


Comparing BitTorrent Clients in the Wild: The Case of Download Speed

Marios Iliofotou
University of California, Riverside

Georgos Siganos, Xiaoyuan Yang, and
Pablo Rodriguez
Telefonica Research, Barcelona

April 27th ♦ IPTPS 2010 ♦ San Jose, CA



Motivation

- BitTorrent (BT) is the most popular P2P file-sharing protocol
 - 20 million daily users in Pirate Bay alone
- Today, there exist many BT clients
 - How much do we know about them?



Vuze

uTorrent



Transmission

etc.

Question: Are there differences between these clients?


Contributions

- Large scale measurement study of real BT clients
 - 10 million users over one month
 - From 6,000 ISP
- **Observation 1: uTorrent** achieves on average **16%** higher download speeds compared to **Vuze**
- **Observation 2: We identify four implementation differences**
 - Neighborhood management
 - Opening of new connections
 - Termination of connections
 - Upload bandwidth distribution

Outline

- Introduction
- **Dataset collection**
- Part A: Speed comparison
- Part B: Implementation differences
- Conclusions and future work

Download speed inference using Apollo



$$Estimated - Speed = \frac{\# Pieces \times Size}{Time}$$

Dataset description

■ Collection process

- Apollo connects to the top 600 torrents from Pirate Bay (every hour)
- It connects multiple times with as many BT hosts possible

■ BT Clients used: Vuze and uTorrent

- 75% of all BT hosts use these clients (Pirate Bay)

■ Summary statistics

- Data cover 1 month (a representative week is used in the paper)
- 10 millions BT hosts
- 6,000 ISPs

Fair comparison between Vuze and uTorrent

- External factors that can effect download speed
 - **The torrent:** A local torrent can be faster (e.g., a Spanish film)
 - **The ISP:** An ISP can be faster than others
- **Solution:** Select same # hosts from each {torrent, ASN} pair

- Data after processing
 - Overall equal number of Vuze and uTorrent samples
 - 10K up to 80K samples per ASN for each client

Outline

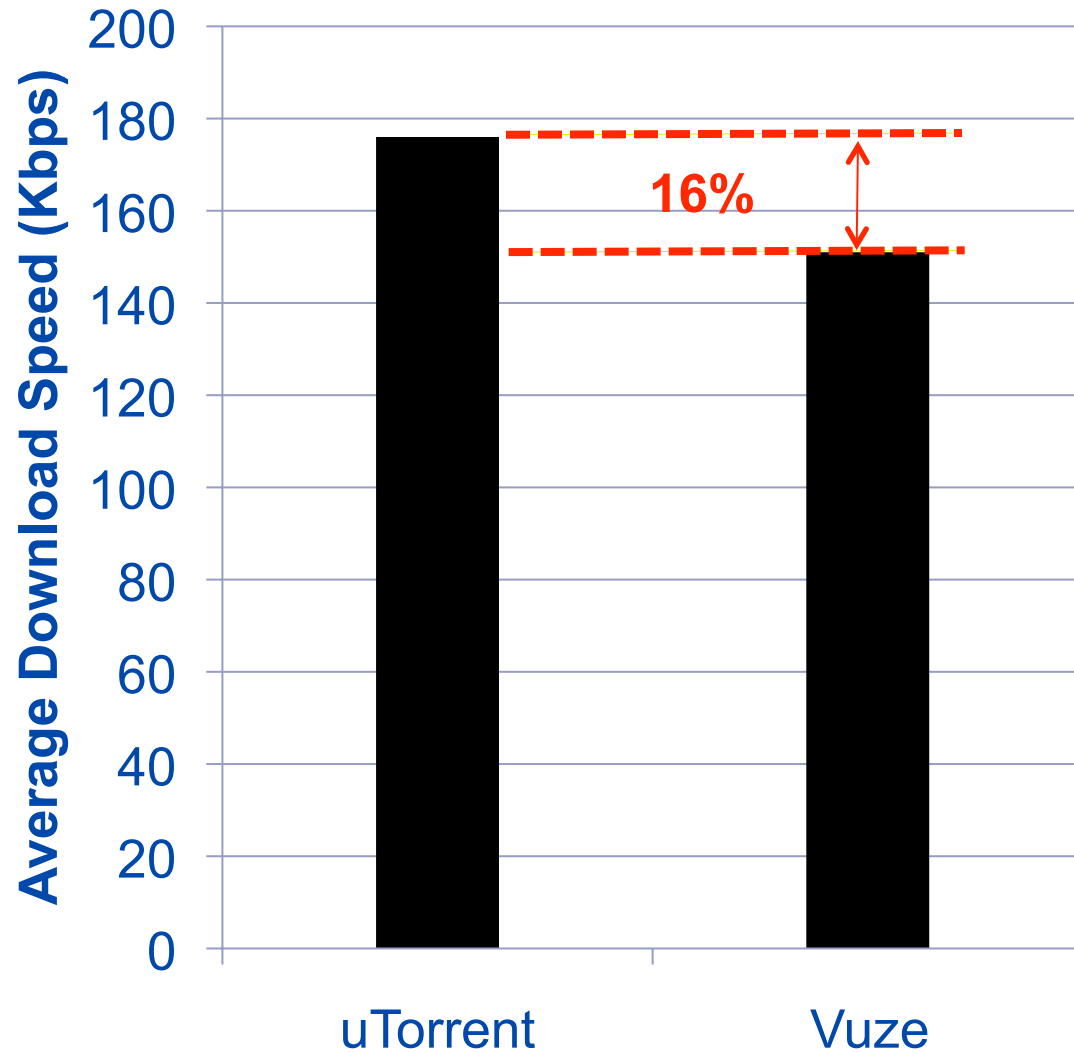
- Introduction
- Dataset collection
- **Part A: Speed comparison**
- Part B: Implementation differences
- Conclusions and future work

Overall speed differences

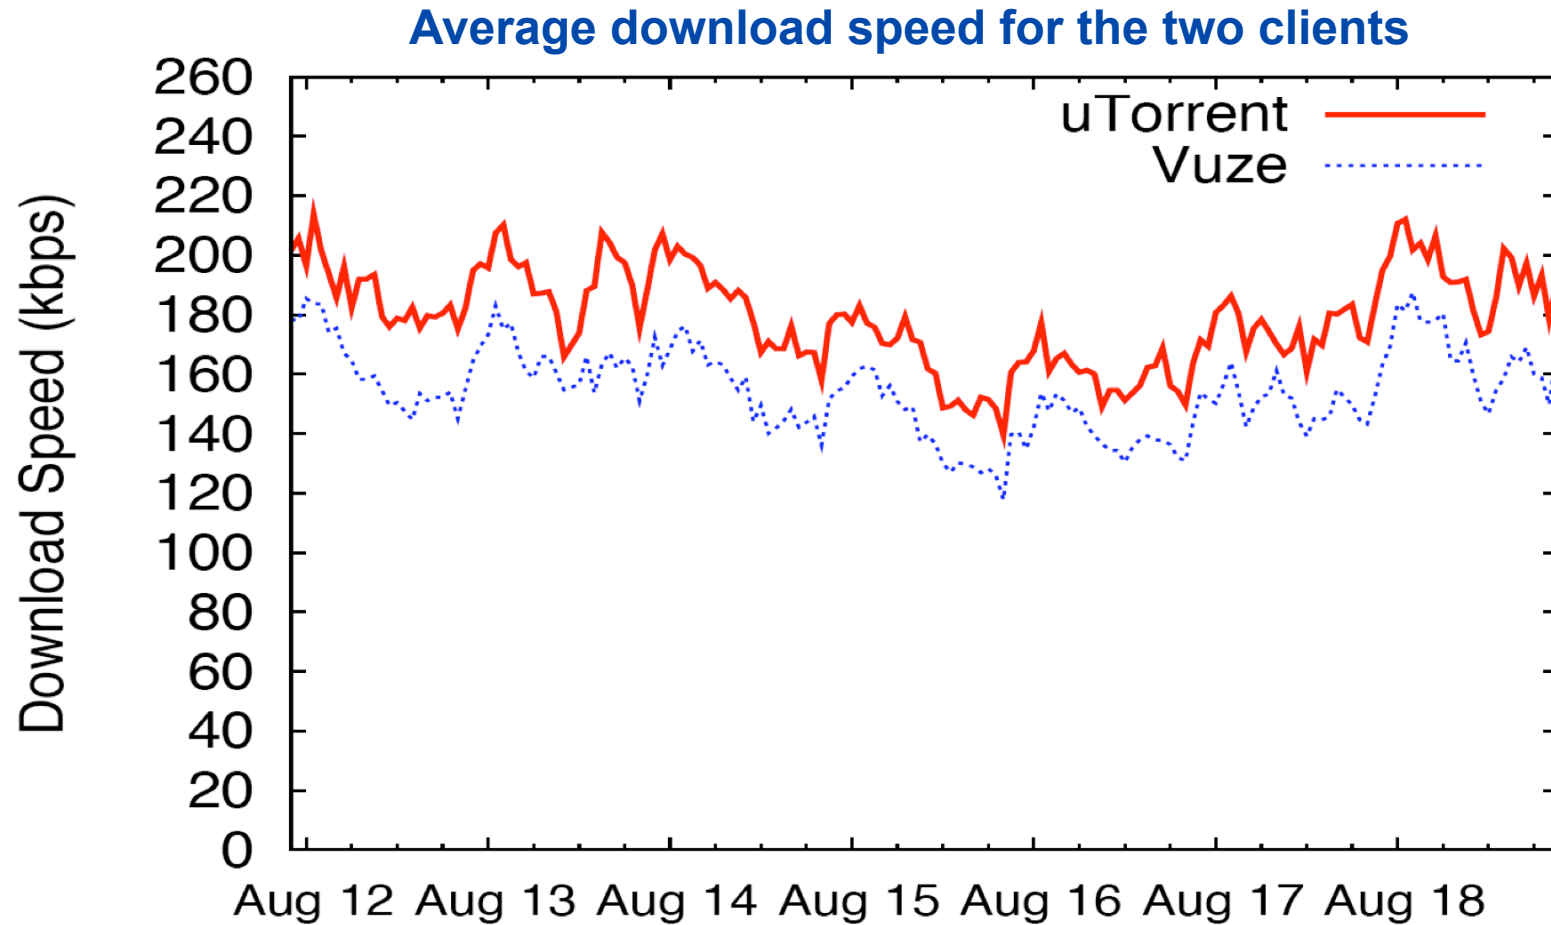
Perc.	uT (Kbps)	Vuze (Kbps)	Ratio
30 th	65	61	1.07
50 th	90	81	1.11
60 th	119	106	1.12
70 th	165	147	1.13

Two main observations:


- uTorrent faster by 16%
- uTorrent is faster for the higher percentiles



Consistent behavior over time



Comparing speeds for different ISPs




Outline

- Introduction
- Dataset collection
- Part A: Speed comparison
- **Part B: Implementation differences**
 - Neighborhood management
 - Opening of new connections
 - Termination of connections
 - Upload bandwidth distribution
- Conclusions and future work

Data collection

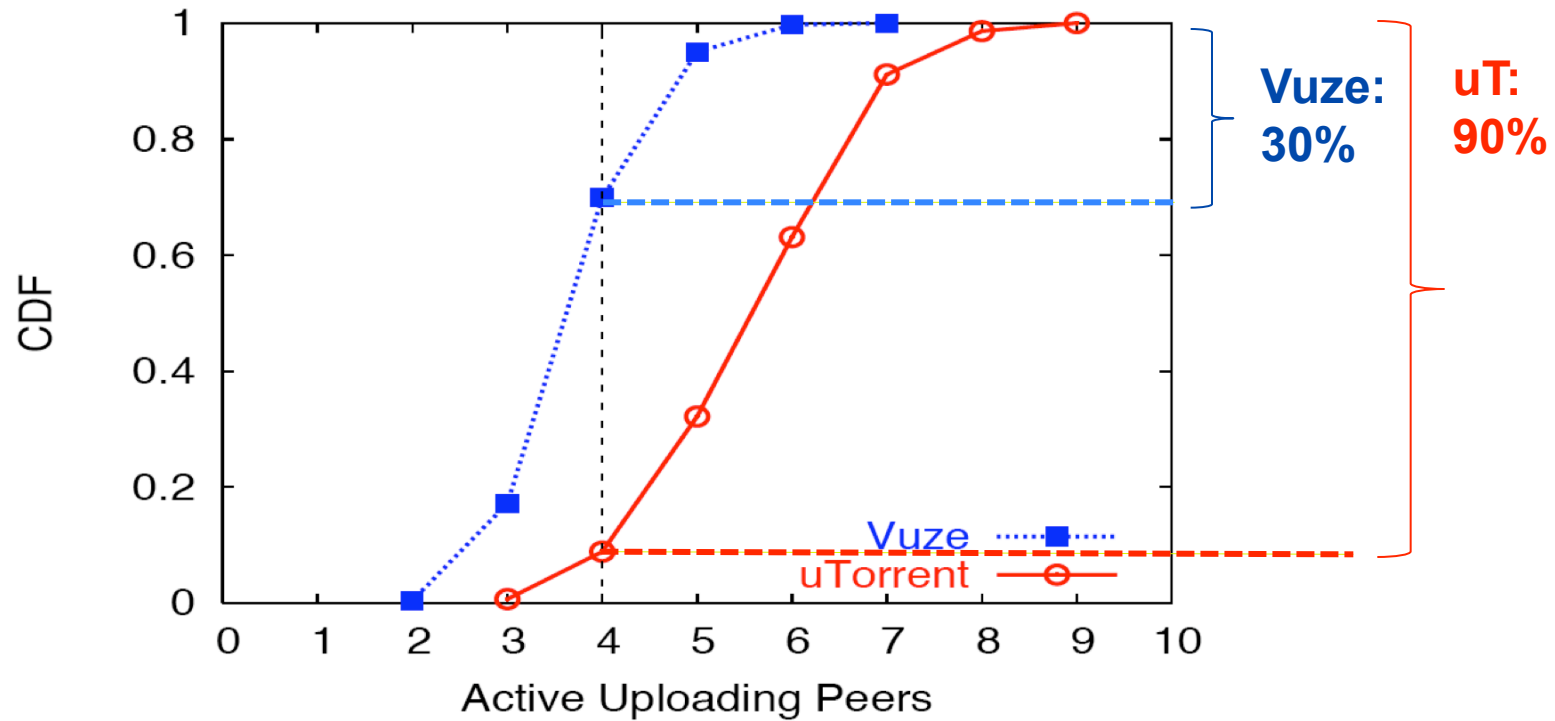
- Trace driven reconstruction of internal state
 - Keep track of active connections
 - Estimate upload and download rate of each connection
- Why this approach?
 - It can be applied to any client (open or close source)
 - We have the same measurement methodology (no biases)
- Dataset
 - Experiments ranged over two months
 - Different times of the days, weekdays and weekends
 - Always downloading a popular torrent (top 20 in Pirate Bay)

Neighborhood set over time



- Both clients are set to have 50 open connections
 - Vuze follows this closely
 - uTorrent shows higher variability
- Vuze has more ephemeral connections (open for less than 5 mins)

Upload bandwidth distribution



- uTorrent simultaneously uploads to more peers
 - uTorrent 90% of the time uploads to more than 4
 - Vuze only 30% of the time uploads to more than 4

Outline

- Introduction
- Dataset collection
- Part A: Speed comparison
- Part B: Implementation differences
- **Conclusions and future work**

Conclusions

- BT is popular, but how much do really we know about it?
 - Our measurements open many interesting questions
- Our goal is to bring to the attention of BT designers that
 - Some design choices can have a significant effect in performance
 - More effort should be added in evaluating design choices
- We don't claim that uTorrent is the way to go
 - We see our work opening the door for new research efforts to measure and evaluate real BT clients in the wild

Future work

- Compare more BT clients
 - E.g., transmission, bitcommet, etc.
- Test different implementation choices in a controlled environment
 - Testbed
 - Planetlab
- Study the effect of the ISP and their offered services to the performance of different clients

Thank You!
Questions/Discussion?