Improving ISP Locality in BitTorrent Traffic via Biased Neighbor Selection

Ruchir Bindal, Pei Cao, William Chan

Stanford University

Jan Medved, George Suwala, Tony Bates, Amy Zhang

Cisco Systems, Inc.

P2P and ISPs: Not Friends

- P2P applications are notoriously difficult to "traffic engineer"
  - ISPs: different links have different monetary costs
  - P2P applications:
    - Peers are all equal
    - Choices made based on measured performance
    - No regards for underlying ISP topology or preferences

P2P and ISPs: Can't Be Foes

- ISPs: need P2P for customers
- P2P: need ISPs for bandwidth

- Current state of affairs: a clumsy co-existence
  - ISPs "throttle" P2P traffic along high-cost links
  - Users suffer

Can They Be Partners?

- ISPs inform P2P applications of its preferences
- P2P applications schedule traffic in ways that benefit both Users and ISPs

→ This paper gives an example for BitTorrent

Outline

- Review of BitTorrent
- Biased Neighbor Selection:
  - Design and Implementations
  - Evaluations
- Comparison with Alternatives

BitTorrent File Sharing Network

Goal: replicate K chunks of data among N nodes

- Form neighbor connection graph
- Neighbors exchange data

BitTorrent: Neighbor Selection

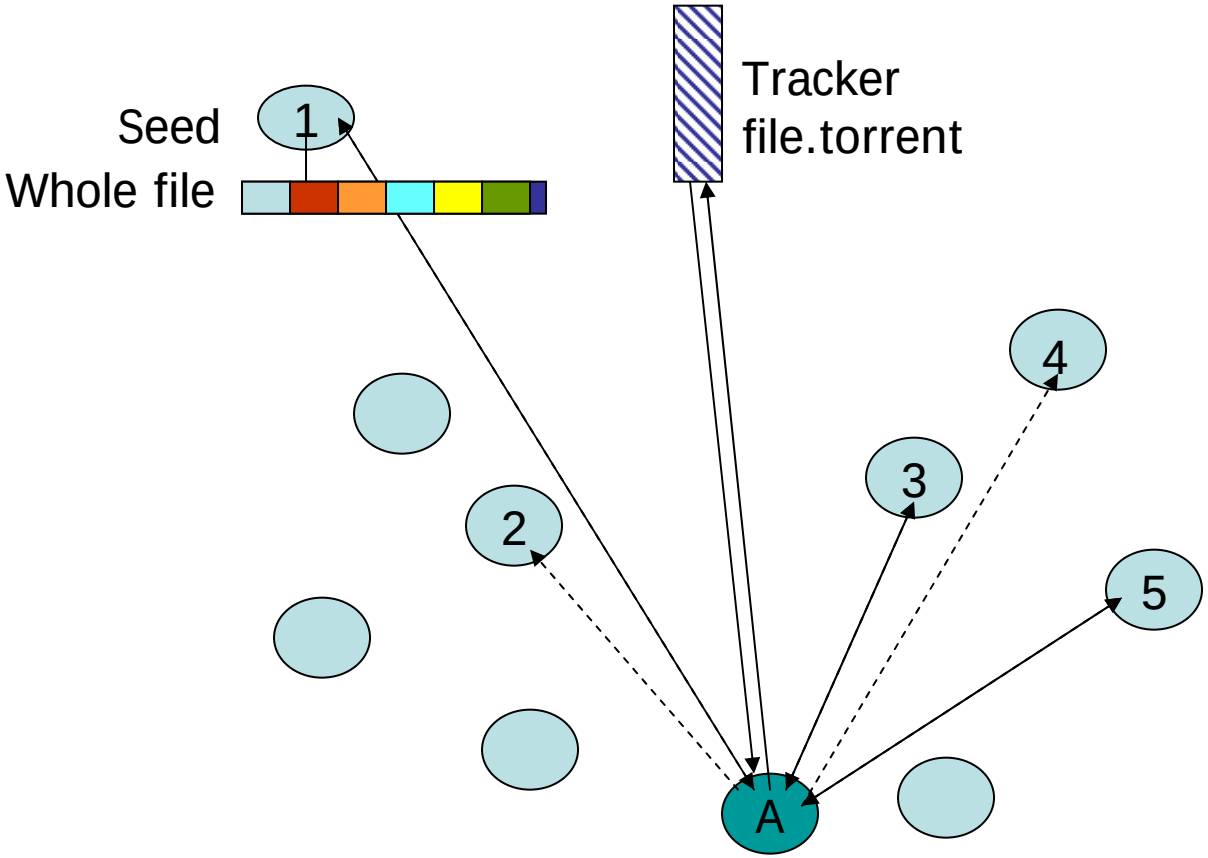

BitTorrent: Piece Replication

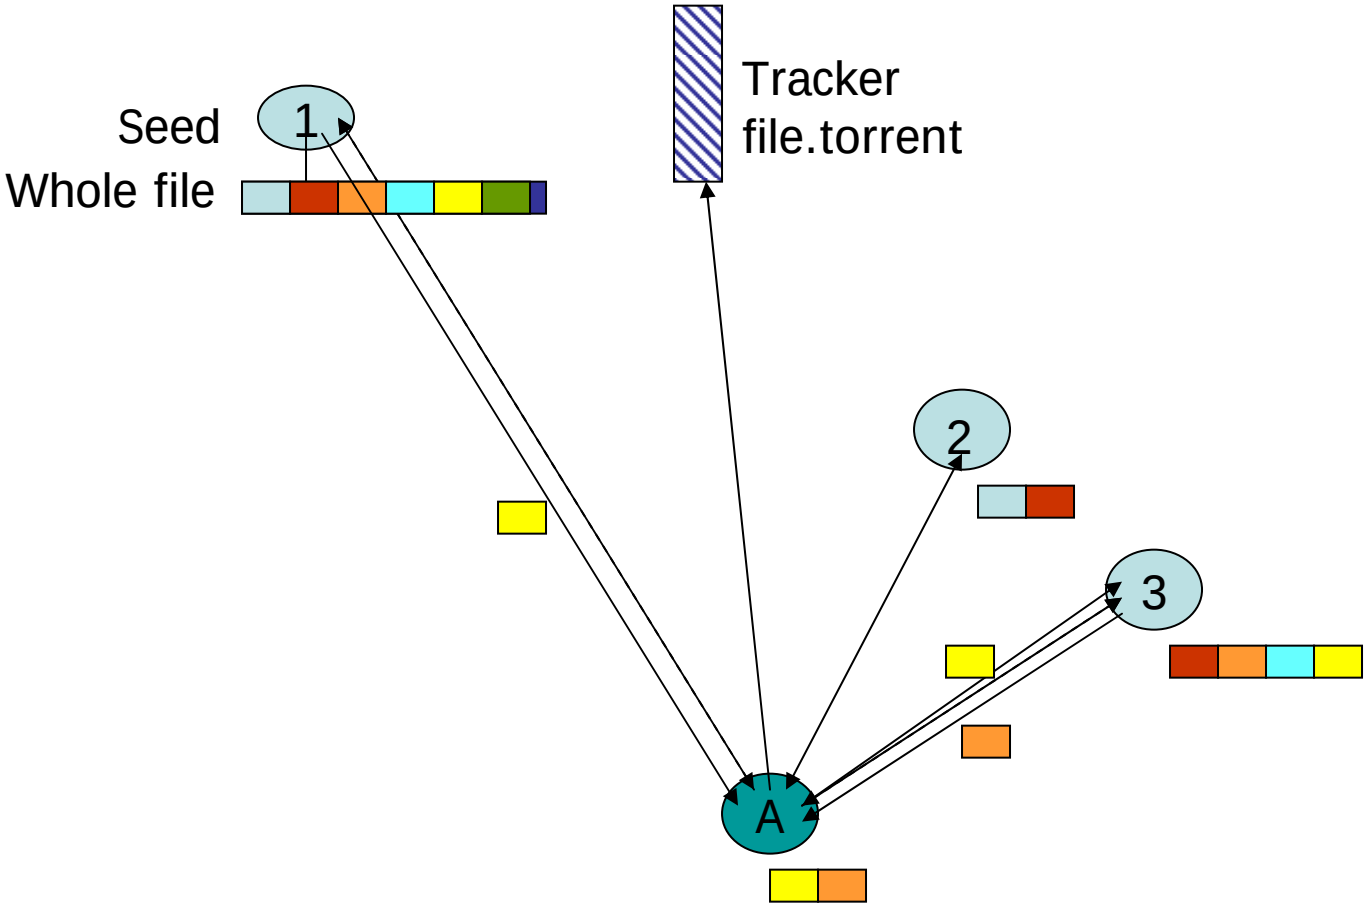

BitTorrent: Piece Replication Algorithms

- “Tit-for-tat” (choking/unchoking):
  - Each peer only uploads to 7 other peers at a time
  - 6 of these are chosen based on amount of data received from the neighbor in the last 20 seconds
  - The last one is chosen randomly, with a 75% bias toward new comers
- (Local) Rarest-first replication:
  - When peer 3 unchokes peer A, A selects which piece to download

Performance of BitTorrent

- Conclusion from modeling studies: BitTorrent is nearly optimal in idealized, homogeneous networks
  - Demonstrated by simulation studies
  - Confirmed by theoretical modeling studies
    - Intuition: in a random graph,

Prob(Peer A's content is a subset of Peer B's) ≤ 50%

Random Neighbor Selection

- Existing studies all assume random neighbor selection
  - BitTorrent no longer optimal if nodes in the same ISP only connect to each other
- Random neighbor selection → high cross-ISP traffic

Q: Can we modify the neighbor selection scheme without affecting performance?

Biased Neighbor Selection

- Idea: of N neighbors, choose N-k from peers in the same ISP, and choose k randomly from peers outside the ISP

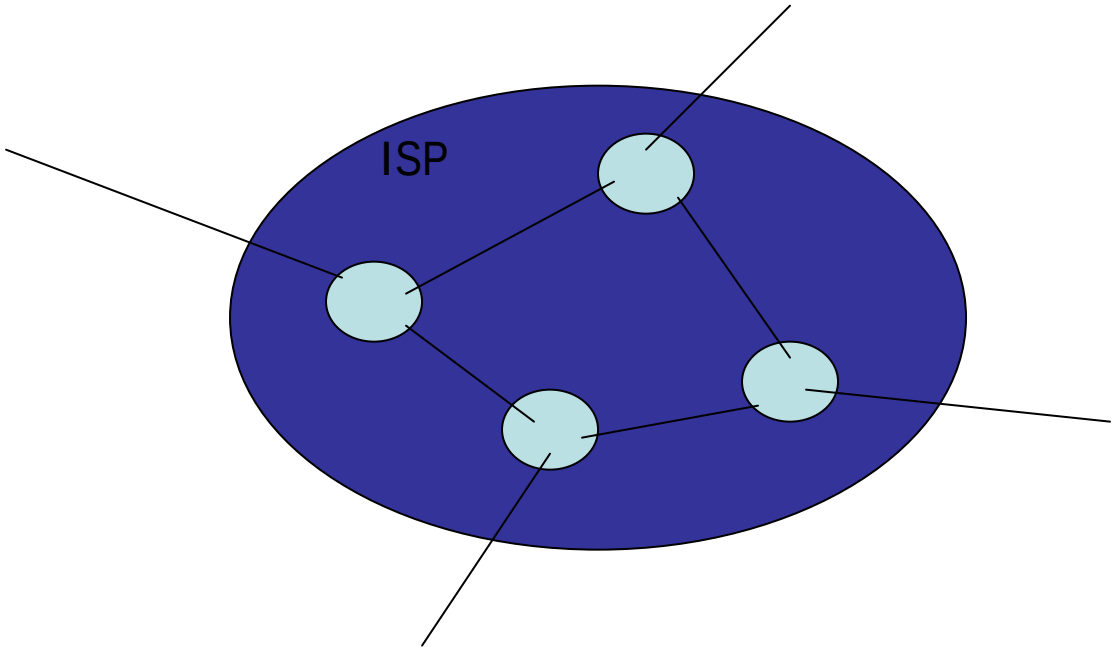

Implementing Biased Neighbor Selection

- By Tracker
  - Need ISP affiliations of peers
    - Peer to AS maps
    - Public IP address ranges from ISPs
    - Special “X-” HTTP header
- By traffic shaping devices
  - Intercept “peer → tracker” messages and manipulate responses
  - No need to change tracker or client

Evaluation Methodology

- Event-driven simulator
  - Use actual client and tracker codes as much as possible
  - Calculate bandwidth contention, assume perfect fair-share from TCP
- Network settings
  - 14 ISPs, each with 50 peers, 100Kb/s upload, 1Mb/s download
  - Seed node, 400Kb/s upload
  - Optional “university” nodes (1Mb/s upload)
  - Optional ISP bottleneck to other ISPs

Limitation of Throttling

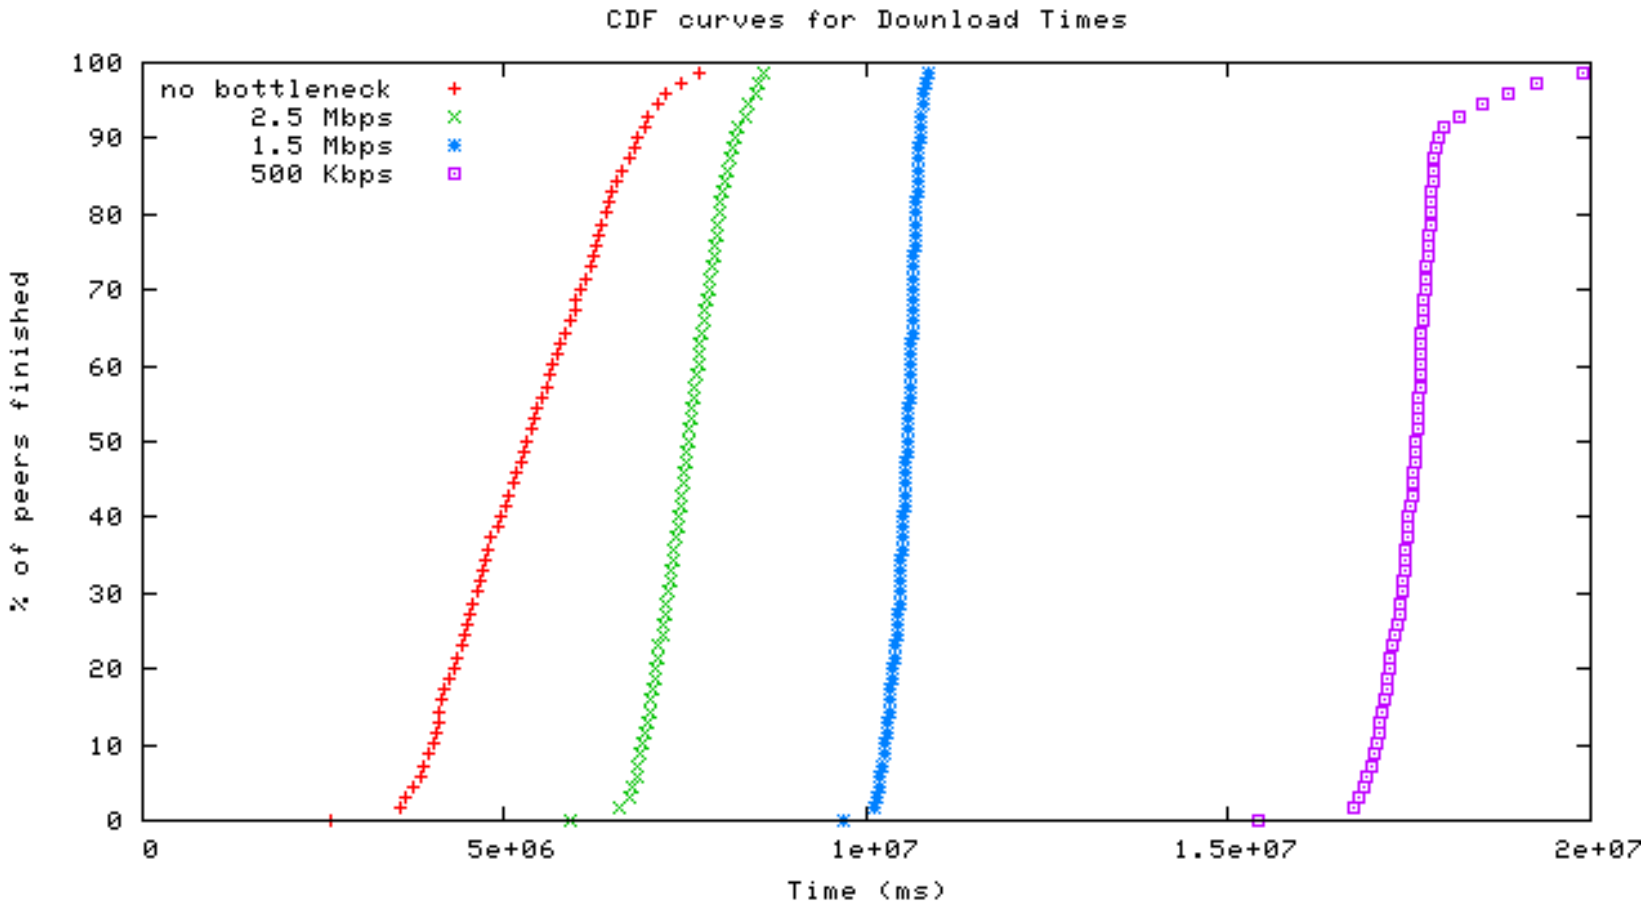

Throttling: Cross-ISP Traffic

Redundancy: Average # of times a data chunk enters the ISP

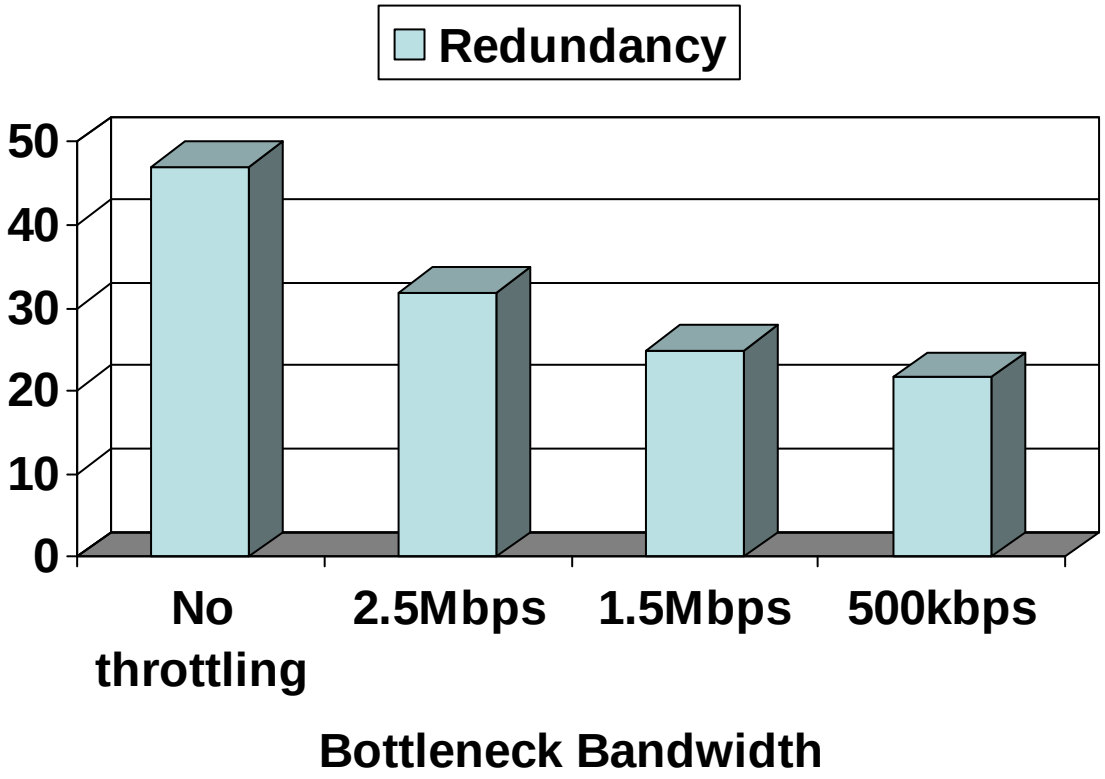

Biased Neighbor Selection: Download Times

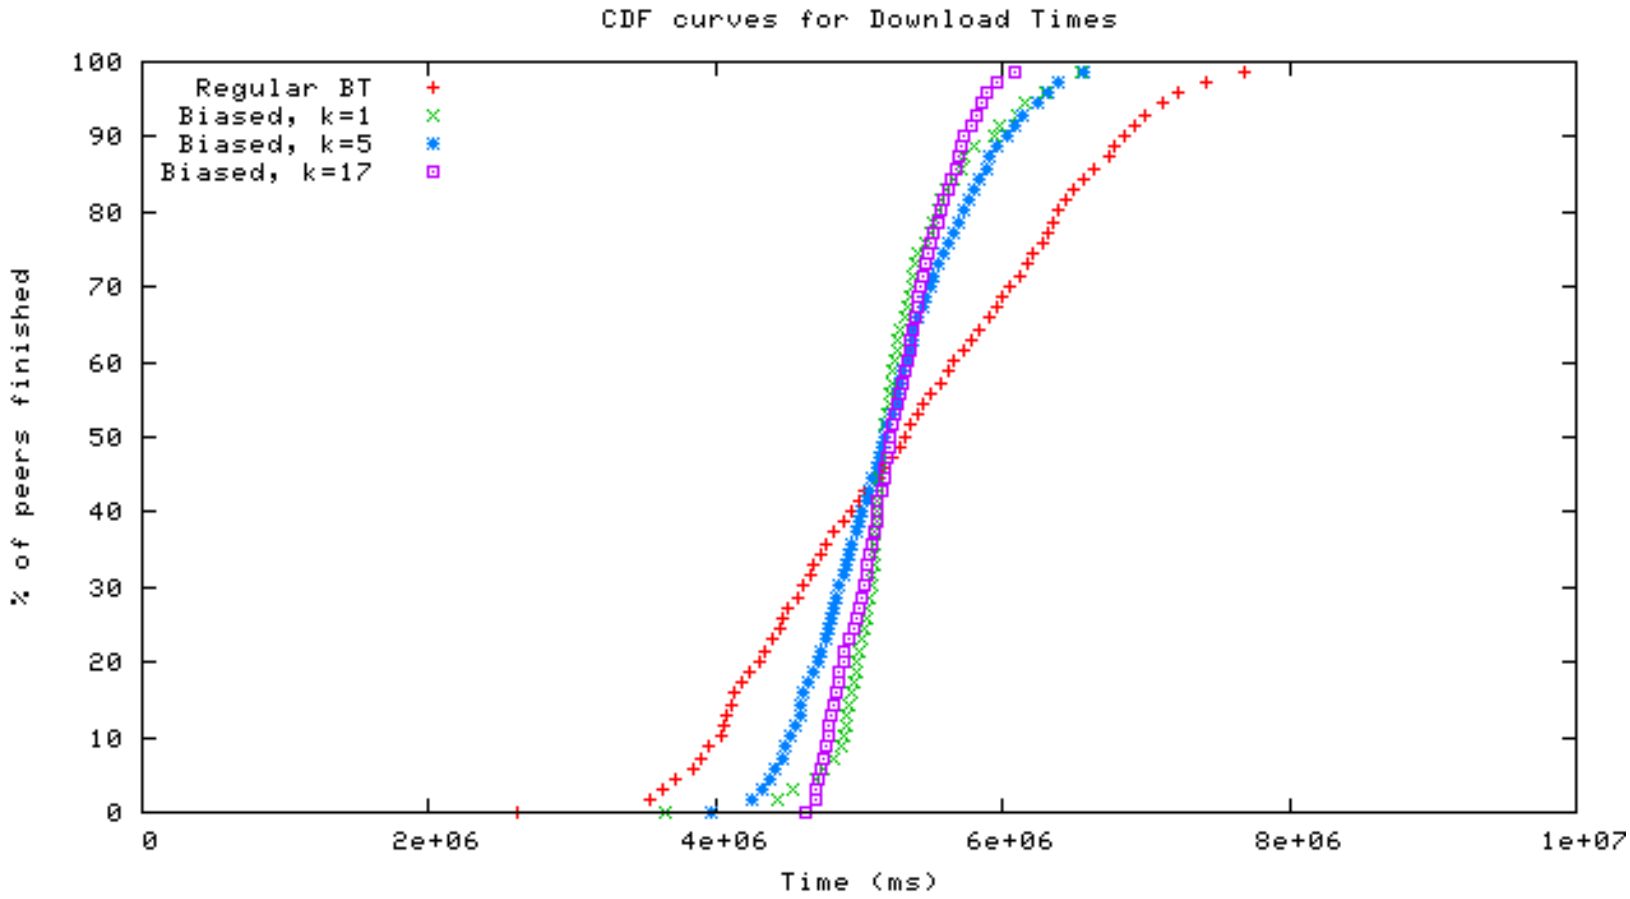

Biased Neighbor Selection: Cross-ISP Traffic

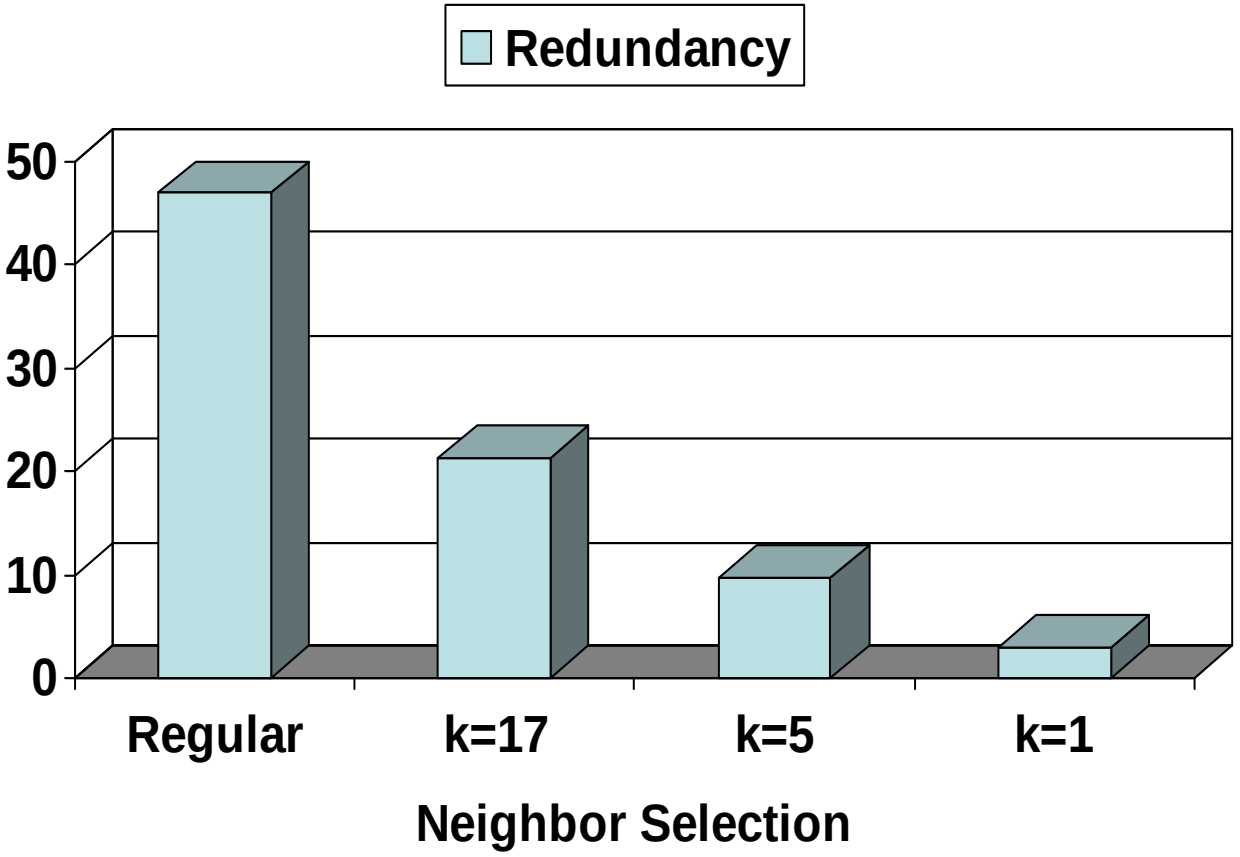

Importance of Rarest-First Replication

- Random piece replication performs badly
  - Increases download time by 84% - 150%
  - Increase traffic redundancy from 3 to 14
- Biased neighbors + Rarest-First → More uniform progress of peers

Biased Neighbor Selection: Single-ISP Deployment

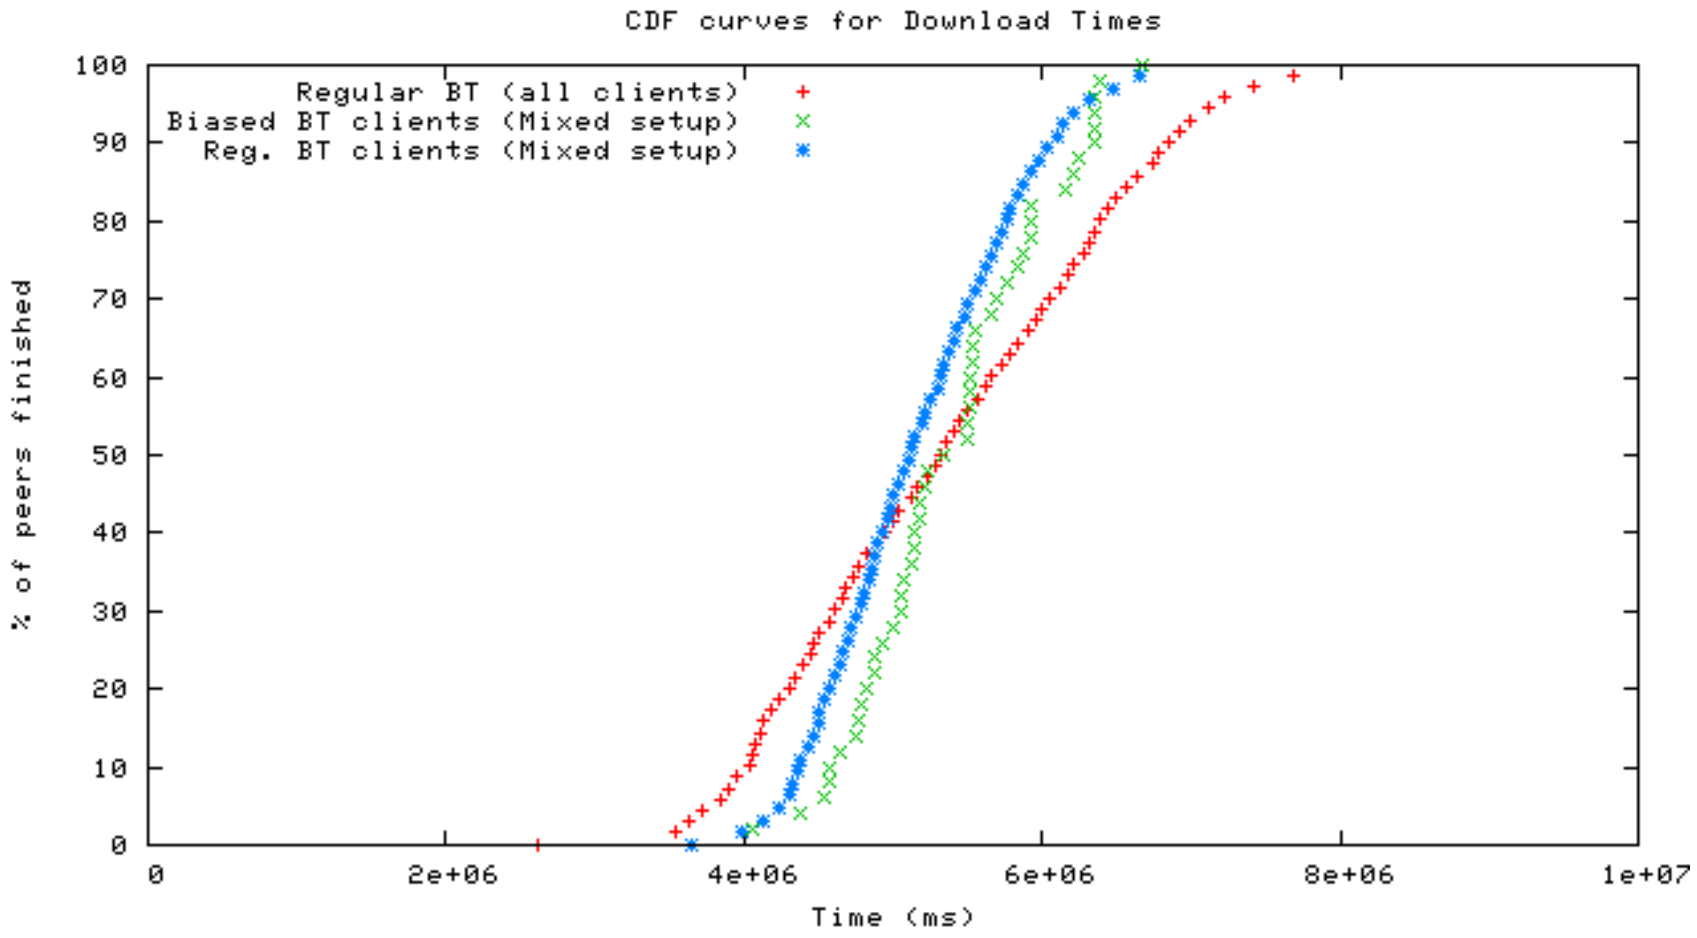

Presence of External High-Bandwidth Peers

- Biased neighbor selection alone:
  - Average download time same as regular BitTorrent
  - Cross-ISP traffic increases as # of “university” peers increase
    - Result of tit-for-tat
- Biased neighbor selection + Throttling:
  - Download time only increases by 12%
    - Most neighbors do not cross the bottleneck
  - Traffic redundancy (i.e. cross-ISP traffic) same as the scenario without “university” peers

Comparison with Alternatives

- Gateway peer: only one peer connects to the peers outside the ISP
  - Gateway peer must have high bandwidth
    - It is the “seed” for this ISP
  - Ends up benefiting peers in other ISPs
- Caching:
  - Can be combined with biased neighbor selection
  - Biased neighbor selection reduces the bandwidth needed from the cache by an order of magnitude

Summary

- By choosing neighbors well, BitTorrent can achieve high peer performance without increasing ISP cost
  - Biased neighbor selection: choose initial set of neighbors well
  - Can be combined with throttling and caching

→ P2P and ISPs can collaborate!

Related Work

- Many modeling studies of BitTorrent
- Simulation studies
- Measurements of real torrents

Future Work

- Implementation of tracker-side changes and experiments
- Theoretical modeling of biased neighbor selection
- Dynamic biased neighbor selection for “global congestion avoidance”