

Academic Torrents: Scalable Data Distribution

Henry Z. Lo*, Joseph Paul Cohen*

October 10, 2015

1 Introduction

As competitions get more popular, transferring ever-larger data sets becomes infeasible and costly. For example, downloading the 157.3 GB 2012 ImageNet data set incurs about \$4.33 in bandwidth costs per download. Downloading the full ImageNet data set takes 33 days. ImageNet has since become popular beyond the competition, and many papers and models now revolve around this data set. For sharing such an important resource to the machine learning community, the sharers of ImageNet must shoulder a large bandwidth burden.

Academic Torrents reduces this burden for disseminating competition data, and also increases download speeds for end users ¹. By augmenting an existing HTTP server with a peer-to-peer swarm, requests get re-routed to get data from downloaders. While existing systems slow down with more users, the benefits of Academic Torrents grow, with noticeable effects even when only one other person is downloading.

*All authors contributed equally.

¹<http://academictorrents.com>. Academic Torrents is run by a pending nonprofit.

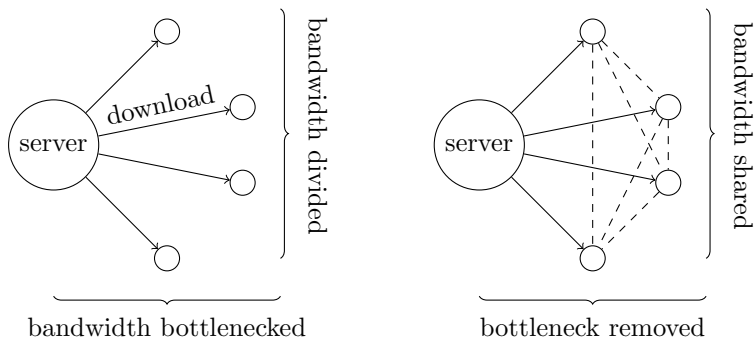


Figure 1: HTTP (client-server) data sharing vs. HTTP + peer-to-peer. Peer-to-peer swarms provide more data sources for downloaders, reducing server load.

Challenge	Upload Bandwidth (100 DLs)			Download Speed		
	HTTP	AT	Savings	HTTP	AT	Savings
Whale ⁴	873.00 GB	20.68 GB	\$23.36	4.85 h	0.07 m	4.78 h
Diabetes ⁵	8.22 TB	0.20 TB	\$220.68	45.66 h	0.67 m	44.99 h
ImageNet	15.73 TB	0.37 TB	\$422.29	87.39 h	1.28 h	86.11 h

Table 1: Cost savings (for uploading) and downloading using Academic Torrents. Values are projections based on actual measures from the Reddit public comments data set. Upload calculations are for 100 downloads. Time is based on 34 MB/s download speed.

2 Case Study

Reddit public comments² is a 160.68 GB data set (when compressed) of text currently linked to on Kaggle. Since being shared on Academic Torrents in May 2015, the original seeder has uploaded 366.68 GB, while the entire community has downloaded 15.43 TB. The upload/download ratio is:

$$U/D = \frac{\text{366.68 GB}}{\text{Sharer uploaded}} / \frac{\text{15.43 TB}}{\text{Total downloaded}} = \mathbf{42.067} \quad (1)$$

For every byte uploaded, the community matched that contribution 41 times. In a standard HTTP upload/download system, sharing the same amount of data would incur the seeder 42.067 times more bandwidth.

This U/D ratio translates to real costs. To the uploader, the Reddit data set costs \$4.42 for each download³. For the 96 downloads, the associated HTTP bandwidth would mean \$424.32, while the bandwidth used on Academic Torrents reduces the bill to \$10.09.

Downloaders also get their data faster. From our university server, we are currently downloading the entire 1.2 TB ImageNet data at 500 KB/s, resulting in a 33-day download. On Academic Torrents, we have previously reached speeds of 34 MB/s; a download at this speed would thus finish in 9.8 hours. This download speed is limited only by the bandwidth of the pipe.

Using this data, we cost-project the benefits of Academic Torrents with the same U/D ratio and speeds that we saw on our previous downloads (Table 1).

3 Future

Academic Torrents is maintained by the pending-nonprofit Institute for Reproducible Research. Support for terms of service for data sets was recently added. Academic Torrents currently indexes 10.12 TB of data, serves 30000 unique visitors a month, and facilitates over 900 GB a day. Future efforts will scale up the platform and improve its user interface to make it more accessible.

²https://www.reddit.com/r/datasets/comments/3bxlg7/i_have_every_publicly_available_reddit_comment/

³Assuming US Amazon S3 data transfer costs of \$0.0275 per GB.