

Experimental Assessment of BitTorrent Completion Time in Heterogeneous TCP/uTP swarms

Claudio Testa¹, Dario Rossi¹, Ashwin Rao², Arnaud Legout²

¹ Telecom ParisTech, Paris, France – `first.last@enst.fr`

² INRIA Planete, Sophia Antipolis, France – `first.last@inria.fr`

Abstract. BitTorrent, one of the most widespread used P2P application for file-sharing, recently got rid of TCP by introducing an application-level congestion control protocol named uTP. The aim of this new protocol is to efficiently use the available link capacity, while minimizing its interference with the rest of user traffic (e.g., Web, VoIP and gaming) sharing the same access bottleneck.

In this paper we perform an experimental study of the impact of uTP on the torrent completion time, the metric that better captures the user experience. We run BitTorrent applications in a flash crowd scenario over a dedicated cluster platform, under both homogeneous and heterogeneous swarm population. Experiments show that an all-uTP swarms have shorter torrent download time with respect to all-TCP swarms. Interestingly, at the same time, we observe that even shorter completion times can be achieved under mixtures of TCP and uTP traffic, as in the default BitTorrent settings.

1 Introduction

Though some might argue that network congestion control is a problem that has been studied to death—to which we tend to agree, at least concerning the large amount of literature on the topic—yet the network architecture and usage are undergoing profound changes, that make the study of congestion control issues once more necessary.

As far as the architectures are concerned, recent research has, e.g., addressed the TCP incast problem in data center networks. As far as the usage is concerned, we have lately witnessed to an explosion of new application-layer flow and congestion control algorithms [1–3], which are usually implemented at the application-layer over either TCP [1] or UDP [2, 3]. Depending on the application they have been built for, these protocols have rather different goals that deeply influence their design.

In this work, we focus on the BitTorrent filesharing protocol that recently replaced TCP by uTP¹ [4, 5], a lower than best effort protocol for data transport on top of UDP. uTP starts from the observation that nowadays the Internet bottleneck is typically at the user ADSL access link: hence, congestion typically happens between different flows of the same user. Moreover, since ADSL modems have rather long buffers (up to few seconds [6, 7]), using TCP for non-interactive but massive data downloads has a possibly negative impact on interactive communication (e.g., Skype, gaming, etc.). In other

¹ Notice that this new protocol has two names: uTP in the BitTorrent community [4], and LED-BAT in the IETF community [5], that we may use interchangeably in the following.

words, as TCP fills the buffer, the self-induced congestion translates into high latency and losses, that possibly harms other interactive application. To avoid the troubles of self-induced congestion and at the same time be efficient for massive data download, uTP tries to limit the end-to-end delay (by reaching a fixed amount of *target delay* in the access buffer) while maximizing the utilization of the access capacity. As the queuing delay is bounded with uTP, this improves users experience for interactive applications.

While uTP has sound and appealing goals, it is clearly understood that users will be the ultimate judge of BitTorrent performance, as in BitTorrent’s own words “*unless we can offer the same performance [of TCP], then people will switch to a different BitTorrent client*” [8]. Our recent work [9] suggests that uTP performs better than standard TCP, as the use of uTP practically limits the queuing delay to a small target, this translates into faster signaling as a side effect. However, results in [9] are based on ns2 simulation: it becomes thus imperative to assess whether the observed phenomena also happens in practice, which is precisely the scope of this work.

We run BitTorrent applications in a flash crowd scenario over the Grid’5000 platform [10], with special attention to the main user-centric metric, the torrent completion time. Results of our experiments confirm our previous simulation results, in that, as observed in [9], uTP can reduce the torrent download time.

Yet, this experimental work brings new insights beyond [9]. Currently, the default settings of BitTorrent yield to the use of a *mixture of TCP and uTP traffic*. Hence, in this work we evaluate how this choice performs compared to the cases in which all peers use only a single protocol between TCP and uTP. In this case, our experimental results show that completion time under heterogeneous swarms can be even lower than all uTP (and, clearly, all TCP) swarms.

The remainder of this paper is organized as follows. First, Section 2 discusses the related work. Section 3 reports preliminary insights on low-level BitTorrent settings gathered from a small local testbed, which are instrumental to our experiments, whose results are reported in Section 4. Finally, we summarize the paper and discuss future work in Section 5.

2 Related work

Two bodies of work are related to this study: on the one hand, we have work focusing on congestion control aspects [1, 2, 11–13], and on the other hand work focusing on BitTorrent [7, 9, 14–20]. First, congestion control literature already proposes several protocols aiming, as LEDBAT, to achieve lower-than-TCP priority, of which TCP-LP [12], NICE [11], 4CP [13] are notable examples. Yet, we pinpoint a recent tendency toward moving congestion and flow control algorithms *from the transport layer to the application layer*, of which uTP [3] for background file-sharing, Skype [2] for interactive VoIP and YouTube [1] for interactive VoD are again notable examples. Unlike transport-layer congestion control, that applies to classes of applications, these application-layer congestion control protocols are usually built for single applications, with specific requirements in mind: these “one of a kind” deployments will in our opinion need further attention in the future. Recently LEDBAT also got the attention of Apple’s developers, resulting in a implementation for MAC OS (preliminary tests available at [21]).

Second, BitTorrent literature dissected many aspects of this successful P2P protocol, from the pioneering time of [14]. While our own previous works, such as [15, 16], already study BitTorrent download performance by means of either passive measurements or experimental tests (as in this work), however they report on performance at a time when BitTorrent was using TCP, and should thus be updated in light of BitTorrent recent evolution. More generally, though related work on uTP exists [7, 9, 17, 17–20, 22], it does however adopt a congestion control perspective (with the exception of [9]). In particular, an experimental approach is adopted in [6, 7, 17]: [17] attacks the problem of clock drift in uTP, while [7] performs a black box study of initial proprietary versions of the protocol and [6] focuses on the interaction of uTP and active queue management techniques that are becoming commonplace in modern home gateways. A simulative approach is instead adopted in [18–20, 22]: a fairness issue of uTP is revealed in [18] and solved in [19], while [20] compares (i) the level of low-priority of TCP-LP [12], NICE [11] and uTP [3] and (ii) the fairness of TCP and uTP, and finally [22] investigates policies for dynamic parameter tuning.

The only previous work addressing the impact of uTP on BitTorrent completion time is our own recent work [9], that however employs ns2 simulations unlike in this work. Interestingly, some of the observations of this study are in agreement with [9], e.g., showing a larger completion time for increasing buffer occupancy on the data plane. Yet, we point out that [9] does not consider an hardcoded preference for uTP, nor bidirectional uTP connections: hence, an interesting difference with respect to the current work is that [9] forecasted heterogeneous performance for heterogeneous swarms (i.e., larger completion time for TCP peers) that we have shown not to hold on practice.

3 Preliminary Insights

As previously stated, the uTP protocol aims at jointly (i) being efficient by fully exploiting the link capacity when no other traffic is present, and (ii) being low priority by yielding to other competing traffic on the same bottleneck. In order to achieve both these goals, uTP needs to insert only a limited amount of packets in the bottleneck buffer: on the one hand, since the queue is non empty, the capacity is fully exploited. On the other hand, as the queuing delay in the buffer remains bounded, this does not harm interactive applications.

uTP exploits the ongoing data transfer to measure the One-Way Delay (OWD) on the forward path. While measuring the OWD is notoriously tricky among non-synchronized Internet hosts, uTP is interested in the *difference* between the current OWD and the minimum OWD ever observed (used as an approximate reference of the base propagation delay). In turn, this OWD difference yield to a measure of the current *queuing delay*, that is used to drive the congestion window dynamics: when the measured queuing delay is below a given *target delay*, the congestion window grows, but when the queuing delay exceeds the target the congestion window shrinks.

The impact of this new protocol on the performance of BitTorrent can be affected by essentially two different settings. At a single flow level, uTP is primarily driven by the uTP *target delay* setting. At a swarm level, peers *relative preference* for TCP vs uTP protocols plays an important role as well. Hence, before running a full-fledged set of

experiments, we need to get some preliminary insights on the settings of the above two parameters. In more details, these are: (i) *net.utp_target_delay*, that tunes the value of uTP delay of each flow, and (ii) *bt.transp_disposition*, that drives TCP vs uTP preference of the BitTorrent client.

To this aim, we perform a battery of tests in a completely controlled environment involving one seed and two leechers, all running the latest version of the uTorrent client available for Linux (3.0 build 25053, released on March 2011). Clients in the local testbed are interconnected by a 100Mbps LAN and we use the Hierarchical Token Bucket (HTB) of *netem* to emulate an access bottleneck on the PC running the seed, whose uplink capacity is then capped at 5 Mbps. Any additional delay was added on the testbed links. The tracker is private within the testbed and is used to announce a set of three different torrent files having different file and chunk sizes (file size of 10, 50 and 100 MBytes, and chunk size of 256, 512 and 1024 KBytes respectively).

To understand the BitTorrent settings, we tweaked the default configuration² aiming at (i) verifying the compliance of *net.utp_target_delay* to the imposed delay and (ii) understanding how *bt.transp_disposition* settings, which controls when uTP is used, impacts the performance of BitTorrent.

In the case of (i) *net.utp_target_delay*, usually a single experiment is sufficient to verify its compliance, since every flow obeys its own setting. Conversely multiple experiments were necessary in the case of (ii) *bt.transp_disposition*, since the behavior of a peer is affected by the *bt.transp_disposition* value of other peers as well.

In the following we summarize the most relevant findings of the local testbed experiments. Overall, in these tests we captured about 2 GBytes of packet level traces, that we make available to the scientific community at [23]. The remainder of the experiments, performed on the Grid'5000 platform, are reported in Section 4.

3.1 *net.utp_target_delay*: Target delay settings

The *net.utp_target_delay* parameter stores the value of the uTP target delay in milliseconds, and its default value is equal to 100 ms as stated in the BEP29 [4]. Furthermore, uTP is also standardized at the IETF LEDBAT Working Group [3], which specifies 100 ms as a *mandatory upper bound value* (while earlier version of the draft referred to a 25 ms delay target).

Yet, the GUI of the Windows client allows to modify its default value, opening the way for competition between legacy applications. This behavior is confirmed by Fig. 1, where we show two experiments, performed at different times, where a single uTP flow sends data on a 5 Mbps bottleneck, with different values of the uTP target delay. We see in Fig. 1 that for both target delay settings, BitTorrent is using the entire available capacity, and the end-to-end delay corresponds to the target that we set by means of the *net.utp_target_delay* parameter.

² uTorrent clients store their configuration in a file which is not directly editable (as it also contains an hash value on the configuration content, performed by the client itself) and moreover Linux GUIs do not offer the possibility of modifying the default settings. However, the configuration file format used by Linux clients is the same as the one of the Windows clients: hence, we used Windows GUIs to produce a pool of configuration files, that we later loaded in Linux.




Fig. 1. uTP Target Settings: RTT and throughput for a single flow with *net.utp_target_delay*=25 ms (left) and *net.utp_target_delay*=100 ms (right).

As the uTP/LEDBAT specifications [3,4] refer to a mandatory target value, we comply to the standard and focus in the remainder of this paper on the study of swarms with the same default value for *net.utp_target_delay*. At the same time, we point out that as a future work, it would be interesting to see whether, by tweaking the *net.utp_target_delay* value, some peers (or some applications) can gather an advantage over the rest of the swarm.

3.2 *bt.transp_disposition*: TCP vs uTP settings

The second parameter, namely *bt.transp_disposition*, controls which protocol is used for the incoming and outgoing data connection of the client. As reported in the online uTorrent manual³, *bt.transp_disposition* is a bitmask that sums up the following behaviors:

- 1: attempt outgoing TCP connections
- 2: attempt outgoing uTP connections
- 4: accept incoming TCP connections
- 8: accept incoming uTP connections
- 16: use the new uTP header format

uTorrent default value is 31, which means that the client will accept both TCP and uTP flavors, for either sending or receiving data, possibly using the new uTP header format.

To understand the implications of *bt.transp_disposition* settings, we perform a number of tests with heterogeneous settings of the client. Notice that the parameter space we explore and that we make available at [23] is larger than the one reported in Tab. 1. Yet, for the sake of simplicity, we only report in the table the cases that we later study in Section 4, which already conveys some interesting information. Notice also that in all the experiments, the seed is set with the default value *bt.transp_disposition*=31.

In *Case 1*, the two leechers A and B have different setting for the *bt.transp_disposition* parameter: more precisely, A should attempt data connection only using TCP while B

³ http://www.bittorrent.com/help/manual/appendixa0212#bt.transp_disposition

Table 1. TCP vs uTP BitTorrent transport disposition.

case	peer	attempt	accept	disp	B → A	A → B	comment
1	A	TCP	*	13	uTP	uTP	Hardcoded uTP preference
	B	uTP	*	14			
2	A	TCP	TCP	5	TCP	TCP	Legacy BitTorrent implementations
	B	*	*	31			

should use uTP (and both will accept every flavor in reception). Our experiments show that in this scenario, peer B sends data to peer A using the uTP protocol, which is the expected behavior. However, peer A sends data to peer B using the uTP protocol too, which happens consistently over all repetitions, and irrespectively of file and chunks size. The reason is that when a uTP connection from B to A is opened, peer A can use this opened bidirectional connection to send data to peer B. Besides, as confirmed by Arvid Norberg, one of the main BitTorrent developers, the uTorrent client has a *hard-coded uTP preference*, so that in case both a TCP and a uTP connection will be successfully established, the former will be closed and only the latter will be used. As we will see in Section 4, this preference has some important (and beneficial) consequences on the overall swarm completion time.

In *Case 2*, leecher A attempts and accepts only connection via TCP (as a legacy BitTorrent implementation would do), while leecher B maintains the default value for *bt.transp_disposition* (which means to attempt and accept both protocols). In this scenario, any communication between the two peers are performed using the TCP protocol, which is consistent and expected for backward compatibility with older BitTorrent clients.

Other cases, not shown in Tab. 1, yield to different shares of traffic between TCP and uTP. At the same time, since the number of leechers is small, the exact value of the breakdown is heavily influenced by the seeder flavor as well. As such, we defer a quantitative analysis of such a breakdown in the next section.

4 Experimental Results

We now report the experimental results on the impact of uTP and TCP transport on the torrent completion time. First we briefly describe the Grid’5000 experimental platform and then focus on two case studies, namely (i) homogeneous and (ii) heterogeneous swarms, depending on the *bt.transp_disposition* settings for the leechers.

Homogeneous settings refer to scenarios where all peers have either a TCP-only preference (*bt.transp_disposition=5*, which mimics the behavior of old uTorrent versions or legacy applications), or a uTP-only preference (*bt.transp_disposition=10*, in case uTP will prevail over TCP), or are able to speak both protocols (*bt.transp_disposition=31*, the current default behavior, though with an hardcoded preference for uTP as we have seen in Section 3.2).

Homogeneous settings provide a useful reference, but we must consider also experiments with heterogeneous scenarios that correspond to what is observed in the Internet with clients that do not support uTP at all, or that support uTP but as a fallback choice rather than the default one.

We therefore investigate heterogeneous settings as well, considering scenarios with different ratios of peers using uTP and TCP. More precisely, we consider the case where peers are able to accept any incoming protocol, but have different preferences for the uplink protocol (*bt.transp.disposition*=13 for TCP, and *bt.transp.disposition*=14 for uTP). We consider the case where the preference splits are 50/50, 25/75, or 75/25 to mimic scenarios where TCP vs uTP preferences are fairly balanced, or biased toward one of the two protocols.

Notice that, while there may be several uTP implementations available, different BitTorrent applications use different default settings (i.e., sticking to TCP preference or embracing uTP) depending on the success of the new protocol (and the existence of readily available libraries for different operating systems).

4.1 Grid'5000 Setup

We performed experiments on a dedicated cluster of machines that run Linux as the host operating system and using the uTorrent 3.0 client as before. Hosts of the Grid'5000 platform are interconnected by an high-speed 1 Gbps LAN, and we emulate realistic bandwidth restrictions and queueing of home gateways by using the `netem` module for the Linux kernel. As noted in [3] and experimentally confirmed by [6, 7], ADSL modems can buffer up to a few seconds worth of traffic: in our experiments, we set the buffer size B according to the uplink capacity C so that $B/C = 1$ second worth of traffic.

We instrumented the Linux kernel to log the queue size Q in bytes and packets after each dequeue operation, logging also the cumulative number of packets served and dropped by the queue. During our experiments, we disabled the large segment offloading [24] which ensured that the maximum segment size of the TCP and uTP packets never exceeded the maximum transfer unit (MTU). In each experiment we used the Cubic flavor of TCP, the default for Linux kernels: in reason of our previous work [7], we may expect Cubic to be more aggressive with respect to the standard TCP NewReno flavor, and more similar to the default TCP Compound flavor adopted in recent versions of Windows.

We use 76 machines on the Grid'5000 platform and consider an Internet flash crowd scenario, where a single seed is initially providing all the content (a 100 MBytes file) to a number of leechers all arriving at the same time (and never leaving the system). Furthermore, each BitTorrent peer (i) experiences an ADSL access bottleneck [25] and (ii) encounters a self-induced congestion, unrelated to the cross-traffic [3].

As for (i), we start by considering 3 simple homogeneous capacity scenarios in which we limit the leechers and seed uplink capacity to $C = 1, 2$ and 5 Mbps with the Hierarchical Token Bucket algorithm. For the sake of simplicity, as the qualitative results do not change for different values of C , in the following we report the results for $C = 1$ Mbps. While it could be objected that Internet capacity are not homogeneous, we argue that homogeneous scenarios are needed as a first necessary step before more

complex and realistic environments are emulated. Additionally, the impact of heterogeneous access capacity is a well known clustering effect [16], that we believe to fall outside of our main aim, i.e., the comparison of TCP vs uTP transport, and that can be studied with a future experimental campaign.

As for (ii), we are forced to map a single peer on each host of the Grid’5000 platform, as otherwise unwanted mutual influence may take place on multiple peers running on the same hosts. Given the number of hosts $N = 76$ we can use, this constrains us on the size of the swarm we investigate. However, we prefer to take a cautious approach, and avoid to introduce the aforementioned mutual influence that could bias in an unpredictable way the results of our experiments (see a discussion on the conclusions).

We repeat the experiments three times for each settings, to smooth out stochastic variability in the experiments due to BitTorrent random decisions (e.g., chunk selection, choke, optimistic unchoking etc). Also, we make results of our campaign available to the scientific community as for the local testbed at [23]. Overall, the volume of collected data in the Grid’5000 testbed amounts to about 10 GBytes. Yet, we point out that, in reason of the large number of experiments and seeds, we were not able to store packet-level traces, but only periodic transport-layer (i.e., TCP, UDP traffic amount), application layer (i.e., tracker) and queuing logs.

4.2 Homogeneous *bt.transp_disposition* settings

Results of the homogeneous case are reported in Fig. 2. For each metric of interest, the figure reports the *envelope* of the gathered results, i.e., the minimum and maximum curves over the three iterations.

We express results in terms of (i) the cumulative distribution function (CDF) of the torrent completion time T , on the right plots and of (ii) the complementary cumulative distribution function (CCDF) of the buffer occupancy Q of the access link of each peer, on the left plots. The buffer occupancy is expressed both in bytes (bottom x-axis) and in terms of the amount of delay an interactive application would experience for the emulated access capacity (top x-axis).

Additional details are reported in the inset of each figure, showing: (iii) the average $E[T]$ and standard deviation $\sigma[T]$ of the torrent completion time; (iv) the byte-wise share between TCP and uTP, with a notation X/Y that specifies that $X\%$ ($Y\%$) of the bytes are carried over TCP (uTP), with $X + Y = 100\%$; (v) the mean queue size $E[Q]$ in KBytes and milliseconds.

In the top row we report the scenario where all peers use default settings (*bt.transp_disposition*=31), i.e., the peers are able to speak both TCP and uTP protocols. Middle plots report the case of an all-uTP swarm (*bt.transp_disposition*=10), while all-TCP swarms (*bt.transp_disposition*=5) are depicted in the bottom row.

Notice that, on the long run, all swarms achieves similar efficiency: looking at the CDF of the buffer occupancy in Fig. 2 we can see that in roughly 80% of the time, after a dequeue operation the queue is non-empty. That efficiency is also tied to the BitTorrent system dynamics (e.g., pipelining of the requests, chunk exchange dynamics, etc.). Also the number of packets remaining in the queue after a packet transmission further depends on the congestion control protocol of choice. As expected, TCP AIMD dynamics tend to fill the buffer, while uTP strives (and manages) to limit the queue size.




Fig. 2. Buffer occupancy CCDF (left) and Completion time CDF (right) for homogeneous swarms: default settings (*bt.transp_disposition*=31, top), uTP only swarm (*bt.transp_disposition*=10, center) and TCP only swarms (*bt.transp_disposition*=5, bottom). The vertical line in the Buffer occupancy plot represent the average of the queue length (in KB and ms).

These behaviors translate into different completion times statistics and, especially, completion times appear to benefit from a mixture of TCP and uTP traffic. We point out that, in the mixed case where BitTorrent peers are able to speak both protocols (*bt.transp_disposition*=31), the following happens: two connections, a TCP and an uTP ones are attempted, and in case uTP is successfully opened, it is preferred over the TCP one. This translates into a traffic mixture where about 80% of the data traffic happens to be carried over uTP.

Notice that the queue size alone cannot explain the difference in the completion time statistics (as otherwise, completion time in all-uTP swarms will be the lowest). Hence, we conjecture this result to be the combination of two effects –on the control

and data plane— that are assisted by the use of uTP and TCP respectively. First, a longer queue size due to TCP can negatively influence the completion time, by hindering a timely dissemination of *control information* (e.g., chunk interest). The longer the time needed to signal out interests, the longer the time prior to start their download, and their subsequent upload to other peers (which harms all-TCP completion time).

Notice indeed that as in the all-TCP case the one way queuing delay may reach up to 400 ms on average, this entails that RTT for signaling exchanges may be on the order of a second, that can possibly slow down significantly the chunk spreading dynamics. Consider then that BitTorrent is using pipelining to avoid a slowdown of the transfer due to the propagation delay of requests for new chunks. From our experiments, it appears that the pipelining used by BitTorrent is not large enough to deal with delays that might be encountered with xDSL connections.

However, as previously said, the completion times statistics are not fully explained in terms of the queuing delay, as otherwise all-uTP swarms should be the winner. Yet, while uTP limits the queue size and avoids to interfere with a timely dissemination of control messages, uTP is also by design less aggressive than TCP. It follows that TCP may be more efficient for rapidly sharing *chunks in the data plane*. This can in turn harm the all-uTP completion time, that is slightly larger with respect to the default settings `bt.transp_disposition=31`.

Interestingly, our previous simulation study [18] shown that a combination of TCP and uTP can increase the efficiency on the case of two flows sharing a bottleneck link. Shortly, this happens because the low-priority protocol is still able to exploit the capacity unused by TCP (as its rate increases when queuing is low), without at the same time increasing the average system queuing delay (as its rate slow down when TCP traffic increases). The experimental results of this work further confirm that a combination of TCP and uTP can be beneficial to the completion time of the whole swarm as well. Moreover, although the exact shape of the completion time CDFs differ across experiments (due to the stochastic nature of BitTorrent chunk scheduling and peer selection decisions), the results are consistent across all iterations.

Unfortunately, latest versions of uTorrent do not allow to export chunk level logs, which could bring further information as the trading dynamics between peers, that remains an interesting direction for future work.

4.3 Heterogeneous `bt.transp_disposition` settings

Having seen that a mixture of TCP and uTP protocols can be beneficial to the completion time, we further investigate different shares of TCP (`bt.transp_disposition=13`) vs uTP peers (`bt.transp_disposition=14`), i.e., peers that prefer one of the two protocols for active connection open, but that can otherwise accept any incoming connections.

We consider three peer-wise shares, namely 25/75, 50/50, and 75/25 (in the X/Y notation, $X\%$ represents the percentage of peers preferring TCP on their uplink, i.e., `bt.transp_disposition=13`). These shares represent three different popularity cases of uTP, that can be either the default in only few implementations of the BitTorrent applications (25/75), or compete equally (50/50) or even be dominant (75/25). We believe these shares to represent illustrative points, covering all relevant scenarios of the possible population repartition.

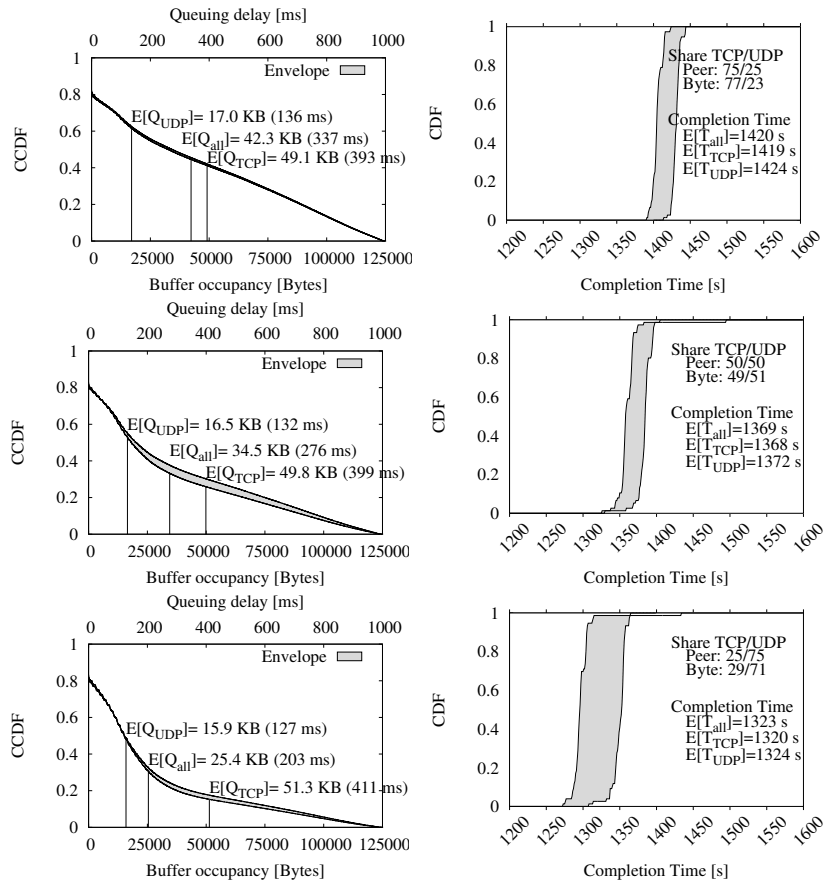


Fig. 3. Buffer occupancy CCDF (left) and Completion time CDF (right) for heterogeneous swarms: prevalence of TCP peers (75/25, top), fair population share (50/50 middle), and prevalence of uTP peers (25/75, bottom).

The plots in Fig. 3 additionally report (i) the average queuing delay, for all the swarm as well as for different peer classes, (ii) the peer- and byte-wise traffic shares, and (iii) the average system completion time, as well as the average completion time for peers of different classes.

As for (i), the average queuing delay statistics are as expected, with an increase of the queuing delay of uTP peers due to bursty acknowledgements in reply to TCP traffic due to TCP peers in the reverse path. As for (ii), the byte-wise share closely follow the peer-wise share. Finally, let us focus on (iii) the completion time statistics. Interestingly, as Fig. 3 shows, while a small amount of TCP traffic is beneficial in reducing the overall swarm completion time (bottom row), a large TCP amount can instead slow down the torrent download for the whole system (top row).




Fig. 4. Completion time as a function of the TCP vs uTP byte-wise share (left) and as a function of the average buffer occupancy (right).

Further, notice that completion times are practically the same for uTP and TCP peers (with a slight advantage for the latter). Hence, differently from our previous simulation work [9], we do not observe an unfairness of completion time between different peer classes within an heterogeneous swarm. This is due to the fact that [9] considered a simpler model for *bt.transp_disposition*, that neither (i) accounted for TCP peers using an already opened uTP connection in the reverse side nor (ii) for the hardcoded uTP preference.

4.4 uTP vs TCP in a nutshell

Fig. 4 present a summary of our results considered so far. T is the completion time (mean and standard deviation) for different iterations with both homogeneous and heterogeneous swarm populations. The T metric is reported as a function of the byte-wise TCP traffic share (left plot) and of the average buffer occupancy (right plot).

Both plots also report, for TCP traffic shares different from zero (non-0 TCP) only, a linear regression of the completion time. Notice that the linear model provides a nice fit to forecast the completion time performance in presence of different TCP vs uTP mixtures.

Furthermore, as observed in [9] by means of simulation, Fig. 4 confirms that the *completion time increases for increasing buffer occupancy*, which in turn generally increases with the amount of TCP traffic exchanged.

As previously argued, this is due to a slow-down of BitTorrent signaling traffic, while the completion time increase of all-uTP swarms is instead likely due to the low-priority of uTP in the data plane. Hence, we also remark a non-monotonous behavior for the completion time, that decreases for decreasing percentages of TCP traffic, and then increases again for all-uTP swarms. As different dynamics takes place, hence the linear dependence only applies in case of uTP and TCP traffic mixtures (i.e., non-0 TCP traffic share).

Finally, notice that the default BitTorrent settings consistently yield to the shortest download time we observed in the experiments, which confirms the soundness of the *bt.transp_disposition* design decision and settings.

5 Conclusions

This work assess the impact of uTP (the new BitTorrent congestion control algorithm for data exchange) on the torrent completion time (the main user QoE metrics) by means of an experimental campaign carried on in a fairly large scale controlled testbed.

Our results show that, in flash crowd scenarios, users will generally benefit of a mixture of TCP and uTP traffic, both in homogeneous and heterogeneous swarms. Interestingly indeed, results with mixed TCP and uTP traffic show consistently shorter download time with respect to the case of homogeneous swarms using either an all-TCP or an all-uTP congestion control. Especially, our results confirm the soundness of default BitTorrent settings, which use both TCP and uTP protocols and lead to the shortest completion time in our experiments.

This results is the combination of two effects, on the control and data plane, that are assisted by the use of uTP and TCP respectively. By keeping the queue size low, uTP yields to a timely dissemination of signaling information, that would otherwise incur in higher delays due to longer queues building up with TCP. At the same time, by its more aggressive behavior, TCP yields to higher efficiency in the data plane, that results in more timely dissemination of chunk content.

This work leaves a number of interesting points open, that we aim at addressing in the future. First, we would like to investigate whether it would be possible to extend the swarm size by running multiple peers per machine, without however incurring in a bias due to the mutual interaction of the traffic injected. Second, we would like to study the impact of heterogeneous target values in the swarm (i.e., to see whether fairness issues can possibly arise) and refine our experimental setup (e.g., including heterogeneous capacities, peer churn, etc.). Third, on a longer timescale, we aim at developing a passive BitTorrent inspector, capable of parsing traffic to produce chunk-level logs, that would greatly enhance our analysis capabilities.

Acknowledgement

We thank Arvid Norberg for the fruitful discussions. Experiments presented in this paper were carried out using the Grid'5000 experimental testbed, being developed under the INRIA ALADDIN development action with support from CNRS, RENATER and several Universities as well as other funding bodies [10]. Project funding for our future work on the topic would be more than welcome.

References

1. Alcock, S., Nelson, R.: Application flow control in youtube video streams. *ACM SIGCOMM Computer Communication Review* **41**(2) (2011) 24–30
2. Bonfiglio, D., Mellia, M., Meo, M., D.Rossi: Detailed analysis of skype traffic. *IEEE Transaction on Multimedia* **11**(1) (January 2009) 117–127
3. Shalunov, S.e.a.: Low Extra Delay Background Transport (LEDBAT). IETF Draft (Oct 2010)
4. Norberg, A.: BitTorrent Enhancement Proposals on μ Torrent transport protocol. http://www.bittorrent.org/beps/bep_0029.html (2009)

5. IETF: LEDBAT Working Group Charter. <http://datatracker.ietf.org/wg/ledbat/charter/>
6. Schneider, J., Wagner, J., Winter, R., Kolbe, H.: Out of my Way – Evaluating Low Extra Delay Background Transport in an ADSL Access Network. In: 22nd International Teletraffic Congress (ITC22), Amsterdam, The Netherlands (Sep 2010)
7. Rossi, D., Testa, C., Valenti, S.: Yes, we LEDBAT: Playing with the new BitTorrent congestion control algorithm. In: Passive and Active Measurement (PAM 2010), Zurich, Switzerland (Apr 2010)
8. Morris, S.: μ Torrent release 1.9 alpha 13485. <http://forum.utorrent.com/viewtopic.php?pid=379206#p379206> (Dec 2008)
9. Testa, C., Rossi, D.: The impact of utp on bittorrent completion time. In: 11th IEEE International Conference on Peer-to-Peer Computing (P2P), Kyoto, Japan (2011)
10. <https://www.grid5000.fr>
11. Venkataramani, A., Kokku, R., Dahlin, M.: TCP Nice: A mechanism for background transfers. In: 8th USENIX Symposium on Operating Systems Design and Implementation (OSDI 2002), Boston, MA (Dec 2002)
12. Kuzmanovic, A., Knightly, E.: TCP-LP: low-priority service via end-point congestion control. *IEEE/ACM Transactions on Networking (TON)* **14**(4) (2006) 752
13. Liu, S., Vojnovic, M., Gunawardena, D.: 4cp: Competitive and considerate congestion control protocol. In: ACM SIGCOMM, Pisa, Italy (Sep 2006)
14. Qiu, D., Srikant, R.: Modeling and performance analysis of BitTorrent-like peer-to-peer networks. *ACM SIGCOMM Comp. Comm. Rev.* **34**(4) (2004) 367–378
15. Rao, A., Legout, A., Dabbous, W.: Can realistic bittorrent experiments be performed on clusters? In: 10th IEEE International Conference on Peer-to-Peer Computing (P2P). (aug. 2010) 1–10
16. Legout, A., Liogkas, N., Kohler, E., Zhang, L.: Clustering and sharing incentives in bittorrent systems. In: Proc. of ACM SIGMETRICS’07, San Diego, CA, USA (June 2007)
17. Cohen, B., Norberg, A.: Correcting for clock drift in uTP and LEDBAT. In: Invited talk at 9th USENIX International Workshop on Peer-to-Peer Systems (IPTPS 2010), San Jose, CA (Apr 2010)
18. Rossi, D., Testa, C., Valenti, S., Muscariello, L.: LEDBAT: the new BitTorrent congestion control protocol. In: 19th IEEE International Conference on Computer Communications and Networks (ICCCN 2010), Zurich, Switzerland (Aug 2010)
19. Carofiglio, G., Muscariello, L., Rossi, D., Valenti, S.: The quest for LEDBAT fairness. In: IEEE Global Communication (GLOBECOM 2010), Miami, FL (Dec 2010)
20. Carofiglio, G., Muscariello, L., Rossi, D., Testa, C.: A hands-on Assessment of Transport Protocols with Lower than Best Effort Priority. In: 35th IEEE Local Computer Network (LCN 2010), Denver, CO (Oct 2010)
21. Padma Bhooma: Measurements using the LEDBAT implementation for MAX OS X. <http://www.ietf.org/mail-archive/web/ledbat/current/msg00532.html>
22. Abu, A., Gordon, S.: A Dynamic Algorithm for Stabilising LEDBAT Congestion Window. In: 2nd IEEE International Conference on Computer and Network Technology (ICCNT 2010), Bangkok, Thailand (Apr 2010)
23. <http://perso.telecom-paristech.fr/~testa/pmwiki/pmwiki.php?n=Main.BTtestbed>
24. Mogul, J.C.: Tcp offload is a dumb idea whose time has come. In: 9th conference on Hot Topics in Operating Systems - Volume 9, Berkeley, CA, USA, USENIX Association (2003) 5–5
25. Akella, A., Seshan, S., Shaikh, A.: An empirical evaluation of wide-area internet bottlenecks. In: Proc. of the 3rd ACM SIGCOMM Conference on Internet Measurement (IMC’03), Miami, FL, USA (Oct. 2003)