# Can We Make BitTorrent ISP Friendly?

Stevens Le Blond and Arnaud Legout
I.N.R.I.A.
Sophia Antipolis France
{Stevens.Le_Blond,Arnaud.Legout}@sophia.inria.fr

***Abstract*— The huge popularity of Peer-to-Peer (P2P) content distribution has modified the workload of the internet and changed the relationship among content providers, Internet Service Providers (ISPs), and end-users. Indeed, content providers used to distribute their contents using dedicated servers or CDNs. Whereas those infrastructures are costly, they respect the hierarchical structure of the internet by minimizing the load on ISPs' peering links. However, the recent shift to P2P content distribution has reduced the cost for content providers at the expense of ISPs, P2P applications generating much of the paid traffic on ISPs' peering links. This shift in cost is one of the main obstacles to the widespread adoption of commercial P2P content distribution solutions.**

**Using large scale BitTorrent experiments on $10,000$ instrumented peers, we show that we can reduce the load on ISPs' peering links by up to $2$ orders of magnitude without a major impact on the download completion time of the end-users. This gain is obtained with a simple modification of the tracker and does not require any specific support on the end-users clients. Therefore, we argue that P2P content distribution can be cost effective for both content providers and ISPs, while efficient for end-users.**

## I. Introduction

Content distribution is today at the core of the services provided by the internet. However, distributing content to a large audience is costly with classical client-server or CDN solutions. This is the reason why content providers start to move to Peer-to-Peer (P2P) content distribution that allows to significantly reduce the cost for the content providers without penalizing the experience of the end-users. One striking example is the P2P iPlayer used to provide a video-on-demand service of the recent BBC programs.

However, whereas the current P2P content distribution solutions like BitTorrent [4] are very efficient to minimize the end-users download time [7]–[10], they can create a huge increase of traffic on ISPs' peering links, which are the scarcest and most expensive ISPs' resources. Indeed, in BitTorrent, each peer is connected to a random set of neighbors, and it can source the content to any of them. Hence, the lack of locality in BitTorrent neighbor selection favors overhead on the ISPs' peering links.

Therefore, even if current P2P content replication solutions significantly reduce content provider costs, they cannot be promoted as a global solution for content replication as they induce a major cost increase for ISPs. In particular, the current trend of ISPs is to block P2P traffic that saturates peering links [2].

One solution to this problem is to use P2P locality, i.e., to constrain the P2P traffic within the ISP boundaries in order to minimize the amount of traffic that crosses peering links [3], [6]. However, there are three main issues with P2P locality.

- How much can we reduce the traffic on ISPs' peering links? Peer-to-peer traffic creates load on peering links. However, the exact amount of traffic needs to be quantified.
- What is the impact of locality on end-users download time? As we create locality, we constrain the P2P connectivity graph, which might have an severe impact on the experience on the end-users.
- How to implement a locality policy with no cost (in terms of servers to maintain) for the ISPs and no specific support in the client used by the end-users? Peer-to-peer represents a huge saving for content providers, however, ISPs have no specific reason to support P2P. Therefore, a solution that does not require a specific investment from the ISPs (for instance in terms on servers) will be a plus to convince ISPs not to throttle P2P traffic.

We use a locality policy that does not require any support from the end-users and ISPs, but only a modification of the BitTorrent tracker that is usually hosted by the content provider. We have performed BitTorrent experiments with $10,000$ peers running a real BitTorrent client in a controlled torrent to explore the impact of this locality policy. We show that:

- The locality policy reduces by up to two orders of magnitude the load on ISPs' peering links.
- The locality policy only moderately increases the mean time for peers to download the content.
- BitTorrent achieves a remarkable efficiency for torrents up to $10,000$ peers.

The main goal of our work is to get a deeper understanding of the impact of locality on the load of ISPs' peering links and on end-users experience for large torrents. We do not claim that our proposed locality policy is optimal, nor that we have proposed a complimentary architecture to implement it. Our main conclusion is that P2P content replication is a viable and cost-effective solution for both content providers and ISPs for large scale services. The architectural work required to implement an efficient locality policy is orthogonal to this work and is currently addressed in, for instance, the on-going works P4P [5] and Ono [1].

## II. Experimental Results

In this section, we explore the impact of locality on the number of times a $100$ MB file is uploaded across ISPs and the average download slowdowns of peers as compared to the best average download completion time. We refer to those two metrics as content replication overhead and slowdown. Then, we define locality as the percentage of cross-ISP peers that are returned by the tracker. Hence, when a peer $P_A$ from an ISP $A$ asks for a set of peers, the tracker returns peers inside $A$ with probability $\frac{n}{100}$, and outside $A$ with probability $1 - \frac{n}{100}$. We choose a locality value of $99.9$ for our experiments, meaning that one connection on $1,000$ will be cross-ISP in average. We also vary the number of ISPs and limit ourselves to a scenario with a single initial source that uploads at $20$ kB/s and stays connected for the entire duration of the experiment. Finally, peers source the content for $5$ before disconnecting.

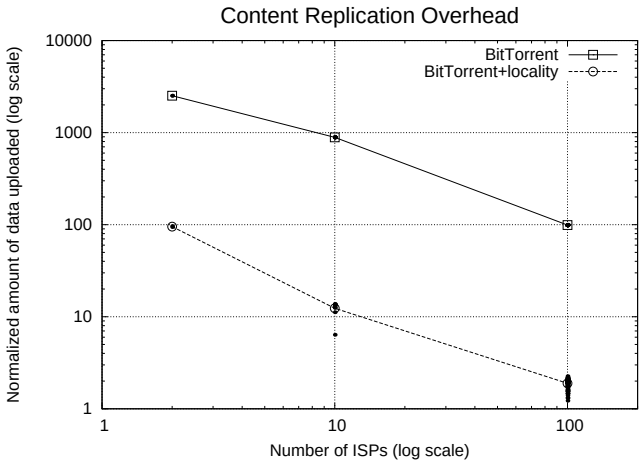


Fig. 1. Content replication overhead for 2, 10, 100 ISPs with a single source. Each dot represents the content replication overhead for a single ISP. The square (BitTorrent) and the circle (BitTorrent+locality) represents the content replication overhead averaged over all ISPs for a given experiment.

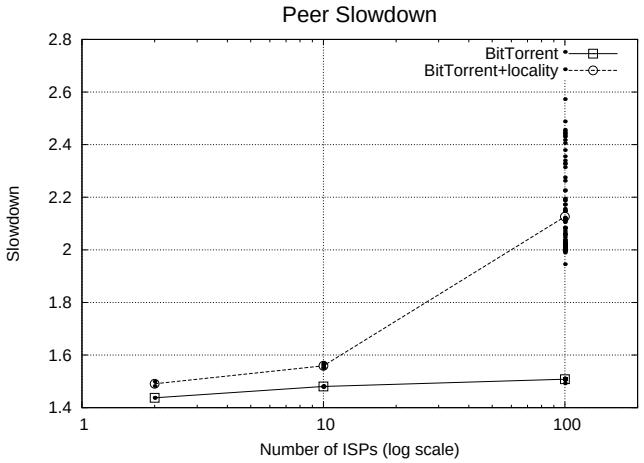


Fig. 2. Peer slowdown for 2, 10, 100 ISPs with a single source. Each dot represents the peer slowdown for a single ISP. The square (BitTorrent) and the circle (BitTorrent+locality) represents the peer slowdown averaged over all ISPs for a given experiment.

Fig. 1 shows that the amount of data uploaded on the ISPs' peering links is reduced by up to 2 orders of magnitude with locality as compared to regular BitTorrent. Indeed, with locality, the content replication overhead of BitTorrent is divided by a factor $26.5$, $74$, and $55$ for respectively 2, 10, and 100 ISPs. We also observe that the results are consistent among the ISPs, as the dots representing the per-ISP content replication overhead are close to each other. The larger the number of ISPs, the lower the amount of data uploaded per ISP. This is because, as we keep the torrent size constant, increasing the number of ISPs reduces the number of peers per ISP, thus the traffic uploaded from the ISPs.

This huge bandwidth saving for the ISPs is at the cost of a negligible to moderate increase in the peers download completion time. Indeed, we see in Fig. 2 that the peer slowdown is $4.1\%$, $4.7\%$, and $41\%$ larger with locality than with regular BitTorrent for respectively 2, 10, and 100 ISPs. Therefore, we see that for a small number of ISPs, locality has almost no impact on the peer slowdown. For a large number of ISPs, the slowdown is moderately increased, but it has to be pondered with the huge bandwidth saving for the ISPs.

We observe in Fig. 2 that with $100$ ISPs and locality, there is larger heterogeneity in the per-ISPs peer slowdown than for the other experiments. This is because our locality policy is probabilistic and thus may lead to some ISPs having more cross-ISP-connections than others. This explains the observed heterogeneity, because when there is few connections among ISPs, the random variations in the inter-ISP connections lead to significant variations in the peer completion time.

In summary, we have shown in this section that locality can dramatically reduce the load on ISPs' peering link with only a negligible to moderate increase in the download completion time experienced by the end-users.

## III. Conclusion

We have performed P2P experiments with up to $10,000$ peers running an instrumented BitTorrent client in order to explore the impact of locality on the load of ISPs' peering links and on the end-users download completion time.

Using a simple locality policy that only requires a modification of the tracker we have shown that:

- The proposed policy allows to reduce the load on ISPs' peering links by up to 2 orders of magnitude.
- The proposed policy slightly increases the peer download completion time, but this increase is small enough to do not significantly impact end-users experience for an asynchronous content distribution service like the one offered by BitTorrent.

Finally, as a consequence of our experiments, we have also validated that BitTorrent achieves a remarkable efficiency for torrents up to $10,000$ peers.

In conclusion, we have shown that our simple locality policy achieves huge bandwidth saving on ISPs' peering link with a minor to moderate impact on the peers download completion. Moreover, this result is consistent for a large variety of scenarios. We expect this work to provide important insights on the impact of a locality policy on large torrents.

## Acknowledgments

Experiments presented in this paper were carried out using the Grid'5000 experimental testbed, an initiative from the French Ministry of Research through the ACI GRID incentive action, INRIA, CNRS and RENATER and other contributing partners (see https:// www.grid5000.fr)

## References

[1] http://www.aqualab.cs.northwestern.edu/ projects/ono.html.
[2] http://www.azureuswiki.com/index.php/bad_isps.
[3] R. Bindal, P. Cao, W. Chan, J. Medved, G. Suwala, T. Bates, and A. Zhang. Improving traffic locality in bittorrent via biased neighbor selection. In *Proc. of ICDCS'06*, Lisboa, Portugal, July 2006.
[4] B. Cohen. Incentives to build robustness in Bittorrent. In *Proc. of the Workshop on Economics of Peer-to-Peer Systems*, Berkeley, CA, USA, June 2003. http://bitconjurer.org/BitTorrent.
[5] A. K. H. Xie, A. Silberschatz, and Y. R. Yang. P4P: Explicit communication for cooperative control between p4p and network providers, 2007. www.dcia.info/documents/P4P_Overview.pdf.
[6] T. Karagiannis, P. Rodriguez, and K. Papagiannaki. Should internet service providers fear peer-assisted content distribution? In *Proc. of IMC'05*, Berkeley, CA, USA, October 2005.
[7] A. Legout, N. Liogkas, E. Kohler, and L. Zhang. Clustering and Sharing Incentives in BitTorrent Systems. In *SIGMETRICS'07*, San Diego, CA, USA, June 2007.
[8] A. Legout, G. Urvoy-Keller, and P. Michiardi. Rarest First and Choke Algorithms Are Enough. In *IMC'06*, Rio de Janeiro, Brazil, October 2006.
[9] D. Qiu and R. Srikant. Modeling and performance analysis of bittorrent-like peer-to-peer networks. In *Proc. of SIGCOMM*, Portland, Oregon, USA, August 2004.
[10] X. Yang and G. de Veciana. Service capacity of peer-to-peer networks. In *Proc. of INFOCOM*, Hong-Kong, China, March 2004.