

On a Balanced Neighbor Selection Strategy for Tracker-Based Peer-to-Peer Networks

Sándor Laki and Tamás Lukovszki

Faculty of Informatics, Eötvös Loránd University
Pázmány Péter sétány 1/C, H-1117 Budapest, Hungary
Email: {lakis,lukovszki}@inf.elte.hu

Abstract—In this paper we introduce a novel neighbor selection strategy for tracker-based peer-to-peer systems like BitTorrent that can uniformly distribute the load among peers in the network. Our method is based on a balanced multiple choice algorithm which takes into account not only the actual load of a peer, but the possibility as well that it will be selected in the future. We first analyze the properties of the constructed overlay topology theoretically, proving that, the maximum degree in the constructed graph is $O(1)$ while the diameter remains $O(\ln n)$, with high probability, where n is the number of nodes. Considering a randomized upload policy, we show that the full distribution of b blocks on the network generated by our neighbor selection strategy takes $O(b + \ln n)$ phases only, with high probability, which is optimal up to a constant factor. This result improves the previous upper bound of $O(b + (\ln n)^2)$ by Arthur and Panigrahy (SODA'06). In order to adapt our algorithm in real BitTorrent networks only a slight modification of the tracker is necessary without any change in the clients. Besides theoretical analysis, thorough simulations have been done to validate our algorithm and show its applicability in real BitTorrent networks. To this end, we have extended the BitTorrent implementation of the PeerSim simulation framework with a new tracker using our balanced neighbor selection strategy and demonstrated that it can speed up the file-sharing process in heavy loaded situations.

I. INTRODUCTION

In the past decade, tracker-based peer-to-peer networks like BitTorrent (BT) [1] and Tribler [2] have emerged as popular solutions in the area of not only simple file-sharing, but video-on-demand services as well. These applications are still showing an increasing interest, generating a significant part of the overall Internet traffic.

The selection of neighbors is an important design decision of peer-to-peer systems. In tracker-based peer-to-peer networks, each peer that enters the network, first has to connect to a central component called tracker to obtain a peer set representing the initial neighborhood of the joining client. The tracker maintains a list of all nodes in the system, called the swarm, and returns a random subset of the existing nodes. This random neighbor selection may lead to suboptimal overlay topologies. In order to optimize the network, various neighbor selection strategies can be found in the literature that considers different aspects from locality [16] and load balancing [3], [9] to quality of experience [14].

The performance of BT like peer-to-peer systems has been widely analyzed in the past few years from theoretical and practical aspects as well. The empirical results [10], [5] show that the simple routing policy applied by the original BT is

quite effective even in case of a flash crowd setting when a great deal of peers join the network almost at the same time. Besides empirical evidences, this heavy loaded case has already been investigated from a theoretical perspective as well. In [3], several algorithms are demonstrated that share b data blocks among n clients in a network of diameter d and degree D in $O(D(b+d))$ steps with high probability¹, where in one time step, each client can upload one data block to, and download one block from one of its neighbor. For a network used by BT it results in a time bound of $O(b \ln n)$ time steps. They propose a neighbor selection strategy which improves this bound to a near-optimal $O(b + (\ln n)^2)$ steps.

In this paper, we improve the neighbor selection strategy resulting in a time bound of $O(b + \ln n)$ steps, which is optimal up to a constant factor. Our method uses the idea of multiple choice [4] and takes into account not only the actual load of the peers, but the possibility as well that a client will be selected in the future. This will ensure an overlay network of constant degree and logarithmic diameter with high probability. The constructed overlay topologies are examined from both theoretical and practical perspectives as well. We model these overlay networks as a graph, whose vertices are the peers and neighboring peers are connected by an edge. We analyze the key graph properties of the proposed network, showing that the maximum degree in our overlay topology is $O(1)$, with high probability, while its diameter still remains logarithmic in the number of peers n . In such a network, the randomized upload policy will share b data blocks among n clients in $O(b + \ln n)$ time steps with high probability, which is optimal in networks of n vertices, in which the degree of the vertices is bounded by a constant. Besides the theoretical analysis thorough simulations have been performed to validate the different properties of the constructed overlay networks. In addition to the analysis of the degree and diameter distributions, we also show how our balanced algorithm could work in a tracker-based file sharing system such as BitTorrent to accelerate the delivery of the data blocks in the whole network. To this end, we have implemented a modified tracker for the BT simulation of the PeerSim peer-to-peer simulation framework [11], and carried out experiments to validate that our balanced overlay construction could work not only in

¹An event E is said to occur with high probability, if given $n > 1$, $Pr[E] > 1 - 1/n^c$, where $c > 1$ is a constant [12].

theory, but in practice as well.

The rest of the paper is organized as follows: in Section II, we briefly overview the related works. In Section III, we describe our model and the different neighbor selection methods. Section IV details our theoretical analysis and includes the proof of our theoretical upper bounds. In Section V, we experimentally analyze the different properties of the constructed overlay networks and also show how the proposed method can improve the performance of file sharing in case of a flash crowd setting. Finally, we include some concluding remarks in Section VI.

II. RELATED WORKS

A great deal of theoretical and empirical studies have emerged in the past decade to analyze the performance of existing BT like peer-to-peer networks and propose new neighbor selection methods to optimize the constructed overlay topology. At the empirical end, Izal et al. [10] and Pouwelse et al. [15] present measurement based studies of BT which are based on tracker logs of different torrents. Their analysis shows that the simple mechanisms that can be found in the original BT makes this file-sharing system very efficient.

Bharambe et al. [6] conducted simulations to confirm that BT performs near-optimally in terms of uplink bandwidth utilization, and download time except under certain extreme conditions. They have also found that the rate-based tit-for-tat policy is not effective in preventing unfairness, which means that low bandwidth peers can download more than they upload to the network when high bandwidth peers are present. To solve this issue, they propose some slight changes to the tracker and a stricter tit-for-tat policy.

Bindal et al. [7] examine a new approach to enhance BT traffic locality, in which a peer chooses the majority, but not all, of its neighbors from peers within the same ISP. In this way the traffic costs at ISPs can significantly be reduced.

Besides empirical works there are several theoretical ones as well. Zhang et al. [17] formulate an optimization problem to solve for the optimal peer selection strategy to maximize the global system-wide performance. They also derive a purely distributed algorithm that is provably globally optimal. Besides the tracker, the proposed solution requires some changes in the ordinary peers as well. We also mention the classical graph theoretic results by Bollobás and Fernandez de la Vega [8]. They proved that, for $r \geq 3$, almost every random r -regular graph has logarithmic diameter.

Arthur and Panigrahy [3] propose a mathematical framework to model the distribution of individual data blocks. They examine several properties of BT like networks and discuss a number of extensions to them, including a new neighbor selection strategy that can easily be implemented in the trackers to achieve near-optimal performance in the distribution of data blocks. In our work, we use the same analytical framework introduced in [3].

III. SYSTEM MODEL

In this section, we briefly outline our system model, which is in accordance with what is proposed by Arthur and Panigrahy

in [3], and describe the neighbor selection strategies to be examined.

For the sake of simplicity, we model the constructed overlay topologies as directed graphs where each vertex represents a client in the peer-to-peer network. We also assume equal bandwidth and delay among all the peers, and ignore other BT specific mechanisms such as tit-for-tat and optimistic unchoke. This simplified model does not take care of the process of clients joining and leaving the network during the file-sharing. Furthermore, the file sharing process can be considered as routing data blocks on the directed graph over discrete time steps. To this end, Arthur and Panigrahy [3] proposed the randomized upload policy where each vertex attempts to upload a block to a random neighboring client during each time step. They proved the following theorem.

Theorem 3.1 ([3]): Suppose a vertex u begins with a copy of every block. Let D denote the maximum out-degree in a directed graph consisting of n vertices, and suppose the distance from u to every other vertex is at most d . If we route on this graph using the randomized upload policy, then $T \leq 4D(4d + b)$ with probability at least $1 - 2n \exp(-\frac{d}{2})$, where b denotes the number of distinct blocks to be distributed and T is the number of time steps before the routing completes.

We introduce a neighbor selection strategy, which results in a network of constant degree and logarithmic diameter. Applying Theorem 3.1 to this network we obtain a bound of $O(b + \ln n)$ time steps for the completion of the routing of the b blocks, with high probability. Note that every graph of n vertices with constant degree has a diameter $\Omega(\ln n)$. Thus the routing of a block to the farthest vertex takes $\Omega(\ln n)$ steps. Since it can receive one block in a step receiving all the blocks takes $\Omega(b + \ln n)$ steps. Thus, our result is optimal up to a constant factor.

As we mentioned before, in BT each joining client first sends a request to the tracker that returns a peer set chosen uniformly at random among the existing peers. These nodes form the initial neighborhood of the new client. To model such a peer-to-peer network, Arthur and Panigrahy [3] propose the BitTorrent- C graph ($C \geq 2, C \in \mathbb{N}$) which is a directed graph constructed by the following method.

Definition 3.2: BitTorrent- C graph (abbr. BT- C graph) [3]:

- 1) At the beginning it consists of C vertices, v_1, \dots, v_C and edges from v_j to v_i if and only if $j < i$.
- 2) While the total number of vertices is less than n , add a vertex and add directed edges from C existing vertices chosen uniformly at random to the new vertex.

This graph has been analyzed thoroughly in [3]. It has been shown that, with high probability, the maximum out-degree in a BT- C graph is at most $3C(1 + \ln n)$, while the diameter is at most $3 \lg n$. Furthermore, based on Theorem 3.1 they also proved that the required time steps for distributing b blocks with the randomized upload policy can be bounded by $O(\ln n(b + \ln n))$, with high probability.

To improve the performance of the original BT, especially in case of flash crowd settings, they recommend a practical

variant of BT- C , called Smoothed-BT- C graph that can be constructed as follows.

Definition 3.3: Smoothed-BitTorrent- C graph (Smoothed-BT- C graph) [3]:

- 1) At the beginning it consists of C vertices, v_1, \dots, v_C and edges from v_j to v_i if and only if $j < i$.
- 2) While the total number of vertices is less than n , add a vertex and add directed edges from C existing vertices to the new vertex, but instead of choosing each previous vertex uniformly at random, select two previous nodes and connect the one with higher index to the new vertex.

In [3] it has been proved that by using the randomized upload policy in a Smoothed-BT- C graph with n vertices, the routing completes in at most $O(b + (\ln n)^2)$ time steps, with high probability.

In this paper, we propose two novel neighbor selection strategies that use the idea of multiple choice to ensure an overlay network of constant out-degree, with high probability. We first introduce the MultipleChoice-BT- C graph that can be built up in the following manner.

Definition 3.4: MultipleChoice-BitTorrent- C graph.

- 1) At the beginning it consists of C vertices, v_1, \dots, v_C and edges from v_j to v_i if and only if $j < i$.
- 2) We add the remaining vertices one by one. Let $t \geq 2, t \in \mathbb{R}$, be a constant. When we add the i th vertex, $C < i \leq n$, we choose $t \log n$ vertices from the vertex set $\{v_j : j \in [\min(i/2, i - C), i - 1]\}$ uniformly at random and connect the lowest degree vertex to the new vertex. We repeat this C times to add C directed edges to the new vertex.

Estimating $\log n$ can be done in various ways, some simple methods are discussed e.g. in [13]. We will show in Section IV that the multiple choice policy guarantees constant out-degrees with high probability. The disadvantage of this strategy is that at the insertion of the i th vertex it tends to select the neighbors from the last $\frac{1}{2t \log n}$ fraction of the previous vertices, since the newer vertices have more likely a lower degree. The consequence of this will be a super-logarithmic diameter.

To remedy this problem, we present a modification of this multiple choice neighbor selection policy, which guaranties an overlay topology with $O(C)$ maximal degree and logarithmic diameter with high probability. This improved variant of BT- C is called Balanced-BT- C graph. It is similar to the MultipleChoice-BT- C graph with the following modification. Instead of choosing the vertex with the smallest out-degree from $t \log n$ previous vertices, we also take into account the expected number of times it will be selected in the future. Before the insertion of the i -th node, for a vertex $v_j, j < i$, let $\delta(v_j)$ denote the out-degree of v_j and $w(v_j) := \sum_{i < t \leq 2j} \frac{2C}{t}$. The value $w(v_j) \cdot t \log n$ expresses the expected number of times v_j will be chosen after the insertion of v_i . For $j < i/2$, $w(v_j) = 0$ and for $j \geq i$, $w(v_j)$ is the expected number of future edges. Let $\delta^*(v_j) := \delta(v_j) + w(v_j)$

Definition 3.5: Balanced-BitTorrent- C graph.

- 1) At the beginning the graph consists of C vertices, v_1, \dots, v_C and edges from v_j to v_i if and only if $j < i$.

- 2) We add the remaining vertices one by one. Let $t \geq 2, t \in \mathbb{R}$, be a constant. When we add the i th vertex, $C < i \leq n$, we choose $t \log n$ vertex from the vertex the set $\{v_j : j \in [\min(i/2, i - C), i - 1]\}$ uniformly at random and connect the vertex v_j with the lowest $\delta(v_j) + w(v_j)$ to the new vertex. We repeat this C times in order to add C directed edges to the new vertex.

Using this neighbor selection strategy ensures that at the moment we insert the i th vertex, the expected values of $\delta^*(v_j), j \in [i/2, i - 1]$ are approximately the same. Thus the selected neighbor is distributed evenly among $\{v_j : j \in [i/2, i - 1]\}$. As a consequence, the median neighbor is selected expectedly from the middle of the interval $[i/2, i - 1]$, which results a network with logarithmic diameter.

IV. THEORETICAL ANALYSIS

In order to give a bound on the diameter of the network, the term of median depth has been introduced in [3], and defined as follows.

Definition 4.1: Median depth: For any $i \in \mathbb{N}$, consider the set of integers j such that there is an edge from v_j to v_i . Let $m(i)$ denote a median element in this set. Recursively, let $m^k(i) = m(m^{k-1}(i))$ for $k > 0$ and $m^0(i) = i$. The median depth of v_i is defined as the smallest k such that $m^k(i) = 1$.

Clearly, the distance from v_1 to v_i is at most the median depth of v_i .

Lemma 4.2: For $t \geq 2, t \in \mathbb{R}$, the maximum out-degree in a MultipleChoice-BT- C graph is at most $2C$ with probability of at least $1 - \frac{C}{n}$.

Proof: Consider the insertion of the i th vertex. To create a maximum out-degree greater than $2C$, at least one edge to be inserted from a previous vertex with out-degree of at least $2C$. The neighborhood of the i th vertex consists of peers from the interval $I = [\frac{i}{2}, i - 1]$. We also know that peers from I may previously have been selected by at most $\frac{i}{2} - 1$ vertices as neighbors, so the number of edges originated from I is at most $C(\frac{i}{2} - 1) < C\frac{i}{2}$. Consequently, the interval contains less than $\frac{i}{4}$ peers with out-degree $2C$. Thus the probability that we select a vertex from I having out-degree $2C$ is at most $\frac{1}{2}$. The probability that all the $t \log n$ vertices have out-degree $2C$ is at most $(\frac{1}{2})^{t \log n} = \frac{1}{n^t}$. Since all the vertices have C ingress edges, the probability that after inserting the i th vertex there is a peer with out-degree $2C + 1$ is at most $\frac{C}{n^t}$. Since $t \geq 2$ using the union bound over the error probabilities of all the vertices proves the lemma: $\sum_{i=1}^n \frac{C}{n^t} \leq \frac{C}{n^{t-1}} \leq \frac{C}{n}$. ■

Note that a linear increase in t leads to an exponential decrease of the failure probability.

Although this neighbor selection strategy guarantees a constant degree with high probability, the median depth of the resulting graph will be super-logarithmic. The reason is that newer vertices tend to have lower degree and the neighbors of the new vertex are preferentially selected among the last $\frac{1}{2t \log n}$ fraction of the previous vertices. When we insert the n th vertex, a vertex $v_j, j \in [n/2, n - 1]$ has been chosen before the n th insertion expectedly $t \log n \sum_{j < i < n} \frac{2}{i}$ times.

Thus newer vertices have been chosen less frequently and more likely have lower degree.

Lemma 4.3: For $t \geq 2, t \in \mathbb{R}$, the maximum out-degree in a Balanced-BT- C graph is at most $6C$ with probability of at least $1 - \frac{C}{n}$.

Proof: Consider the insertion of the i th vertex. The neighborhood of the i th vertex consists of peers from the interval $I = [\frac{i}{2}, i - 1]$. We prove that after the insertion of the i th vertex, for each $v_j, j \in I, \delta^*(v_j) \leq 6C$, with high probability. It will imply the claim of the lemma. To create a vertex v_j with $\delta^*(v_j) > 6C$, all of the $t \log n$ chosen vertices must have a $\delta^*(\cdot)$ value of at least $6C$. We know that peers from I may previously have been selected by at most $\frac{i}{2} - 1$ vertices as neighbors, so the number of edges originated from I is at most $C(\frac{i}{2} - 1) < C\frac{i}{2}$. Furthermore, we know that, for each $j \in I, w(v_j) = \sum_{i < \ell \leq 2j} \frac{2C}{\ell} \leq (2j - i)\frac{2C}{i} \leq 2C$. Since I contains $\frac{i}{2}$ vertices, $\sum_{j \in I} w(v_j) \leq Ci$. Therefore, $\sum_{j \in I} \delta^*(v_j) < C\frac{3i}{2}$. Consequently, less than $\frac{i}{4}$ peers with a $\delta^*(\cdot)$ value $6C$. Thus the probability that we select a vertex from I having a $\delta^*(\cdot)$ value $6C$ is at most $\frac{1}{2}$. The probability that all the $t \log n$ vertices have a $\delta^*(\cdot)$ value $6C$ is at most $(\frac{1}{2})^{t \log(n)} = \frac{1}{n^t}$. Since all vertices have C ingress edges, the probability that, after inserting the i th vertex, there is a peer with a $\delta^*(\cdot)$ value greater than $6C$ is at most $\frac{C}{n^t}$. Since $t \geq 2$ using the union bound over the error probabilities of all the vertices proves the lemma: $\sum_{i=1}^n \frac{C}{n^t} \leq \frac{C}{n^{t-1}} \leq \frac{C}{n}$. ■

One can observe that our balanced neighbor selection strategy ensures that at the moment we insert the i th vertex, the expected values of $\delta^*(v_j), j \in [i/2, i - 1]$ are approximately the same, resulting that the selected neighbor is distributed evenly among $\{v_j : j \in [i/2, i - 1]\}$. In this case, the median neighbor is expectedly from the middle of the interval $[i/2, i - 1]$, which leads to a network with logarithmic median-depth.

Applying Theorem 3.1, it can be stated that in a Balanced-BT- C graph with n vertices, routing b data blocks with the randomized upload policy completes in at most $O(b + \ln n)$ time steps, with high probability, which is provable optimal up to a constant factor.

V. EXPERIMENTAL RESULTS

In order to adapt our theoretical result in real BT networks only a slight modification of the tracker is necessary without any change in the clients. Simulations have been performed to experimentally examine the properties of the overlay networks constructed by the different neighbor selection strategies. In addition to the analysis of the degree and median-depth distributions, we also show how our balanced algorithm could work in a tracker-based file sharing system such as BT to accelerate the delivery of the data blocks in the whole network. To this end, we have implemented a modified tracker for the BT simulation of PeerSim [11], and carried out experiments to validate that our balanced overlay construction could work not only in theory but in practice as well.

1) *Properties of BT- C Overlay Networks:* Simulations have been done to analyze the key properties of the overlay networks generated by the proposed methods. In the previous section, high probability upper bounds have been proved for the maximum out-degree and the diameter of the different networks. In addition, simulations enable us not only to validate these theoretical results but to reveal more information on the distributions themselves and explore the practical strength of the above bounds.

The parameters of the constructed overlay networks have been varying in a reasonable wide range: C was chosen from the range $[2, 90]$, while n from $[500, 100000]$. Due to page limitation, results for $C = 2$ and 30 are only presented in this paper. $C = 2$ demonstrates well the behavior of different strategies for constant size neighbor set, while $C = 30$ corresponds to the neighbor set whose size is rather logarithmic in the number of peers in our simulations. For each setting, the simulations have been repeated 100 times to obtain statistically enough data for further analysis. The constant t in our multiple choice methods has been fixed to 2 during the analysis.

First, we pay attention to the maximum out-degrees observed in the different networks. In a file sharing network, the out-degree of a peer determines the maximum load on it, indicating the maximum number of other nodes that the given peer can upload data to. Figure 1 illustrates the maximum out-degrees for two different C values. Each dot in these figures indicates the maximum out-degree observed in a given experiment. Since for each setting and algorithm, the simulations have been repeated 100 times, we can see 100 individual dots for each n value and graph construction. The fitted curves for the different networks are marked by different colors, expressing the connection between the maximum out-degree values and the number of clients. Looking at Figure 1(a), one can observe that, not surprisingly, the original BT- C results the highest maximum out-degrees and shows a logarithmic correlation between the maximum values and the network size. In case of a such small C , Smoothed-BT- C behaves very similarly to the previous overlay construction. Although it provides less maximum out-degrees for large n values, it is almost the half of what we can see in BT- C , but the above relationship is still logarithmic. In accordance with the theory, both our MultipleChoice-BT- C and Balanced-BT- C algorithms aim at keeping the out-degrees at a constant level. For $C = 2$, the maximum values are 3 and 4, respectively. These values correspond to the theoretical high probability upper bounds. Considering larger neighborhood sizes in Figure 1(b), where C is 30, similar correlations can be identified, but there are some slight differences we have to shed light on. First of all, for large peer set sizes (C), the maximum out-degrees in the networks constructed by MultipleChoice-BT- C and Balanced-BT- C never reach the theoretical upper bound and, as it is expected, the former method results in a slightly lower load level on the peers than the latter one. From a practical perspective, when the network size is within a reasonable range (e.g. less than 1 million), the out-degrees in a Smoothed-BT- C graph are only slightly influenced by the




Fig. 1. Maximum out-degrees for the different overlay networks. Each dot represents a maximum out-degree value after the given simulation has been done. For each parameterization and algorithm, the simulations have been repeated 100 times. The fitted curves/lines indicate the relationship between the maximum out-degrees and the number of clients in the network.

number of peers.

Besides the analysis of the maximum values, our experiments enable us to examine the out-degree distributions themselves. Figure 2 shows the complementary cumulative distribution function (CCDF) of the out-degrees for the different algorithms on a semi-log plot. First, we consider $C = 2$, in case of BT-C the out-degrees seem to follow an exponential decay, but for larger C values the results show slightly better load distribution, for the tail of the distribution decays sharply after a certain point. For Smoothed-BT-C, the figures indicate that the CCDF of the out-degrees decreases faster than exponential, for both C values. Not surprisingly, our MultipleChoice-BT-C and Balanced-BT-C methods result in a sudden drop in the CCDF plots at a constant value which is less than or equal to $2C$, indicating that the majority of the nodes have the same constant out-degree, and the load of the peers are much more balanced.

Besides the maximum out-degree, the diameter of the overlay topology also plays a crucial role in data distribution. For example, if we consider a peer-to-peer network containing only one seeder node having the whole file at the beginning, the data blocks need to propagate through the whole network to reach all the peers. It has already been shown in Section IV that besides the number of blocks to be distributed and the maximum out-degree, the diameter of the network has also significant effect on the time required for spreading the blocks of a given file.

Figure 3 depicts the relationship between the maximum median-depth and the number of peers in the different overlay networks. Each dot in the figures represents an individual experiment where the network size can be seen on the horizontal axis, while the observed maximum median-depth on the vertical ones. The fitted curves show the trend of this relationship for the different methods. Looking at the case $C = 2$, the smallest maximum median-depth values are produced by the simple BT-C and our Balanced-BT-C approaches. In accordance with our theoretical results, in both

cases a logarithmic correlation can be identified, with a slight difference between them. Considering larger C values, this difference increases a bit, but it is not significant. For $C = 30$, in BT-C and Balanced-BT-C networks with 100000 clients, the maximum median-depths are 20 and 30, respectively, and they are even less in smaller networks. Smoothed-BT-C also produces a logarithmic relationship with a bit larger base. In the previous example, the maximum diameter resulted by this approach is less than 65. Taking into account the resulted maximum out-degrees and the ease of its implementability, this method could also perform well in practice, providing a good trade-off between reduced load and a bit longer diameter. In addition, we have seen that MultipleChoice-BT-C results in constant maximum out-degree which is less than what we can get in the case of our Balanced-BT-C method. However, the correlation between the diameter and the network size is much worse than logarithmic and seems to be proportional to $(\ln n)^2$. According to Theorem 3.1, both the diameter and the maximum out-degree of a network have significant influence on how much time it takes to fully distribute the data blocks of a given file in the network. In this respect, our improved Balanced-BT-C method outperforms all the other examined approaches, keeping the load of peers at a constant level and producing short paths in the network whose lengths can be bounded by $O(\ln n)$ with high probability.

It can be said that though the multiple choice can reduce the maximum out-degree in the overlay topology, but in itself it results a too regular network with unmanageably huge diameter. It shows the practical meaning of our balancing technique that can remedy this phenomenon.

The histograms of the maximum median-depth values observed in our simulations are presented in Figure 4. The first conspicuous difference we can recognize is that our balanced algorithm shows significantly narrower distribution compared to the other cases. All the values fall between 34 and 47, while in the other networks they show much higher variance, resulting twice or more wider ranges. We can also see that

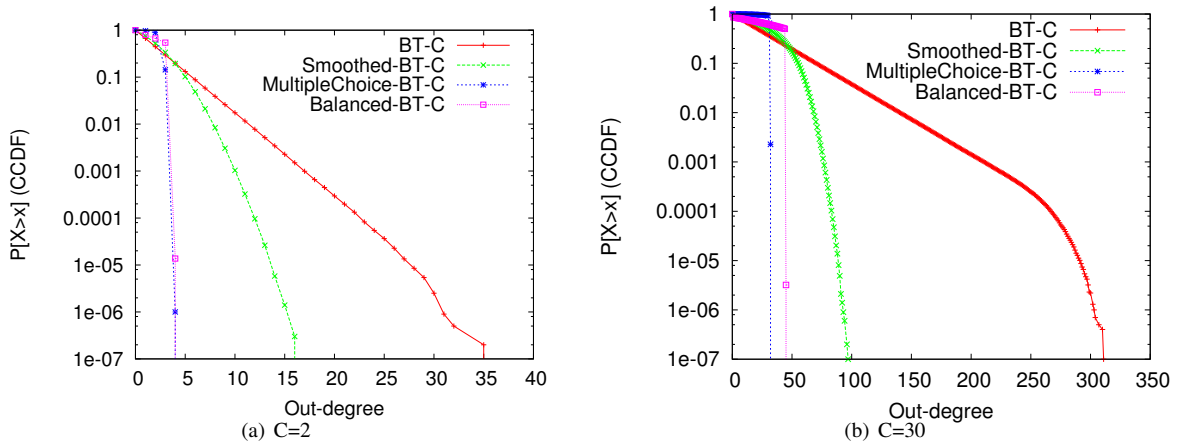


Fig. 2. The CCDF of the out-degrees in a network containing $n = 100000$ peers. Each method has been evaluated 100 times.

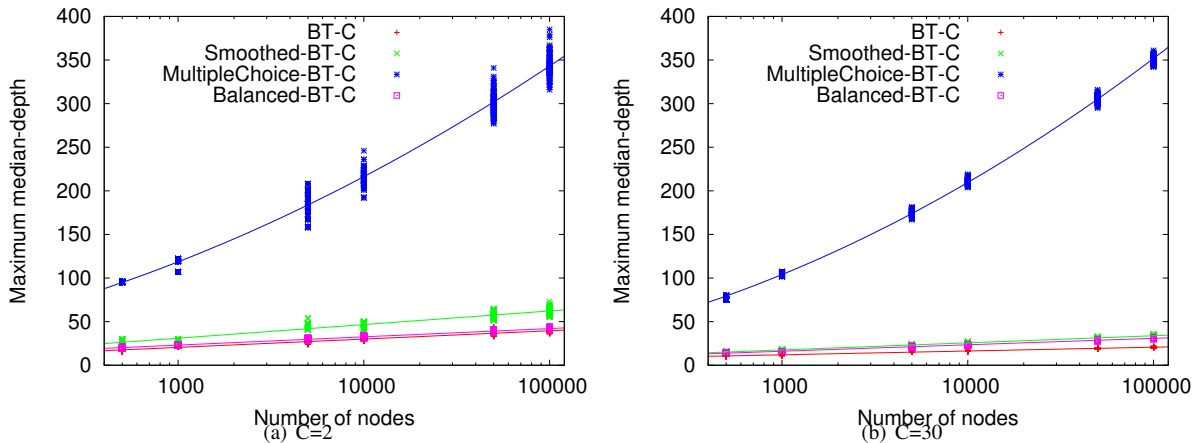


Fig. 3. Maximum median-depths for the different overlay networks. Each dot represents a maximum median-depth value after the given simulation has been done. For each parameterization and algorithm, the simulations have been repeated 100 times. The fitted curves/lines indicate the relationship between the maximum median-depth and the number of clients in the network.

the difference between *BT-C* and *Balanced-BT-C* is even less than what can be derived from Figure 3. The most likely maximum median-depth values are 39 and 42, respectively. However, in case of *Smoothed-BT-C* and *MultipleChoice-BT-C* the observed values cover significantly larger ranges.

Besides the maximum values, we have also examined the distribution of median-depths in the different overlays. Figure 5 depicts the CCDF of the median-depth values where each network consists of 100000 peers. As it is expected, *MultipleChoice-BT-C* provides significantly higher median-depths with a much slower decay than the others. Meanwhile, the other three methods show very similar distributions. In Figure 5(a), where $C = 2$, the slope in the CCDF plot of *Balanced-BT-C* is much sharper than in the case of the other two methods, which indicates that the most likely values are coming from a narrow range. It means that not only the maximum diameters have smaller variance, but the median distances from the source to ordinary peers as well. The CCDF of *Balanced-BT-C* takes place between *BT-C* and *Smoothed-BT-C* in all three figures ($C = 2$ and 30). One

can also recognize that for larger C , the median-depths in our balanced overlay follow almost the same distribution as what can be seen in the case of *Smoothed-BT-C*. As C is increasing, the difference between the two distribution is basically disappearing which suggests that, in terms of median-depths, *Balanced-BT-C* could be at most as worse as the *Smoothed-BT-C* method.

A. BitTorrent Simulation

In Section IV we have shown that the full distribution of b data blocks in the network takes $O(b + \ln(n))$ time steps with high probability in the optimal case. Among the examined methods, only *Balanced-BT-C* can construct such network that can reach this optimum. The proof of this theorem considers a randomized upload policy and assumes that the overlay topology remains invariable during the content sharing and the neighborhood of a given peer can be influenced by the tracker only. Unfortunately, existing tracker-based networks like BitTorrent do not meet these assumptions. For example, BT enables ordinary peers to extend their neighborhood with other

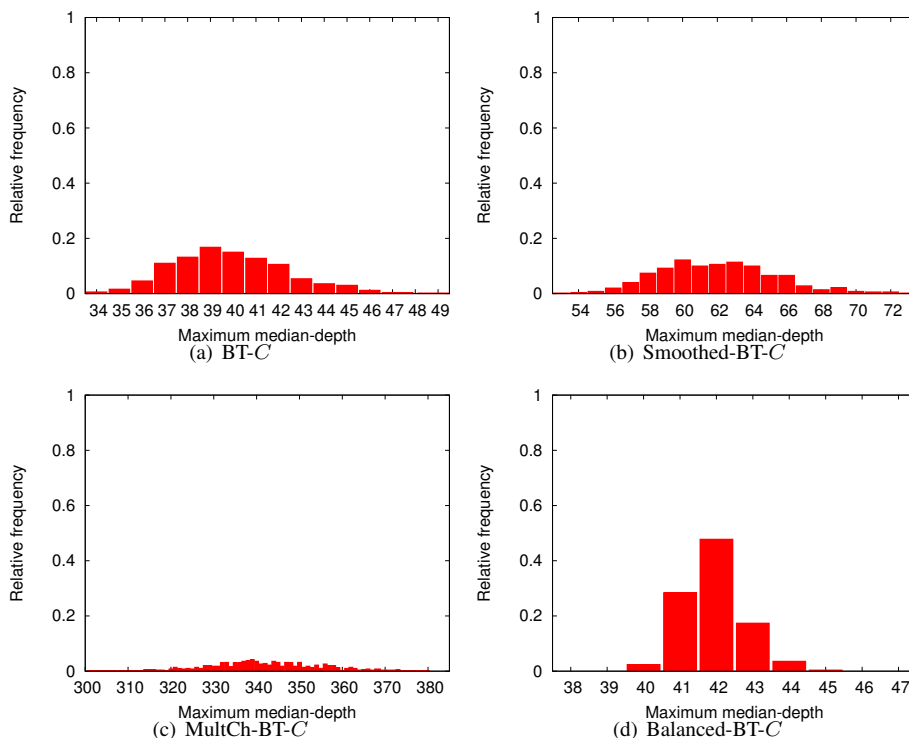


Fig. 4. The empirical distribution of the maximum median-depths observed in our simulations. Each algorithm has been performed 500 times with the parameters $n = 100000$ and $C = 2$.




Fig. 5. The CCDF of the median-depths for different overlay constructions with fixed network sizes $n = 100000$. Each method has been evaluated 100 times.

peers that want to download pieces, and to remove worthless neighbors, as well. This mechanism can dynamically improve the structure of the constructed overlay topology. However, a good initial neighborhood provided by the tracker can improve the speed of data distribution even at the beginning of the download process, especially in flash crowd settings when a sudden increase in the number of peers appears.

To analyze the different neighbor selection strategies in a more realistic environment, we have extended the BitTorrent implementation of the PeerSim simulation framework with two other trackers using the Smoothed-BT-C and the Balanced-BT-

C strategies. Note that other parts of the BT protocol in the simulator (e.g. policies like tit-for-tat, optimistic unchoke, etc.) have not been changed. For each simulation scenario discussed in this section, all the three tracker implementations have been evaluated 20 times under the same conditions corresponding to a flash crowd setting, where simulations start with only one seeder possessing the whole file, and then in every 20 milliseconds a new peer arrives and joins the network until the number of peers reaches 1000. The size of the file to be shared in the simulation is 10MB. With these settings, downloading all the pieces would not take too much time, fostering that

the neighborhoods of the peers are less affected by other BT policies like peer exchange.

In the first simulation scenario, we examine the effect of different tracker strategies on the performance of piece distribution in a homogeneous environment where each peer connects to the network with the same, constant bandwidth and delay. The block download time in PeerSim has been set to 1 ms. We also note that after downloading all the pieces, peers remain in the network, serving other clients as seeders. Figure 6 illustrates the performance of the three neighbor selection techniques for both small and large initial neighborhood sizes ($C = 4$ and 30), depicting the number of peers having downloaded the whole file as function of time. Even in this relatively small network, the improved performance of our balanced algorithm is apparent.

For $C = 4$, one can observe that the Balanced-BT- C and the Smoothed-BT- C trackers overcome the original random neighbor selection strategy of BT, leading to a better overlay routing in which the pieces can be spread faster. Though our Balanced-BT- C provides the fastest growth in the number of complete downloads, the difference to the Smoothed-BT- C variant is not significant. More than 540 out of 1000 peers finish downloading in 300 seconds, if the initial neighborhood is provided by our balanced strategy. For BT- C and Smoothed-BT- C , only 220 and 380 peers become ready during this time, respectively. The difference is more significant at 400 seconds, where the original BT strategy results 500 out of 1000 complete downloads, while for the two advanced techniques approximately it is 800 peers, indicating that 80% of the downloaders have every piece of the file. The full distribution of the file in the network is completed in 550 seconds in average for the two advanced methods which is approximately 2.5 minutes shorter than what we can see for BT- C .

In Figure 6(b), we consider a larger neighborhood size ($C = 30$) advised by the BT specification as a proper initial peer list size. One can observe that all three approaches show very similar performance, and Balanced-BT- C provides only a slightly better result than the other two solutions. This phenomena is in accordance with the theoretical results.

Besides the homogeneous case, we also examine how different strategies perform in a heterogeneous environment where links with different characteristics are used to connect peers to each other. The block downloading delay on these links have been chosen from the range between 10 and 1000 ms uniformly at random. The simulation results are depicted in Figure 7 for both $C = 4$ and 30.

Considering small initial neighborhoods ($C = 4$), one can observe that our Balanced-BT- C starts faster than the other two approaches. At 900 seconds, 620 out of 1000 peers have downloaded the complete file in the network constructed by Balanced-BT- C strategy. For BT- C and Smoothed-BT- C , this number is 340 and 480, respectively, which is a significant difference to our balanced algorithm. However, the full distribution of the pieces in the whole network takes almost the same time for both Balanced-BT- C and BitTorrent- C , while Smoothed-BT- C runs 200 seconds longer in average.

We can also recognize that for BT- C , the slope of the curve is basically the same as the one for our balanced strategy, indicating that after a slower starting period its performance becomes as good as it is for the Balanced-BT- C method.

For $C = 30$, the original BT- C and our balanced strategy show very similar behavior, however the difference between Smoothed-BT- C and the other two approaches is significant. The curve of Smoothed-BT- C starts together with the others, but at a point it slows down, and finally provides an approximately 25% higher overall file distribution time in average.

The preliminary simulation results show that our balanced neighbor selection strategy can work in practice as well, and lead to a more efficient overlay routing speeding up the file-sharing process in heavy loaded situations. Furthermore, it creates a well balanced overlay network even at the beginning of file downloads. The tracker On the other hand, BitTorrent applies numerous methods like peer exchange to optimize the neighborhoods initially provided by the tracker. These mechanisms dynamically modifying the overlay topology, changing its properties, such as degree distribution, diameter, etc...

VI. CONCLUSION

In this paper, we have introduced a novel neighbor selection strategy which uses the idea of multiple choice to improve the performance of spreading blocks in a BitTorrent like peer-to-peer network. Our balanced multiple choice algorithm takes into account not only the current load of a given peer, but the expected value that it will be selected as uploading neighbor in the future. The constructed overlay topology has been analyzed from both theoretical and experimental aspects and it has been proved that this topology has constant degree and logarithmic diameter, with high probability. We have also shown that considering a randomized upload policy, routing of b data blocks in the proposed network requires at most $O(b + \ln n)$ time steps, with high probability, which is optimal up to a constant factor. Besides the theoretical analysis, thorough simulations have been performed to examine the graph properties of the constructed networks and validate the theoretical results as well. In order to adapt our algorithm in BitTorrent networks only a slight modification of the tracker is necessary without any change in the clients. If a new peer joins the network, the tracker has to check the degree of $O(\ln n)$ peers in order to provide a neighbor list for the new peer. We have implemented a modified tracker for the BitTorrent simulation of the PeerSim simulation framework to show that our neighbor selection strategy could work even in practice and achieve higher performance in file-sharing than the original tracker of the BitTorrent.

ACKNOWLEDGMENT

Supported by the EU EIT project SmartUC, the grant EITKIC-12-1-2012-0001 of the National Development Agency, and the EU FP7 OpenLab project (Grant No. 287581). Performed in cooperation with the EIT ICT Labs Budapest Associate Partner Group (www.ictlabs.elte.hu). Sándor Laki also thanks the partial support of MAKOG Foundation.




Fig. 6. Performance evaluation of BitTorrent simulations in a homogeneous environment with constant link delays. Each curve indicates the number of nodes that have downloaded all the pieces of the file to be shared in function of time for different tracker policies. In the left figure, the initial neighborhood size (C) has been set to 4, while the right figure illustrates how the different strategies perform with $C = 30$.




Fig. 7. The performance of different tracker policies in a heterogeneous environment where link delays have been chosen from 10 ms to 1000 ms uniformly at random. Simulations have been carried out for both small and large neighborhood sizes.

REFERENCES

- [1] Bittorrent [online], available: <http://bittorrent.com>.
- [2] Tribler [online], available: <http://www.tribler.org>.
- [3] D. Arthur and R. Panigrahy. Analyzing bittorrent and related peer-to-peer networks. In *Proc. 17th ACM-SIAM Symp. on Discrete Algorithms (SODA)*, pages 961–969, 2006.
- [4] Y. Azar, A. Z. Broder, A. R. Karlin, and E. Upfal. Balanced allocations. *SIAM J. Comput.*, 29(1):180–200, 1999.
- [5] A. R. Bharambe, C. Herley, and V. N. Padmanabhan. Understanding and Deconstructing BitTorrent Performance. Technical Report MSR-TR-2005-03, Microsoft Research, 2005.
- [6] A. R. Bharambe, C. Herley, and V. N. Padmanabhan. Analyzing and improving a bittorrent networks performance mechanisms. In *Proc. 25th IEEE Int. Conf. on Computer Communications (INFOCOM)*, pages 1–12, 2006.
- [7] 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 *ICDCS'06*, pages –1–1, 2006.
- [8] B. Bollobás and W. Fernandez de la Vega. The diameter of random regular graphs. *Combinatorica*, 2(2):125–134, 1982.
- [9] F. V. Hecht, T. Bocek, and B. Stiller. B-tracker: Improving load balancing and efficiency in distributed p2p trackers. In *Proc. 11th IEEE Int. Conf. on Peer-to-Peer Computing (P2P'11)*, pages 310–313, 2011.
- [10] M. Izal, G. Urvoy-Keller, E. W. Biersack, P. Felber, A. Al Hamra, and L. Garcs-Erice. Dissecting bittorrent: Five months in a torrent's lifetime. In *PAM '04*, pages 1–11, 2004.
- [11] A. Montresor and M. Jelasity. PeerSim: A scalable P2P simulator. In *Proc. 9th IEEE Int. Conf. on Peer-to-Peer Computing (P2P'09)*, pages 99–100, 2009.
- [12] R. Motwani and P. Raghavan. Randomized algorithms. *ACM Comput. Surv.*, 28(1):33–37, 1996.
- [13] M. Naor and U. Wieder. Novel architectures for p2p applications: The continuous-discrete approach. *ACM Trans. on Algorithms*, 3(3), 2007.
- [14] M. Picone, M. Amoretti, and F. Zanichelli. An evaluation criterion for adaptive neighbor selection in heterogeneous peer-to-peer networks. In *Proc. 12th IFIP/IEEE Int. Conference on Management of Multimedia and Mobile Networks and Services: Wired-Wireless Multimedia Networks and Services Management*, pages 144–156, 2009.
- [15] J. A. Pouwelse, P. Garbacki, D. H. J. Epema, and H. J. Sips. The bittorrent p2p file-sharing system: Measurements and analysis. In *Proc. Int. Workshop on Peer-to-Peer Systems (IPTPS)*, pages 205–216, 2005.
- [16] H. Wang, J. Liu, B. Chen, K. Xu, and Z. Ma. On tracker selection for peer-to-peer traffic locality. In *Proc. 10th IEEE Int. Conf. on Peer-to-Peer Computing (P2P'10)*, pages 1–10, 2010.
- [17] H. Zhang, Z. Shao, M. Chen, and K. Ramchandran. Optimal neighbor selection in bittorrent-like peer-to-peer networks. In *Proc. ACM SIGMETRICS Joint Int. Conf. on Measurement and Modeling of Computer Systems*, pages 141–142, 2011.