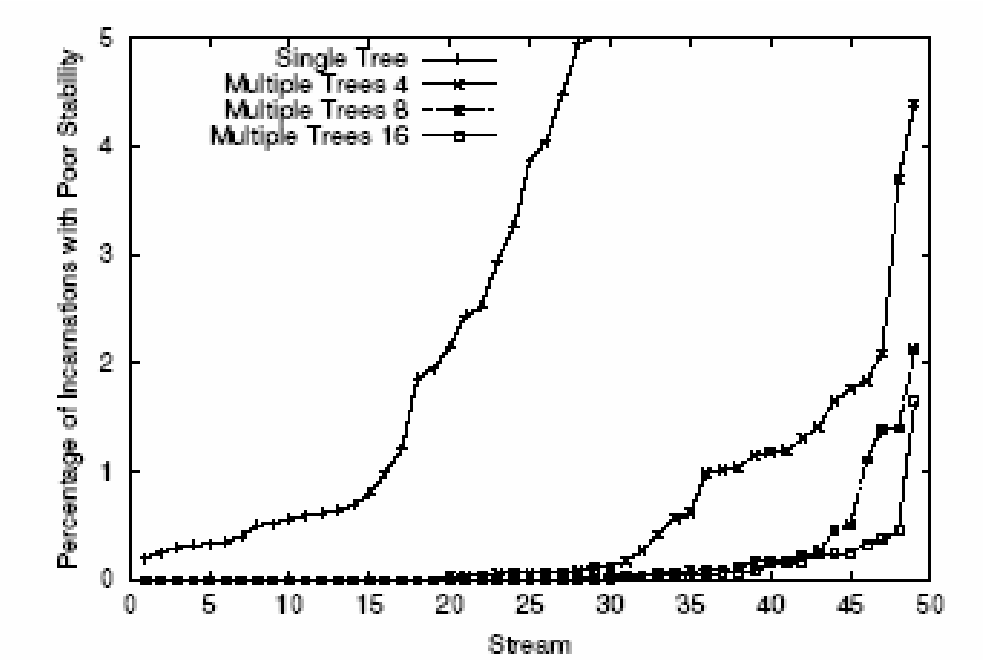

Can Efficient Overlays be Constructed?

- An efficient overlay is one in which the overlay structure closely reflects the underlying IP network.
- The challenge is to enable hosts to discover other *nearby* hosts that may be used as parents.

Can Efficient Overlays be Constructed?

- Clustering Policies
  - Random
    - the clusters are agnostic of network proximity.
  - Delay-based Clustering
    - short delays are reasonably correlated with good bandwidth performance, which is also important for streaming applications.
  - Geographic Clustering
    - roughly approximates network distance.

Can Efficient Overlays be Constructed?

- Clustering Quality

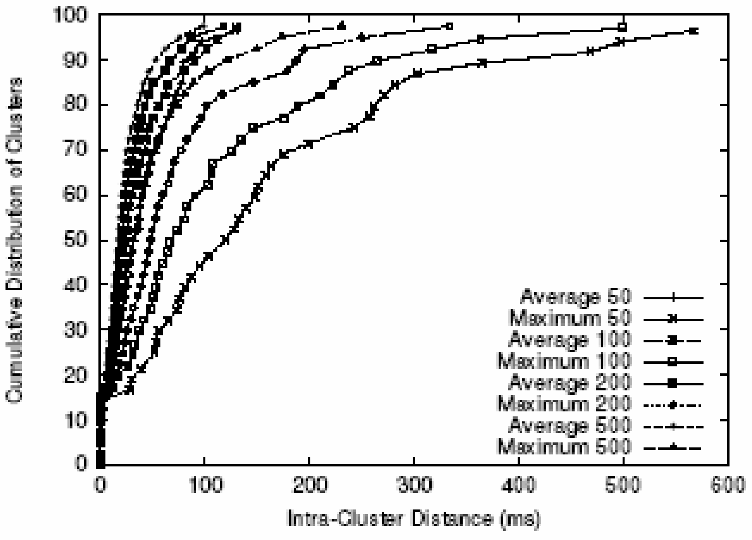

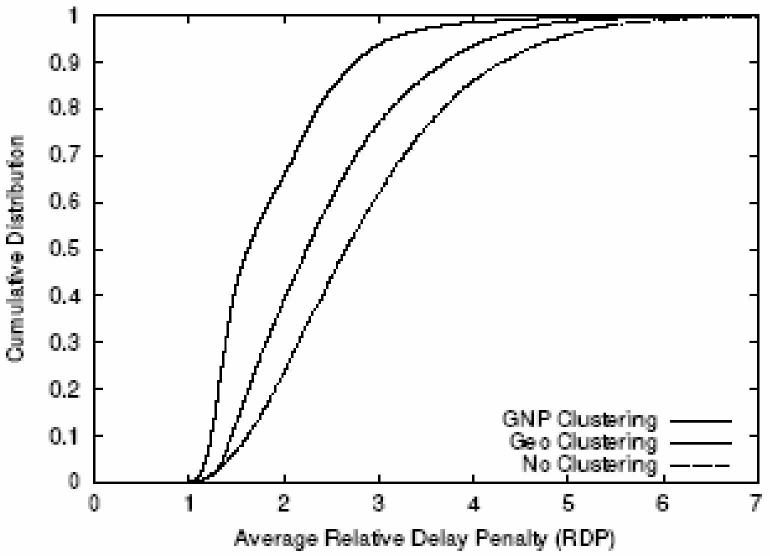

Can Efficient Overlays be Constructed?

- Clustering Quality

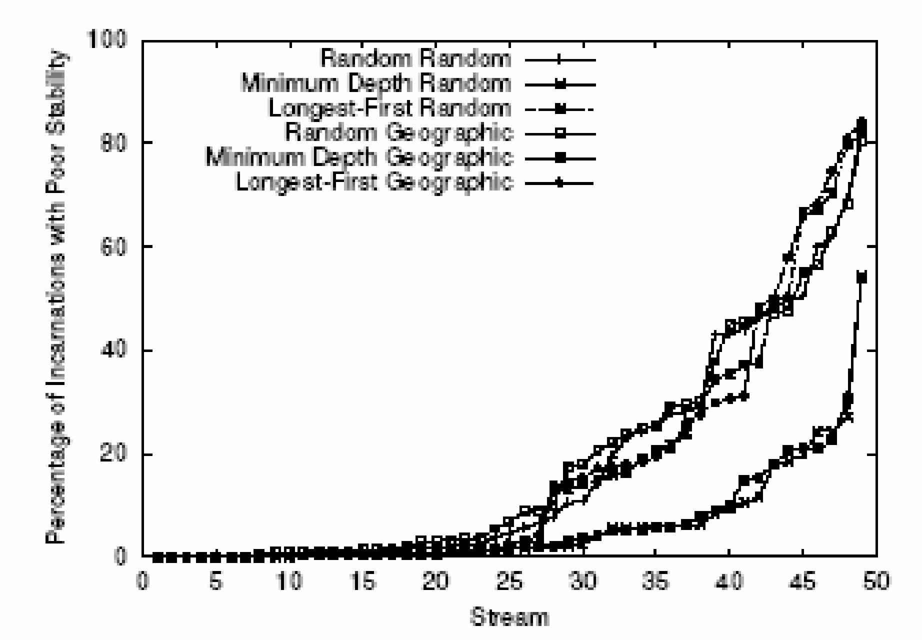

Summary

- Using a single-tree protocol and a single encoding rate, all audio streams have abundant resources and most video streams have enough inherent resources.
- In resource constrained environments, using multiple-tree protocols can increase the supply of resources in the system and improve the situation.

Summary

- Minimizing depth in single-tree protocols provides good stability performance.
- The use of multiple-tree protocols can significantly improve the perceived quality of streams at the expense of an increase in protocol activity, overhead, and complexity.
- Simple clustering techniques improve the efficiency of the overlay structure.

References

- [1] Akamai. http://www.akamai.com/.
- [2] Y. chu, A. Ganjam, T. S. E. Ng, S. G. Rao, K. Sripanidkulchai, J. Zhan, and H. Zhang, Early Experience with an Internet Broadcast System Based on Overlay Multicast. In *Proceedings of USENIX*, 2004.