



www.tribler.org

BitTorrent or BitCrunch: Evidence of a credit squeeze in BitTorrent?

David Hales, Rameez Rahman, Boxun Zhang, Michel Meulpolder, Johan Pouwelse,

Tribler group, Technical University of Delft, Netherlands.

COPS Workshop @ WETICE 2009, Groningen, 29/06/09

Public Trackers

- BitTorrent uses Trackers to index swarms
- Public trackers let anyone join a swarm
- Sharing is incentivised via a form of tit-for-tat
- However there is no incentive for:
 - Seeding (uploading after file is downloaded)
 - Capping (creating and injecting a new file)

Private Trackers

- Private Trackers have emerged more recently
- Only allow registered users to join swarms
- Track upload / download of each user
- Keep centralised accounts for each user
- When users download much more than upload they may be kicked out
- Different schemes: ratio, credits, points etc

Private Trackers - Credit

- Consider a scheme based on credits
 - Uploading 1MB earns one credit
 - Downloading 1MB costs one credit
 - A user with no credits can't download
- Users must be given some initial credit
- In fixed size pop. total credit remains constant
- Similar to a fixed supply of money in an economy (loose analogy!)

Private Trackers - Credit

- How much credit should be put into the system?
- How would it effect the efficiency of the system?
- When do credit squeezes occur?
- How can they be avoided?

We define a credit squeeze as a situation in which, due to lack of credit, the efficiency of the system is significantly reduced.


BitCrunch Model

- Highly abstract and simplified model
 - All nodes have equal upload / download
 - Equally interested in all swarms on the tracker
 - Always uploading to one swarm (seeding)
 - Always downloading from another swarm (leeching)
 - No modelling of tit-for-tat or free-riding
 - Always online, fixed population
 - If run out of credit (broke) must wait until earns some via upload before being allowed to download
 - Swarms assumed to share upload “perfectly”

BitCrunch Model – baseline runs

- Parameters:
 - 500 peers, 100 swarms
 - Peer upload and download capacity = 1 unit
 - Each file shared in each swarm = 10 units size
 - One simulation cycle = each swarm processes one unit of time
 - Run for 20,000 cycles (x10 runs)
 - For initial credit per peer of 1, 10 and 100 units

Typical baseline simulation run



Baseline simulation results

C	T	β	G	φ
1	0.58	0.36	0.87	0.84
10	0.81	0.20	0.77	0.43
100	0.97	0.06	0.59	0.10

C = initial credit

T = total throughput = total number of units uploaded as proportion of maximum possible (infinite credit)

B = proportion of nodes that are “broke” (zero credit)


G = Gini measure (simple measure of inequality of credit)

Phi = turnover of top 10% of peers ranked by credit (credit mobility)

Unequal capacities runs

- To determine what happens when some nodes of different upload capacities
- Parameters (same as baseline runs but):
 - All peers download capacity = 10 units
 - 10% of peers upload capacity = 10 units
 - 90% of peers upload capacity = 1 unit
 - Examined a (1.5 credit) seeding bonus approach to dynamically introduce more credit into the system

Typical unequal capacities run



Unequal capacities simulation results

C	T	β	G	φ
1	0.56	0.39	0.90	0.82
10	0.71	0.32	0.93	0.44
100	0.77	0.29	0.94	0.06
100++	0.97	0.01	0.71	0.00

C = initial credit

T = total throughput = total number of units uploaded as proportion of maximum possible (infinite credit)

B = proportion of nodes that are “broke” (zero credit)

G = Gini measure (simple measure of inequality of credit)

Phi = turnover of top 10% of peers ranked by credit (credit mobility)

100++ indicates initial credit of 100 with 1.5 credit seeding bonus

Observations

Even in a trivial model where all peers have the same capacities and user behaviour, all swarms have equal popularity and all peers start with equal credits, the performance of the system may be inhibited by credit shortages

Observations

Adding extra capacity to the system, in the form of upload and download, can actually reduce the performance. This is highly counter intuitive and something that should be avoided because it implies lack of scalability.

Observations

By injecting new credit into the system in the form of a “seeding bonus” a credit squeeze can be ameliorated when peer capacities are unbalanced.

Statistics from a Private Tracker

Day	T	Δ	Δ_0	δ	S/L
1	48	24	17	0.23	26
2	40	20	15	0.25	26
3	50	25	12	0.16	25
4	67	33.5	17	0.17	25
5	52	26	19	0.24	25
6	46	23	15	0.21	25
7	87	43.5	17	0.13	25
Ave.	56	28	16	0.19	25

Approx. 50,000 peers per day, 10,000 swarms,
access to credit balances of top 10%

T = throughput in TB over all swarms

Δ = total credit increase that day in the entire system

Δ_0 = total credit increase for top 10% of peers

δ = minimum fraction of credit increase that goes to top 10% of peers

S/L = seeder to leecher ratio over all swarms

Statistics from a Private Tracker

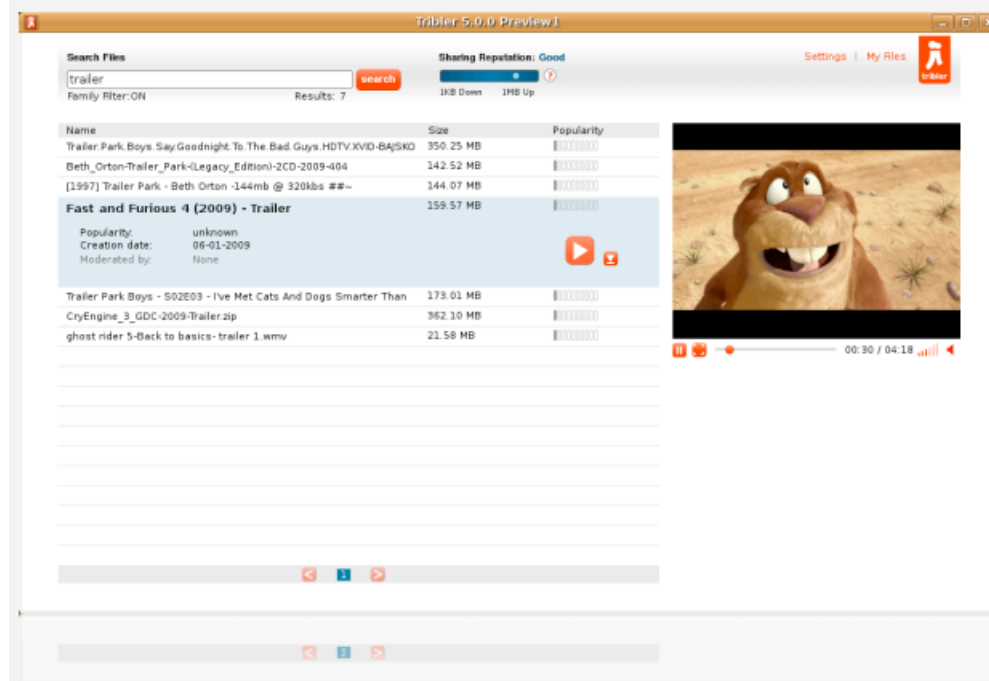
- Indicates “rich getting richer” since top 10% are getting a lot of the new credit
- High Seeder / Leecher ratio suggestive that a credit squeeze is happening for many
- But need more information to verify this
- Would be interesting to see what happened to throughput if there was a “free day” or seeding bonus was increased

Conclusions

- Private trackers using “ratio enforcement” policies appear to be ad-hoc and various
- But can have dramatic effects on efficiency
- Too much credit could encourage free-riding
- Too little creates squeezes = lower efficiency
- These are just initial investigations
- Much more work needs to be done!



Advert / plug / shameless promotion!
Download tribler 5.1 at:
www.tribler.org



Tribler 5.1

We aim to show just how easy file-sharing can be. Due to our simplified search box you can find your files without the need for a website.

- Most easy to use
- Integrated search box
- Video-on-demand support
- Fully distributed

[Download page](#)

[[Windows](#) | [Mac](#)]

Comments / discussion and suggestions on:
forum.tribler.org