Measurement-Based Optimization Techniques for Bandwidth-Demanding Peer-to-Peer Systems

INFOCOM 2003

謝志峰

2003/12/17

Outline

- Introduction
- Bandwidth-Demanding Peer-to-Peer Systems
- P2Cast
- Methodology
- Analyses of Peers Selection Techniques
- Application Case
- Conclusions

Introduction

- **Measurement-based optimization is one important strategy to improve the performance of bandwidth-demanding peer-to-peer systems.**

- A basic strategy is to measure the network:
  - using light-weight techniques
  - using the measured information to organize peers
  - high bandwidth low congestion network paths.

Introduction

- It is to provide systematic quantitative information about **light-weight measurement-based optimization techniques** in the peer-to-peer environment through **trace-based** experiments.

- We study three basic lightweight techniques :
  - RTT(round-trip time probing)
  - 10KB TCP probing, and
  - BNBW(bottleneck bandwidth probing)

Bandwidth-Demanding Peer-to-Peer Systems

- We assume a Napster-like media file sharing system :
  - A client peer issues a request for a desired media file to a directory service.
  - The directory service returns a list of server peers that have the wanted media file.
  - Given a set of candidate server peers to choose from and the TCP bandwidth achievable from the best server peer.

Bandwidth-Demanding Peer-to-Peer Systems

- Peers self-organize into an overlay tree rooted at the source, and data is sent along links in the overlay tree.
  - First, when a member joins the group, it needs to decide which parent it should attach.
  - Second, even after making an initial decision, members still need to periodically re-evaluate their choice of parent.

Bandwidth-Demanding Peer-to-Peer Systems

- A member does not select any descendant as a parent.
- A member can filter out those peers that receive worse application level performance.

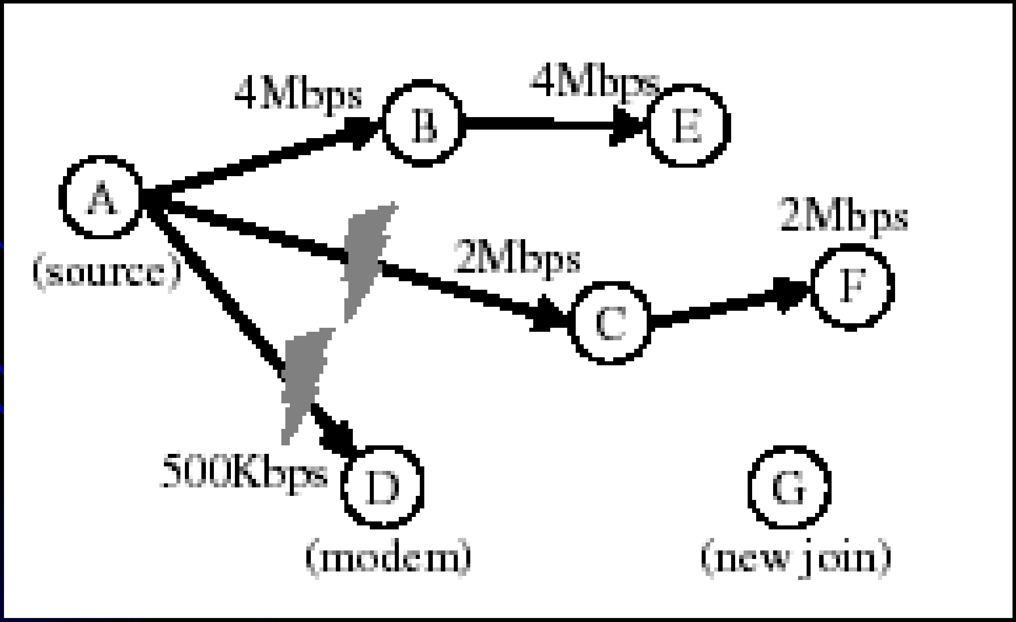

P2Cast

- Base Stream Forwarding:P2Cast clients need to be able to forward the received base stream to other clients so that clients and the server can form an application-level multicast tree over which the base stream is transmitted.
- Patch Serving: P2Cast clients need to have sufficient storage to cache the initial part of the video. A P2Cast client can then serve the patch to other clients.

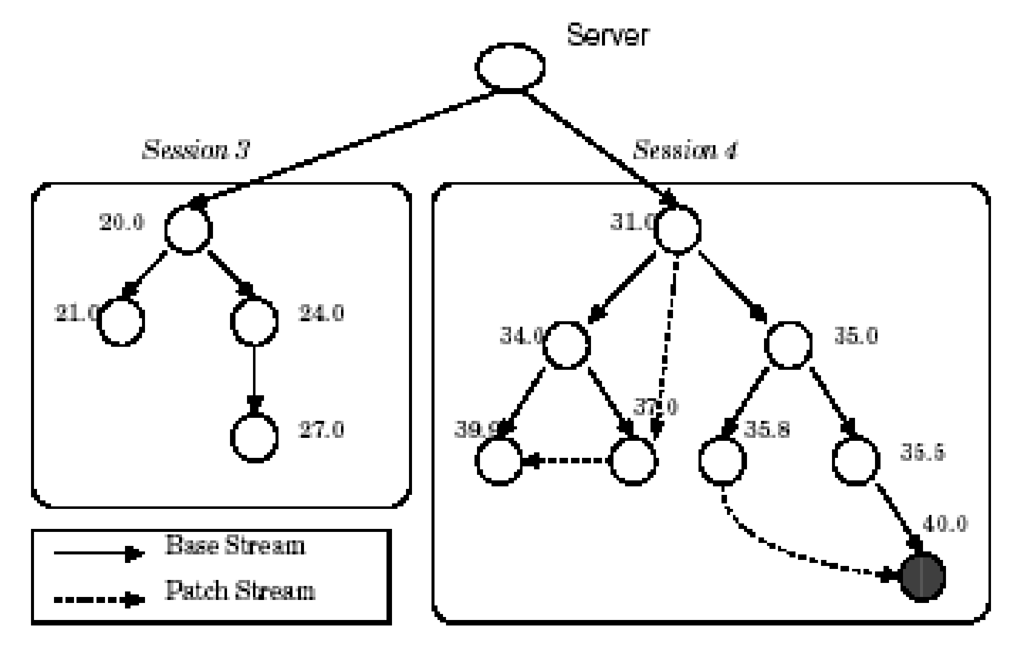

Bandwidth-Demanding Peer-to-Peer Systems

- Performance Metric:
  - The key performance metric is the *convergence time*.
  - We define the convergence time of a peer to be the amount of time after the initial join it takes for the peer to receive more than 95% of the *stable bandwidth* for 30 seconds.

Methodology

- We used 4 data collection hosts located in Carnegie Mellon University (CMU), University of Illinois Urbana- Champaign (UIUC) and University of Alberta, Canada.

Host	Location	Link Speed	# Peers	TCP Avg
1	CMU	10 Mbps	2705	129 kbps
2	CMU	640/90 kbps ADSL	880	111 kbps
3	UIUC	10 Mbps	3116	140 kbps
4	U of Alberta	10 Mbps	3805	130 kbps

TABLE I

COLLECTION HOST PROPERTIES

Analyses of Peers Selection Techniques

- we present analytical results for three basic light-weight peer selection techniques:

  - In RTT probing(round-trip time probing): we measure the RTT to each candidate peer using a single 36 byte ICMP ping message; the candidate peer with the smallest RTT is then selected.

  - In 10KB TCP probing, a TCP connection is opened to each candidate peer and 10KB of data is downloaded; the peer that provides the fastest download is selected.

  - In BNBW probing(bottleneck bandwidth probing): it is used to measure the BNBW to each candidate peer; the peer with the largest BNBW is selected.

Analyses of Peers Selection Techniques

- The goal is not only to quantify the performance of these techniques, but also to understand the techniques' inner-workings, weaknesses, and complementarity.

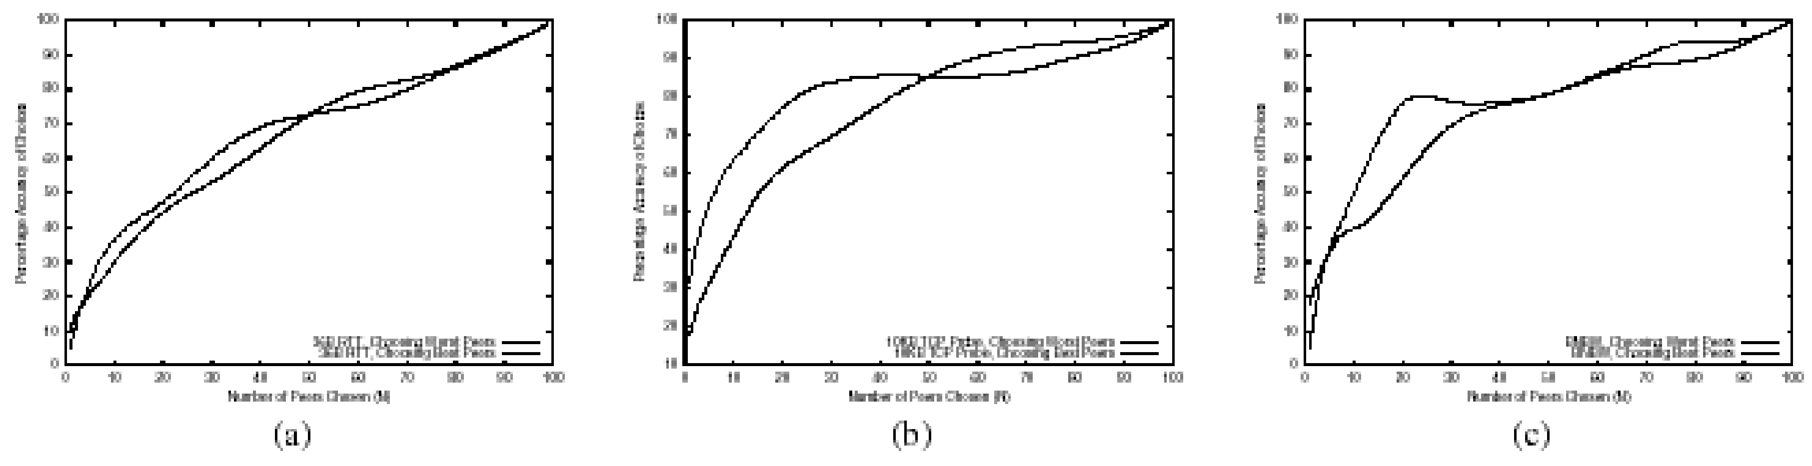

Fig. 3. Percentage accuracy of choices versus number of peers (best or worst) chosen. (a) 36B RTT probing. (b) 10KB TCP probing. (c) BNBW probing.

Analyses of Peers Selection Techniques

- This experiment is repeated 1000 times.
- The main observation is that these basic techniques can typically achieve 40 to 50% of optimal performance (or 3 to 4 times better than random selection).

Host	Location	Random	36B RTT	10KB Probe	BNBW
1	CMU	0.13	0.42	0.47	0.48
2	CMU (ADSL)	0.24	0.52	0.66	0.33
3	UIUC	0.13	0.27	0.43	N/A
4	U of Alberta	0.15	0.40	0.48	N/A

TABLE II

AVERAGE O.R. OF BASIC TECHNIQUES (95% CONFIDENCE INTERVALS < 0.02)

Analyses of Peers Selection Techniques

- Table IV shows the average O.R. of the chosen peer if an adaptive application is able to observe the performance of the top 5 candidate peers recommended by each technique.
- Table V shows the results when the top candidate set size is increased to 10.

Host	Location	Random	36B RTT	10KB Probe	BNBW
1	CMU	0.33	0.82	0.80	0.83
2	CMU (ADSL)	0.52	0.91	0.90	0.58
3	UIUC	0.33	0.58	0.81	N/A
4	U of Alberta	0.35	0.69	0.80	N/A

TABLE IV

AVERAGE O.R. OF BEST PEER AMONG TOP 5 CANDIDATES (95% CONFIDENCE INTERVALS < 0.02)

Host	Location	Random	36B RTT	10KB Probe	BNBW
1	CMU	0.46	0.86	0.91	0.95
2	CMU (ADSL)	0.67	0.95	0.96	0.69
3	UIUC	0.44	0.80	0.95	N/A
4	U of Alberta	0.47	0.83	0.91	N/A

TABLE V

AVERAGE O.R. OF BEST PEER AMONG TOP 10 CANDIDATES (95% CONFIDENCE INTERVALS < 0.02)

Analyses of Peers Selection Techniques

- The idea is to compute the percentage of time where each technique is selecting a peer that is within 90% of the optimal peer, and how often the techniques agree in picking the same 90% optimal peer.

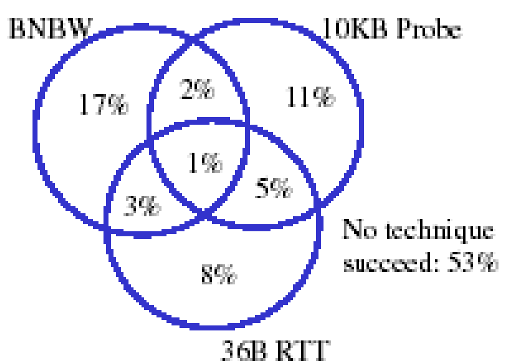

Fig. 5. Percentage of cases where each technique chooses a peer within 90% of optimal.

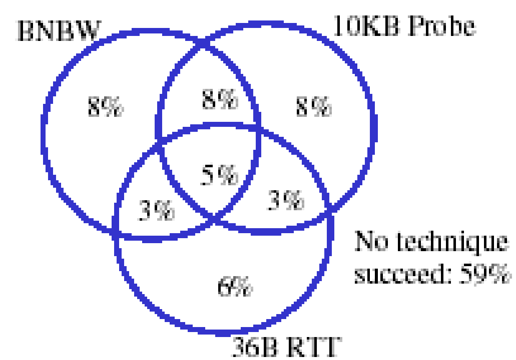

Fig. 6. Percentage of cases where, after using RTT to select 5 candidates, each technique chooses a peer within 90% of optimal.

Application Case

- Media File Sharing:
  - We propose joint ranking as one plausible way of exploiting the complementarity of the basic techniques and evaluate its effectiveness in the media file sharing application.
  - we first rank the candidate peers based on each technique. For each candidate peer, we sum up the rank values assigned by the basic techniques.
  - We then choose the peer with the smallest sum of ranks.

Application Case

- The best technique we have found turns out to be case (o) in Table VI, where RTT is used to first select 5 candidates, then the three basic techniques are used in joint ranking.

Technique		O.R.
100 candidates	(a) Random	0.13
	(b) 36B RTT	0.42
	(c) 10KB Probe	0.47
	(d) BNBW	0.48
100 candidates	(e) 36B RTT joint BNBW	0.49
	(f) 36B RTT joint 10KB Probe	0.54
	(g) BNBW joint 10KB Probe	0.50
	(h) 36B RTT joint BNBW joint 10KB Probe	0.55
5 out of 100 candidates by 36B RTT	(i) Oracle	0.82
	(j) 10KB Probe	0.53
	(k) BNBW	0.50
	(l) 36B RTT joint BNBW	0.51
	(m) 36B RTT joint 10KB Probe	0.55
	(n) BNBW joint 10KB Probe	0.55
	(o) 36B RTT joint BNBW joint 10KB Probe	0.58

TABLE VI

AVERAGE O.R. OF JOINT RANKING TECHNIQUES (95% CONFIDENCE INTERVALS < 0.02)

Application Case

- In this section, we evaluate the effectiveness of different peer selection techniques in overlay multicast streaming.
- An overlay multicast protocol that adapts to both bandwidth and latency, to use various peer selection techniques to optimize performance.
  - RTT: This technique uses single-packet RTT probe to select peers.
  - RTT filter + 10K: This technique selects at most 5 of the candidate peers based on the best RTT estimates and perform 10KB of file transfer using TCP in parallel.
  - RTT filter + 1-bit BNBW: This technique selects at most 5 peers with the lowest RTT estimates and then chooses the peer with the highest bottleneck bandwidth.

Application Case

- This is shown in Figure 7, which plots the mean bandwidth as a function of time.
- We observe that for all techniques, the overlay reaches similar stable average performance within 30 seconds.

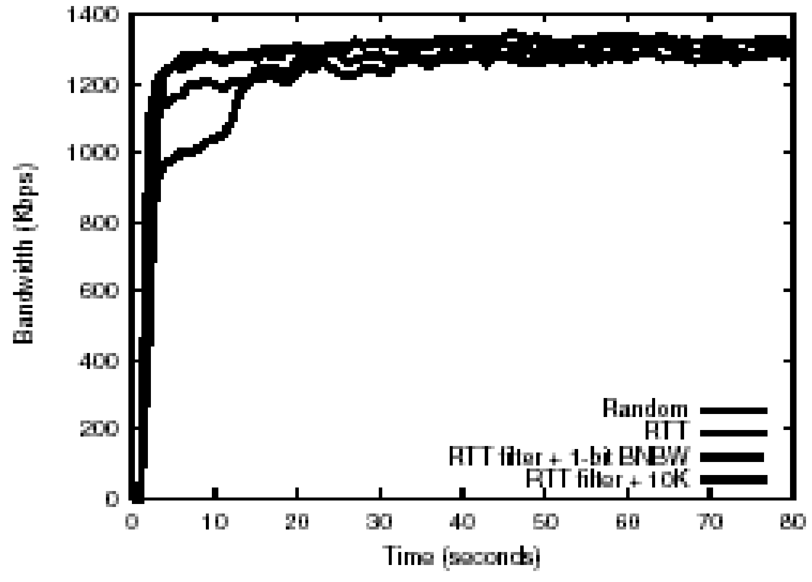

Fig. 7. Mean receiver bandwidth as a function of time at 1.5 Mbps source rate.

Conclusions

- Light-weight measurement-based techniques such as RTT probing, 10KB TCP probing and BNBW probing can perform reasonably well, achieving 40 to 50% optimal performance in media file sharing.

- In our experiments, typically an 80% optimal peer can be found by trying less than 5 candidates.

- We show that by combining all techniques together in a simple fashion, we can boost the performance of media file sharing to 60% optimal and double the worst case performance.