Improving Traffic Locality in BitTorrent via Biased Neighbor Selection

From ICDCS'06

Presented by Ching-Lan Wang

October 20, 2006

Outline

- Introduction
- Related work
- Problem
- Biased neighbor selection
  - Modifying trackers and client
  - P2P traffic shaping devices
- Performance
  - A homogeneous network
  - A heterogeneous network

Introduction

- To distribute a file via BitTorrent, the provider first generates a separate file containing some torrent files and runs a tracker which keeps track of all nodes downloading this file.

- Each peer upon first joining the network, contacts the tracker. The tracker then **randomly** selects N nodes, out of all the nodes in the network, and hands the list back to peer.

Introduction (cont.)

- As long as the original seed (the host that starts sharing the file) has moderately high upload bandwidth.
- BitTorrent ignore traffic costs at ISPs and generate a large amount of cross-ISP traffic.
- Since each BitTorrent node prefers to upload to neighbors which have been giving data to it at a good rate. (the “tit-for-tat” mechanism)
- The “rarest first replication” algorithm is key to the success of solving last block problem.

Related work

- Many analytical and simulation studies have shown that the existing BitTorrent algorithm is nearly optimal in terms of user experienced download time.

- However, all these studies assume that a peer's neighbors are selected randomly from the set all neighbors.

Problem

- All the studies assume that a peer's neigbors are selected randomly from the set all neighbors.
- However, is random neighbor selection just a sufficient condition for performance optimality; or is it also a necessary condition?

Biased neighbor selection

- In biased neighbor selection, a peer chooses most of its neighbors from the same ISP as itself, and only a few from other ISPs.

- Specifically, a parameter k is associated with the scheme, where for each peer, all but k neighbors are from the same ISP, and only k neighbors are chosen from outside the ISP.

Uniform random neighbor selection vs. biased neighbor selection

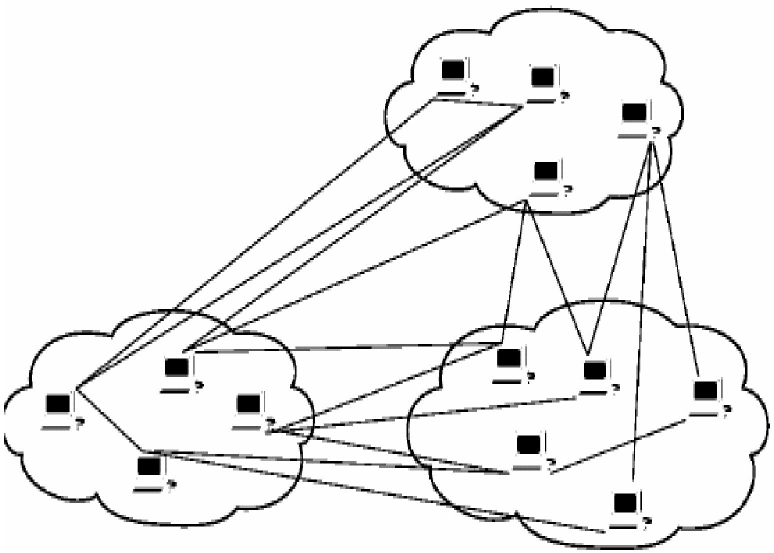

Uniform random neighbor selection

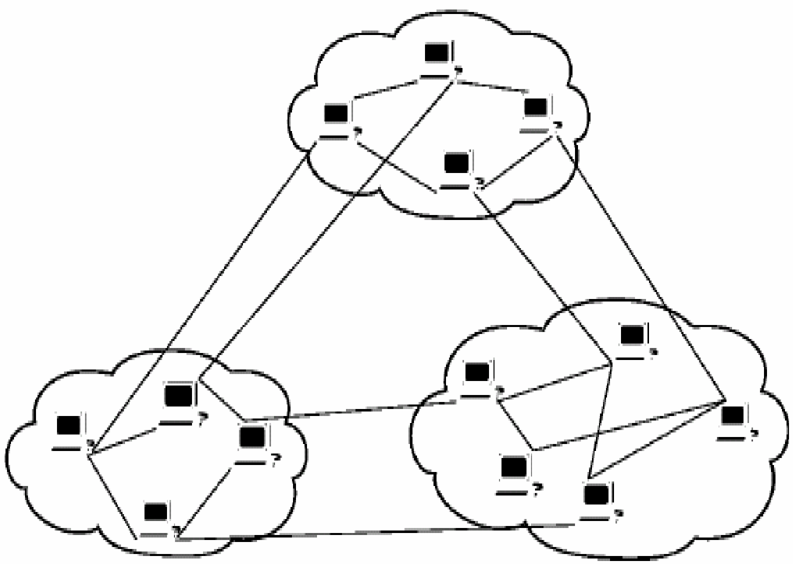

biased neighbor selection

Biased neighbor selection

- There are two ways to implement biased neighbor selection:
  - Modifying trackers and client
    - Biased neighbor selection can be implemented easily by changing the tracker and the client.
  - P2P traffic shaping devices
    - A P2P shaping device can implement biased neighbor selection without much obstacle.

Modifying trackers and client

- The tracker selects $35 - k$ internal peers and k external peers to hand back to the client.
- If there are less than $35 - k$ internal peers, the tracker also notifies the client to contact it again after a certain duration.

- The challenge here lies in informing the tracker of the ISP locality.

The tracker of the ISP locality

- The tracker can use Internet topology maps or IP to Autonomous System (AS) mappings to identify ISP boundaries.
- ISPs wishing to preserve traffic locality can also publish their IP address ranges to trackers.
- Finally, since BitTorrent tracker-client communication protocols run over HTTP, an ISP's HTTP proxy can append a new header "X-Topology-Locality" that contains a locality tag.

P2P traffic shaping devices

- When a peer joins the network for a file, the device intercepts and modifies the responses from the tracker to the peer, substituting outside peers with internal peers.

- When it is necessary to change a peer's neighbors, the device inserts TCP RESET on the connections between the peer and its external neighbors, forcing the peer to contact the tracker to obtain new neighbors.

Performance

- The experiments examine two network settings:
  - A homogeneous network consisting of 700 cable modem nodes spread among 14 ISPs.
  - A heterogeneous network consisting of a number of university nodes (have point-to-point links with each ISP and also with each other) and 700 cable modem nodes.
- In both cases, the original seed is a separate node whose bandwidth varies.

Biased neighbor selection (Modifying trackers and client) in homogeneous networks

Neighbor Selection	50^{th} percentile	95^{th} percentile	Traffic redundancy
Regular BitTorrent	5,312	7,152	46.9
Biased $k = 1$	5,168	6,206	3.04
Biased $k = 5$	5,172	6,281	9.74
Biased $k = 17$	5,220	5,872	21.38

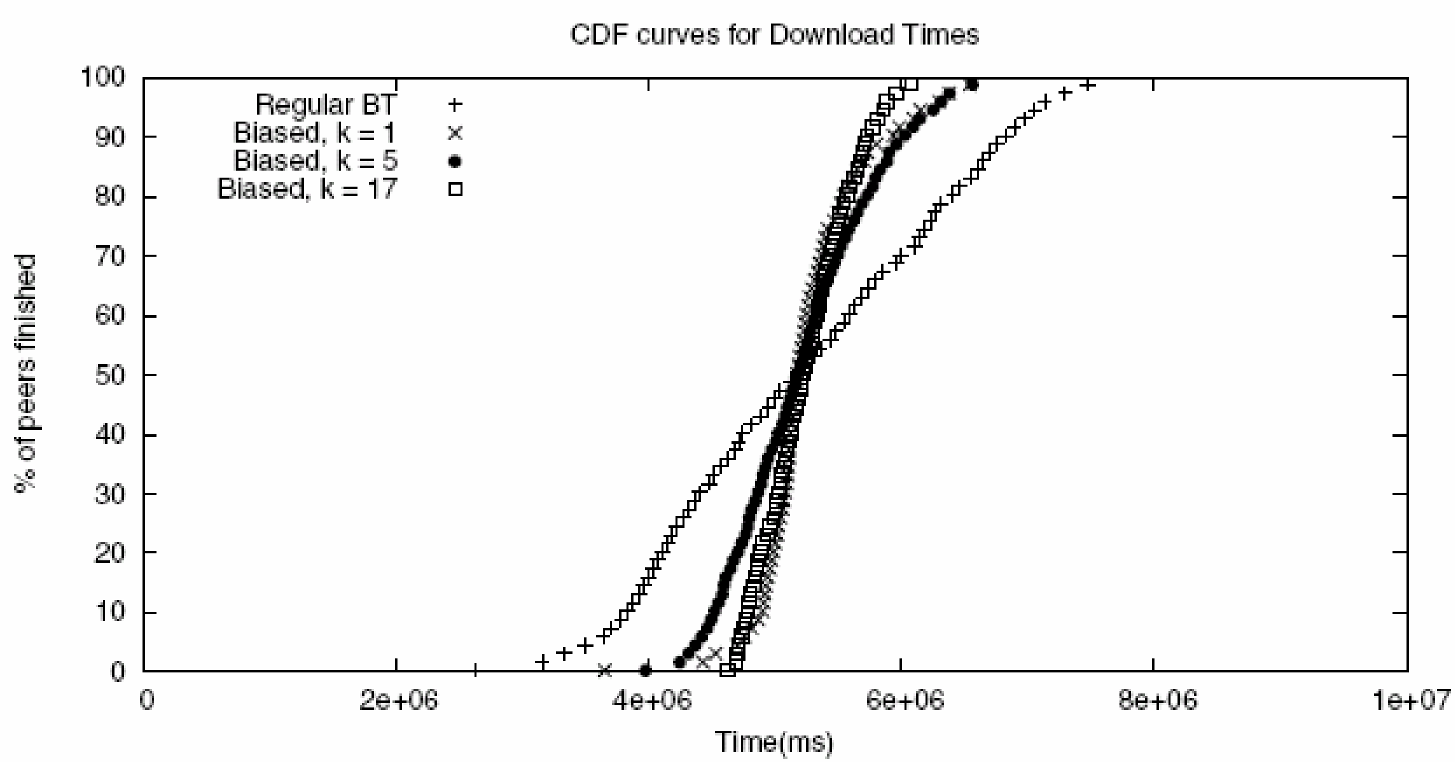

Biased neighbor selection (Modifying trackers and client) in heterogeneous networks

Extra University Nodes	Regular BitTorrent			Biased BitTorrent (k=1)		
	50^{th} percentile	95^{th} percentile	Traffic redundancy	50^{th} percentile	95^{th} percentile	Traffic redundancy
0	1.0	1.34	46.9	0.97	1.16	3.04
7	1.0	1.33	47.06	0.94	1.12	4.19
15	1.0	1.37	46.98	1.01	1.01	7.81
31	0.93	1.28	47.06	0.83	1.06	8.21

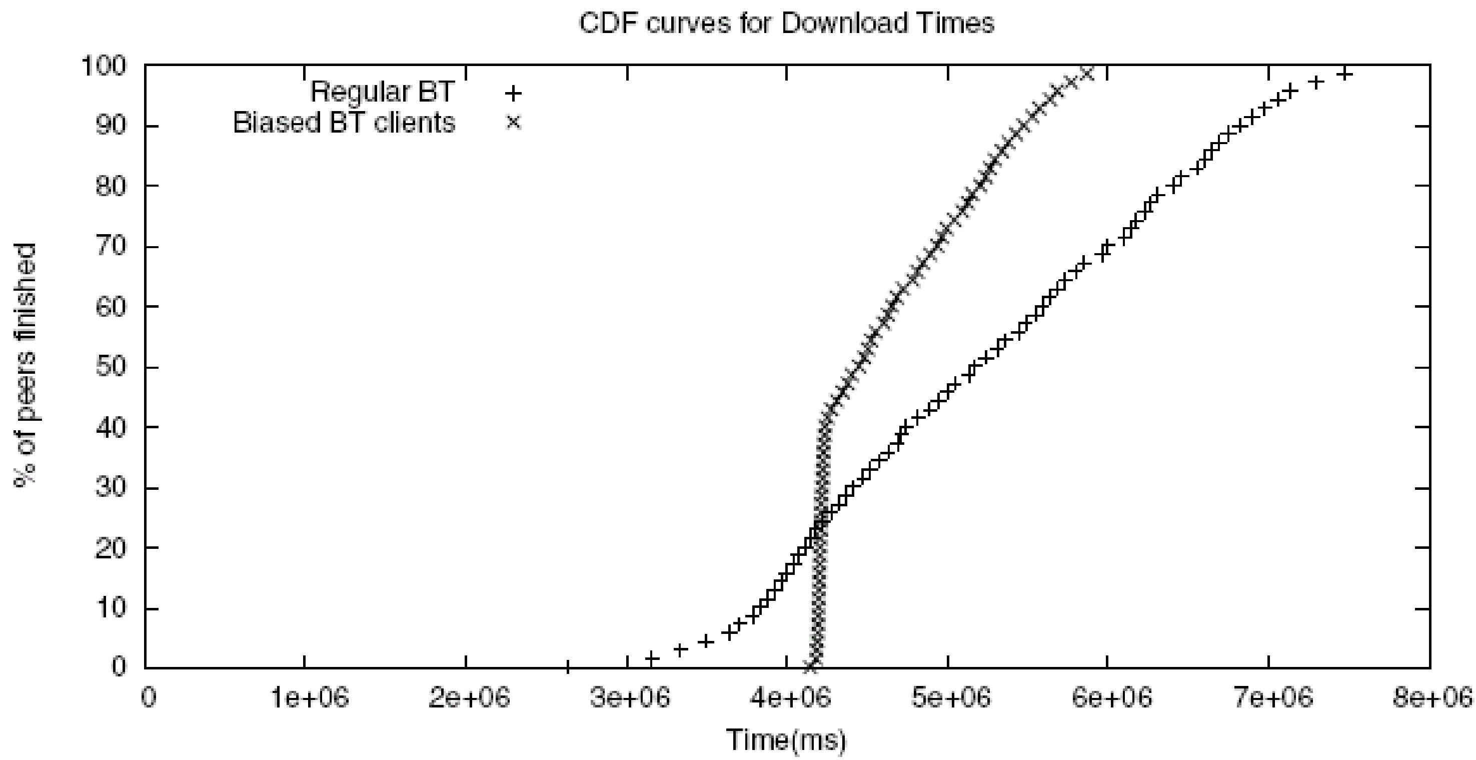

Biased neighbor selection (P2P traffic shaping devices) in homogeneous networks

ISP bottleneck	50^{th} percentile	95^{th} percentile	Traffic redundancy
no bottleneck	1.0	1.35	46.9
2.5Mbps	1.43	1.59	31.76
1.5Mbps	2.01	2.05	24.88
500Kbps	3.33	3.53	21.65

Normalized download time and traffic redundancy under bandwidth throttling.

Piece Selection	50^{th} percentile	95^{th} percentile	Traffic redundancy
Random	1.84	2.51	14.4
Rarest first	1.0	1.20	3.04

Effect of piece selection algorithms on biased BitTorrent. The download time of 1.0 means 5,168 seconds.

Discussion

- However, this method assumes that all peers in this network download under the static state.
- If the method was implemented under the dynamic state, what result will be?

Conclusion

- This main conclusion is that biased neighbor selection, in which a peer chooses the majority, but not all, of its neighbors from peers within the same ISP, can reduce cross-ISP traffic significantly while keeping the download performance nearly optimal.