

The BitTorrent File-Sharing System: Measurements and Analysis

J.A. Pouwelse, P. Garbacki, D.H.J. Epema, H.J. Sips

International Workshop on Peer-to-Peer Systems. Springer, Berlin, Heidelberg, 2005

Presented by Steven Smith - CSC 7970

Outline

- ▶ Overview of Study
- ▶ Introduction
- ▶ BitTorrent System - BitTorrent/Suprnova
- ▶ Experimental Setup
- ▶ Results
- ▶ Discussion and Conclusion
- ▶ Critique and Issues
- ▶ Take Aways

Overview

- ▶ Study Focused On Important Aspects That Lead to Popular Adoption of BitTorrent:
 - ▶ Availability
 - ▶ Integrity
 - ▶ Handling of Flashcrowds
 - ▶ Download Performance
- ▶ Considers BitTorrent and Suprnova.org as a system.
- ▶ Study took place over 8 months - June 2003-March 2004.
- ▶ Considered over 2 thousand global components.

Introduction - Terminology

- ▶ Flashcrowding - The effect caused by a sudden huge increase of peers for a popular new file.
- ▶ Peers - Clients downloading a file and distributing portions they have downloaded.
- ▶ Seeds - Clients with the full file who simply upload to other peers.
- ▶ Barter - The tit-for-tat process peers go through to negotiate downloading and uploading chunks of files that is meant to prevent parasitic behavior.
- ▶ Injector - User who uploads .torrent file and provides the first seed for the file.
- ▶ Pollution - The presence of fake, malicious, or corrupted files.

Introduction - Popular BitTorrent Websites

Site name	Available files	File transfers
Suprnova.org	46,766	2,267,463
Youceff.com	47,137	1,145,889
Piratebay.org	39,294	749,133
Lokitorrent.com	30,957	816,435

Table 1: Popular BitTorrent web sites (Oct 2004).

Introduction - Listed Contributions

- ▶ Adding understanding of the operation of the BitTorrent system and the reasons why it is able to attract millions of users.
- ▶ The results of the paper can help in the mathematical modeling of P2P systems.
- ▶ Findings about the conflict between data integrity and availability when comparing P2P systems with centralized global components vs. those without centralization.

BitTorrent System - BitTorrent Protocol

- ▶ P2P file-sharing protocol - Relies on other components such as websites for users to find files.
- ▶ Peers download via bartering with other peers/seeders.
- ▶ Peers are responsible for maximizing their download rate through selection of suitable peers to download from.
- ▶ Peers become seeds by staying online to distribute files after download.
- ▶ Torrent files are generally found through links on websites
 - ▶ Example: Suprnova.org.
- ▶ Torrent file contents - Metadata and Tracker info.
- ▶ Made up 53% of all P2P traffic in June 2004.

BitTorrent System - Suprnova.org

- ▶ Most popular site for BitTorrent files at the time.
- ▶ Torrent files are downloaded from one of several torrent file servers for load balancing.
 - ▶ These files are not stored on Suprnova or its mirrors.
- ▶ Mirroring System - Used for load balancing of user requests and improving availability.
- ▶ Content Moderation - 2 Levels of Users/Injectors:
 - ▶ Moderated Submitter - Content must first be inspected by moderators.
 - ▶ Unmoderated Submitters - Trusted users may upload without their content being first checked by moderators.
- ▶ Moderators - Trusted users who inspect files, unmoderated submitters may request a promotion to become moderators.

BitTorrent System - Combination of BitTorrent and Suprnova

- ▶ Mirroring for load balancing of website.
- ▶ Meta-data Distribution to different file servers for load balancing.
- ▶ Bartering for fair resource sharing.
- ▶ P2P moderation and meta-data for integrity.

Global Components in BitTorrent/Suprnova System

- ▶ Main Suprnova.org server
- ▶ Suprnova.org Mirrors
- ▶ .torrent file servers and mirrors
- ▶ BitTorrent Trackers
- ▶ All considered to be centralized.

Experimental Setup - Monitoring Global BitTorrent/Suprnova Components

- ▶ Mirrorscript - Measures availability and response time of Suprnova mirrors.
- ▶ HTMLScript - Parses HTML pages of Suprnova to gather new .torrent files every hour.
- ▶ TrackerScript - Parses .torrent files for new trackers to add to a list, and checks the status of known trackers.
- ▶ Goal: Measures uptime of Suprnova mirrors and BitTorrent trackers. Also gather new .torrent files for use in monitoring of peers.

Experimental Setup - Monitoring Peers

- ▶ Made use of 100 Nodes of Distributed ASCII Supercomputer (DAS).
- ▶ HuntScript - Selects a file to follow and initiates a measurement of peers downloading files. Monitors once per minute for new .torrent files.
- ▶ GetPeerScript - Contacts tracker to get IP Addresses of users downloading and seeding the file. Activated after HuntScript selects a file.
- ▶ PeerPingScript - Contacts numerous peers in a parallel fashion through the use of the BitTorrent protocol to obtain download progress and uptime. Also activated after HuntScript selects a file.
- ▶ Goal: Obtain IP addresses of peers that inject new content and get an estimate of average download speeds.

Experimental Setup - Problems

- ▶ Firewalls
 - ▶ Widespread use prevented the PeerPingScript from measuring download speeds. Could only measure download speeds for peers that were not blocked by firewalls.
- ▶ Inability to get all peer IP Addresses
 - ▶ BitTorrent protocol only allows a tracker to return a limited number of randomly selected peers.
- ▶ Peer coverage
 - ▶ Fraction of peers discovered and for which IP Addresses were found. Still managed around 95% in all measurements.
- ▶ Modifications to BitTorrent
 - ▶ Created gaps in traces.

Results - Availability




Figure 1: The number of users downloading or seeding on BitTorrent/Suprnova for one month (Dec'03-Jan'04).

Results - Availability

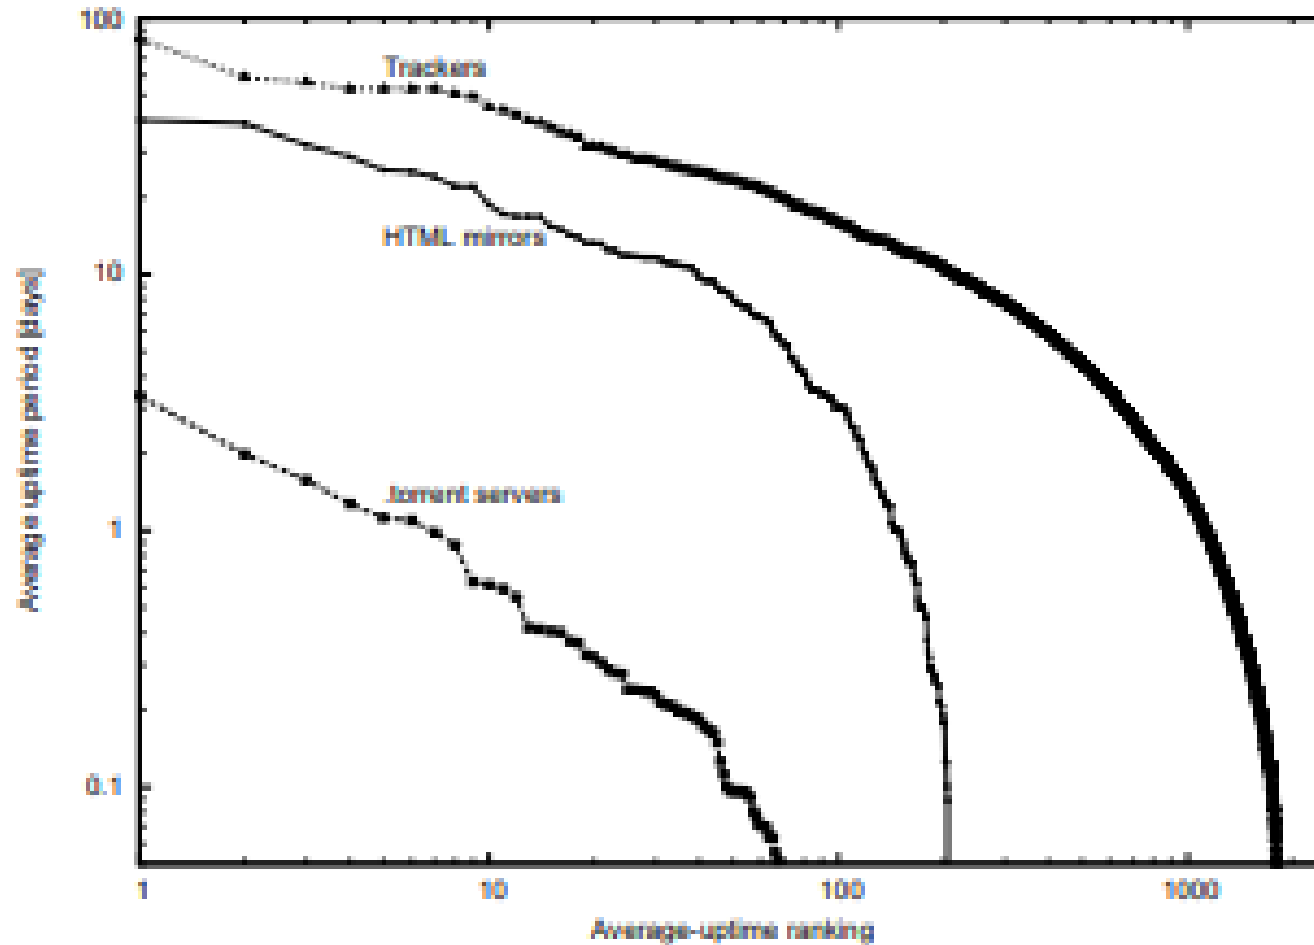


Figure 2: The uptime ranking of three types of BitTorrent/Suprnova global components.

Results - Availability




Figure 3: The uptime ranking of the 53,833 peers downloading "Beyond Good and Evil".

Results - Availability Conclusions

- ▶ Failures in Suprnova mirrors, .torrent file servers, and trackers led to large variation in measurements.
 - ▶ Main supernova.org servers - Often switched IP address or was down.
 - ▶ Suprnova mirrors - Rarely survived more than a few days due to having over 1,200,000 visitors.
 - ▶ Torrent file servers - Sometimes none were available.
 - ▶ Trackers - Frequent targets for DoS attacks. Costly due to daily use of GBs of bandwidth.
- ▶ Number of active users strongly influenced by the availability of global components in BitTorrent and Suprnova.

Results - Availability Conclusions

- ▶ Reliable webhosting of Suprnova pages is a large issue.
- ▶ Trackers show good availability in comparison.
- ▶ Seeds with high availability are rare.
 - ▶ Only 9,219 out of 53,883 (17%) peers on one popular file had an uptime longer than 1 hour after finished downloading. This decreased to 0.34% after 100 hours.
 - ▶ Seeds should be given incentives to increase their uptimes.
- ▶ Unavailability has a significant effect on popularity.
- ▶ Potential need to decentralize components to improve availability
 - ▶ However, BitTorrent and Suprnova's popular features depend on centralized components.

Results - Integrity




Figure 4: The activity of the different content submitters on Suprnova to prevent pollution.

Results - Integrity Conclusions

- ▶ Donated mirror to Suprnova with spyware embedded in the site code to register each .torrent file and could have corrupted the meta-data.
 - ▶ Found the use of donated resources for hosting of meta-data could easily lead to integrity and privacy issues.
- ▶ BitTorrent/Suprnova is considered by many P2P users to be pollution free.
 - ▶ Tested this by actively trying to pollute the system with fake files.
 - ▶ Found direct measurement of corrupted files through manually checking to not be viable.
 - ▶ Content was filtered out by moderators.
- ▶ Found that moderators were effective in remove fake and corrupted files with only a few volunteers - 20 in the study.
 - ▶ However, the system cannot be easily distributed due to moderation requiring a centralized area for checking the files.

Results - Flashcrowds

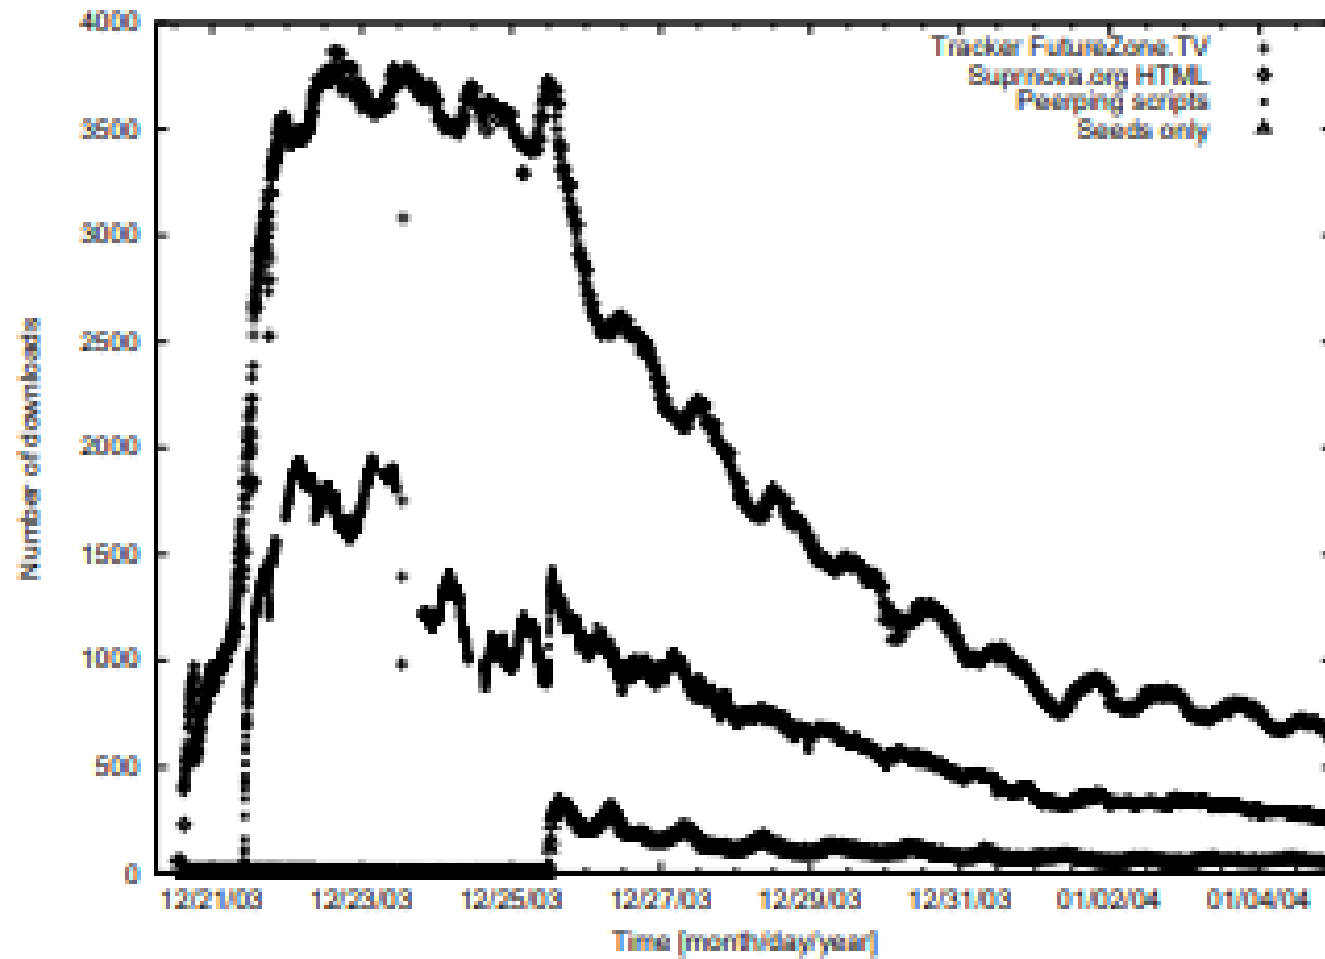


Figure 5: Flashcrowd effect of “Lord of the Rings III”.

Results - Flashcrowds Conclusions

- ▶ Futurezone.tv Tracker used as it provided detailed statistics.
- ▶ The long time period before peers finished downloading could provide an opportunity to identify users violating copyrights.
- ▶ Statistics from Suprnova were in agreement with the total tracker results.
- ▶ Peering script showed a 40% lower amount of downloads due to the firewall problem.
- ▶ Gaps in the Peering results were due to technical issues caused by disk quota problems.
- ▶ Ultimately, conclude that BitTorrent/Suprnova are capable of handling large flashcrowds efficiently.

Results - Download Performance





Figure 6: The average download speed of peers.

Results - Download Performance

- ▶ Huntscript followed first 108 files added to Suprnova on March 10, 2004.
- ▶ 54,845 peers measured over a time frame of 2 weeks.
- ▶ Cumulative Distribution Function (CDF) - Probability that a peer had a download speed with a value of less than or equal to 520 kbps.
 - ▶ $F(x) = (X \leq x)$ for all $x \in R$ [2]
 - ▶ 90% of peers were below 520 kbps.
- ▶ Average download speed of 240 kbps.
- ▶ Relation between average download speed and the number of downloads at that speed.

Results - Download Performance



Results - Download Performance

Conclusions

- ▶ High average download speed allows peers to fetch large files in one day.
- ▶ When popularity drops and the last peer/seed with chunks of files go offline, the content dies.
- ▶ Number of seeds after 10 days is not an accurate measure of content lifetime.
- ▶ BitTorrent implementations punish peers for seeding by making use of their maximum bandwidth capacity to seed to others.
 - ▶ Instead, peers should receive incentives to seed with the chance to pick the amount of bandwidth to use.



Related Work

- ▶ D. Qiu and R. Srikant proposed a fluid model to determine the average download time of a single file. [15]
 - ▶ Assumes Poisson arrival and departure processes, equal upload and download bandwidth for all peers, and no flashcrowd effect.
- ▶ Decentralization of BitTorrent tracker leads to easy corruption of files due to the lack of a central integrity checking mechanism. [18]
- ▶ Studies of Other P2P networks
 - ▶ Most availability studies do not show long term peer behavior due to short time frames. [4, 6, 8]
 - ▶ Integrity not well explored, one study found that up to 70% of popular songs on Kazaa were polluted. [12]
 - ▶ Found nearly 70% of peers in Gnutella networks did not contribute any bandwidth. [3]
 - ▶ Content lifetime not explored.

Discussion and Conclusions

- ▶ Authors published all raw, anonymized data files, measurement software, and documentation for public use.
- ▶ Found overall high levels of integrity in the BitTorrent/Suprnova system with a relatively low amount of moderators.
 - ▶ This comes at a price of centralization, which leads to availability issues.
- ▶ Decentralization would help availability, but make meta-data more vulnerable to corruption and integrity attacks. Needs future research.
 - ▶ Kazaa has many issues with fake files due to its decentralized nature.
- ▶ Creating incentives to seed is important.
 - ▶ Suggest giving seeds more preference for bartering with other files.

Critique

- ▶ Good structure overall.
 - ▶ Conclusion split into results sections made their findings easier to follow.
- ▶ Many issues with their experimental setup.
 - ▶ Few files measured in depth - Only two measured for availability and integrity. Could lead to biases in conclusions.
 - ▶ Issues with firewalls bring download speed results into question.
 - ▶ Claim to be able to corrupt meta-data on a mirror, but do not provide an example of actually doing so successfully.
 - ▶ For content lifetime, only considered lifetime of files over 10 days for files with at least 10 seeds.
- ▶ Do not present many solutions, only vague suggestions like offering incentives through giving preference to barter for other files.
- ▶ Not much info on .torrent file servers.

Potential Improvements

▶ Disk Quota

- ▶ Try to better estimate or measure the amount of users in a flashcrowd situation beforehand to ensure proper disk space is available.
- ▶ This issue calls results from Peering script into question.

▶ Time Frames

- ▶ Consider measuring data over longer periods of time such as months.
- ▶ Study took place over 8 months, but many measurements were only take over a period of days to weeks - Download speed, content lifetime, etc.

▶ Meta-data Corruption

- ▶ Actually corrupt meta-data on mirror and show results.
- ▶ Would users report this? Would moderators be able to respond and take the mirror down?

Potential Improvements

▶ Content Lifetime

- ▶ Consider content lifetime from another angle.
- ▶ With a persistent seed (provided by the study), how long would peers continue to download the file?
- ▶ Also, how long till a random sampling of popular files no longer have seeds or enough file chunks available to recreate the file.
- ▶ How often do files get revived after they no longer have seeds or the full file available?
- ▶ When they do, does the flashcrowd effect occur again, or does the peer count stay low.

▶ Trackers

- ▶ Measure the amount of downtime caused by DoS attacks vs. other reasons.

Potential Improvements

- ▶ Integrity of Files
 - ▶ Download random sampling of files, both popular and unpopular.
 - ▶ Determine if their content is true.
 - ▶ May be able to make a threshold to see if unpopular content tends to have more issues with integrity.

Takeaways

- ▶ BitTorrent is good for providing for file integrity, as long as the centralized components can be trusted.
- ▶ Should check assumptions before going too far in experimental setup.
- ▶ BitTorrent has many potential issues with availability.
- ▶ Seeds should be provided with good incentives to promote better availability.
- ▶ Decentralized integrity checking is a large issue, but could solve the availability issues if a good solution was found.

References

- ▶ [1] Pouwelse, J., Garbacki, P., Epema, D., & Sips, H. (2005, February). The bittorrent p2p file-sharing system: Measurements and analysis. In International Workshop on Peer-to-Peer Systems (pp. 205-216). Springer, Berlin, Heidelberg.
- ▶ [2] Introduction to Probability, Statistics, and Random Processes. Accessed on 2/27/2020. https://www.probabilitycourse.com/chapter3/3_2_1_cdf.php
- ▶ [3] E. Adar and B. A. Huberman. Free riding on gnutella. Technical report, Xerox PARC, August 2000.
- ▶ [4] R. Bhagwan, S. Savage, and G. M. Voelker. Understanding availability. In International Workshop on Peer to Peer Systems, Berkeley, CA, USA, February 2003.
- ▶ [5] S. Byers, L. Cranor, E. Cronin, D. Kormann, and P. McDaniel. Analysis of security vulnerabilities in the movie production and distribution process. In The 2003 ACM Workshop on DRM, Washington, DC, USA, October 2003.
- ▶ [6] J. Chu, K. Labonte, and B. Levine. Availability and locality measurements of peer-to-peer file systems. In ITCOM: Scalability and Traffic Control in IP Networks, Boston, MA, USA, July 2002.
- ▶ [7] B. Cohen. Incentives build robustness in bittorrent. In Workshop on Economics of Peer-to-Peer Systems, Berkeley, CA, USA, May 2003. <http://bittorrent.com/>.
- ▶ [8] K. Gummadi, R. Dunn, S. Saroiu, S. Gribble, H. Levy, and J. Zahorjan. Measurement, modeling, and analysis of a peer-to-peer file-sharing workload. In 19-th ACM Symposium on Operating Systems Principles, Bolton Landing, NY, USA, October 2003.

References

- ▶ [9] M. Izal, G. Urvoy-Keller, E. Biersack, P. Felber, A. Al Hamra, and L. Garces-Erice. Dissecting bittorrent: Five months in a torrent's lifetime. In *Passive and Active Measurements*, Antibes Juan-les-Pins, France, April 2004.
- ▶ [10] T. Karagiannis, A. Broido, N. Brownlee, kc claffy, and M. Faloutsos. Is p2p dying or just hiding? In *Globecom*, Dallas, TX, USA, November 2004.
- ▶ [11] N. Leibowitz, M. Ripeanu, and A. Wierzbicki. Deconstructing the kazaa network. In *3rd IEEE Workshop on Internet Applications (WIAPP'03)*, San Jose, CA, USA, June 2003.
- ▶ [12] J. Liang, R. Kumar, Y. Xi, and K. Ross. Pollution in p2p file sharing systems. In *IEEE Infocom*, Miami, FL, USA, March 2005.
- ▶ [13] T. Oh-ishi, K. Sakai, T. Iwata, and A. Kurokawa. The deployment of cache servers in p2p networks for improved performance in content-delivery. In *Third International Conference on Peer-to-Peer Computing (P2P'03)*, Linkoping, Sweden, September 2003.
- ▶ [14] A. Parker. The true picture of peer-to-peer filesharing, 2004. <http://www.cachelogic.com/>.
- ▶ [15] D. Qiu and R. Srikant. Modeling and performance analysis of bit torrent-like peer-to-peer networks. In *ACM SIGCOMM*, Portland, OR, USA, August 2004.
- ▶ [16] S. Saroiu, P. K. Gummadi, and S. D. Gribble. A measurement study of peer-to-peer file sharing systems. In *Multimedia Computing and Networking (MMCN'02)*, San Jose, CA, USA, January 2002.
- ▶ [17] S. Sen and J. Wang. Analyzing peer-to-peer traffic across large networks. *IEEE/ACM Transactions on Networking*, 12(2):219-232, 2004.
- ▶ [18] R. Sherwood, R. Braud, and B. Bhattacharjee. Slurpie: A cooperative bulk data transfer protocol. In *IEEE Infocom*, Honk Kong, China, March 2004.

THANK YOU!
ANY QUESTIONS?