

Bandwidth Allocation in P2P File Sharing Networks

Creus-Mir, Casadesus-Masanell, Hervas-Drane

Discussion by Benjamin Chiao @ U. of Michigan. Apr 20, 2007

Empirical background: Freeriding in P2P

- ▶ Data from Gnutella
 - ▶ 66% of users do not contribute; 1% is responsible for over 50% shared (Adar and Huberman 2000)
 - ▶ 7% of the users share more than all of the other users combined (Saroiu 2002)
- ▶ A key research direction in the literature: How to allocate bandwidth to sharers and freeriders

Legal background: Freeriding in P2P

- ▶ Recording Industry Association of America (RIAA) had announced that
 - ▶ it is targeting to sue people who share over 800 files
 - ▶ already sued 8,000 file sharers
 - ▶ heavy-handed tactics: sued 83-year-old dead grandmother and 12-year-old girl
 - ▶ statutory damages (up to \$150K per infringement), a settlement in the \$4K range sounds like a "bargain"

Potential good news from this paper

- ▶ Conditional on a fixed supply of files, you don't really have to share because your download bandwidth won't increase much if you free ride instead.

Key Results

Expected benchmark bandwidth obtained by each peer \simeq :

$$\frac{\text{No. of sharers}}{\text{No. of sharers} + \text{No. of freeriders}} \in [0, 1]$$

More precisely:

Sharer's $>$ benchmark

Freerider's $<$ benchmark

The bandwidth differences decreases as network grows.

Comment 1: Can the intuition be further strengthened? Elaborate more on “the more constrained peers are, the better off they end up...[because it] is somewhat equivalent to allocating [bandwidth to] them first...”?

Methodology

Use Monte Carlo method to simulate the bandwidth obtained by sharers and freeriders.

The simulated bandwidth differences show that they are a series of peaks, with peaks decreasing as network grows.

Comment 2: In plot 5, only the LAST peak is shown to decrease with network size while plot 7 simulates the decrease of ALL peaks. Why? Is this valid comparison?

Assumptions and Biases

- ▶ Sharers give bandwidth equally to all peers connected to her

Potential Bias:

- ▶ Things not captured
 - ▶ Reciprocity (e.g. generalized reciprocity by Jian and MacKie-Mason 2006)
 - ▶ Tit for tat (e.g. choking in BitTorrent disallows other users from downloading)
 - ▶ Reputation:
 - ▶ Decentralized Credit System (e.g. eDonkey: Pairwise reputation)
 - ▶ Centralized Credit System (e.g. Kazaa: Score aggregated over all pairs)

Assumptions and Biases

- ▶ Homogeneity in upload and download capacity.

Potential Bias:

- ▶ Ignore preferable peers that have higher upload capacity
 - ▶ Or even higher download capacity (Krishnan et al's potluck/offloading effect 2004).
- ▶ Possible confirmation of heterogeneity:
 - ▶ 66% of Gnutella users do not contribute; 1% contributes over 50% (Adar and Huberman 2000)
 - ▶ 7% of the users share more than all of the other users combined on Gnutella (Saroiu 2002)

Assumptions and Biases

- ▶ Exogeneity of the roles of sharers and freeriders

Potential Bias:

Freeriding is not a decision variable.

But public goods experiments have shown that provisions could more easily breakdown as group size increases.

Then is it too much to hold exogeneity constant while varying network size?

Assumptions and Biases

- ▶ Peers can connect only to one sharer
 - ▶ Authors defense 1: “in real networks...simultaneous connections is generally limited”. 2: “complexity...raises substantially”.

Potential Bias:

In real networks, one can use multiple applications (BitTorrent, etc.) to obtain different segments of a file simultaneously.

Possible Follow-up Paper

An empirical paper that measures the bandwidth obtained in real networks of sharers and freeriders to solve a big question mark:

- ▶ In the simulation, when the network size is 10, the bandwidth difference between sharers and freeriders is at most 10^{-4} . Decrease to 10^{-6} when size is 100.
- ▶ But in the real world, we often observed that bandwidth obtained increases significantly if one shares. Will this simple and nice formula becomes simple, nice and REASONABLE?

$$\text{bandwidth for each peer} \simeq \frac{\text{No. of sharers}}{\text{No. of sharers} + \text{No. of freeriders}}$$

Thank you!