

Adapting BitTorrent to wireless ad hoc networks

Mohamed Karim Sbai, Chadi Barakat, Jaeyoung Choi,
Anwar Al Hamra, and Thierry Turretti

Project-Team Planète, INRIA Sophia Antipolis, France
{mksbai, cbarakat, jchoi, aalhamra, turretti}@sophia.inria.fr

Abstract. BitTorrent is one of the Internet's most efficient content distribution protocols. It is known to perform very well over the wired Internet where end-to-end performance is almost guaranteed. However, in wireless ad hoc networks, many constraints appear as the scarcity of resources and their shared nature, which make running BitTorrent with its default configuration not lead to best performances. To these constraints it adds the fact that peers are both routers and end-users and that TCP-performance drops seriously with the number of hops. We show in this work that the neighbor selection mechanism in BitTorrent plays an important role in determining the performance of the protocol when deployed over a wireless ad hoc network. It is no longer efficient to choose and treat with peers independently of their location. A first solution is to limit the scope of the neighborhood. In this case, TCP connections are fast but there is no more diversity of pieces in the network: pieces propagate in a unique direction from the seed to distant peers. This prohibits peers from reciprocating data and leads to low sharing ratios and suboptimal utilization of network resources. To recover from these impairments, we propose an enhancement to BitTorrent which aims to minimize the time to download the content and at the same time to enforce cooperation among peers. Our solution considers a restricted neighborhood to reduce routing overhead and to improve throughput, while establishing few connections to remote peers to improve diversity of pieces. With the help of extensive NS-2 simulations, we show that these enhancements to BitTorrent significantly improve the file completion time while fully profiting from the incentives implemented in BitTorrent to enforce fair sharing.

Key words: BitTorrent, wireless ad hoc networks, neighbor selection, piece selection, completion time, fair sharing

1 Introduction

Wireless ad hoc networks and P2P file sharing applications are two emerging technologies based on the same paradigm: the P2P paradigm. This paradigm aims to establish large scale distributed services without the need for any infrastructure. Within this paradigm, users have symmetric roles. The global service is ensured thanks to their collaboration. In the case of a wireless ad hoc network, the network is a set of wireless nodes with no central administration or base station. Nodes in such a network operate both as routers and hosts. Multi-hop

routing approaches are used to ensure connection between distant nodes. For P2P file sharing applications, peers collaborate in downloading data and multimedia content. Each peer shares some of its upload capacity by serving other peers. The global capacity of the system grows then exponentially with the number of peers. Gnutella [6], Freenet [7] and BitTorrent [1] are examples of P2P content sharing applications in the Internet.

Both P2P file sharing applications and wireless ad hoc networks are mature fields of research. They have been studied heavily but separately in the literature. Only few works try to study how they perform together (e.g., [12] [13] [14]). These works focus on the content lookup problem in wireless ad hoc networks without studying the efficiency of the content sharing itself. Studying the performance of file sharing applications over wireless ad hoc networks is challenging because of the diverse constraints imposed by the use of wireless channels. Indeed, as nodes are both routers and end-users, the routing overhead must be taken into consideration. Furthermore, the performance of transport protocols such as TCP drops seriously when multi-hop paths are used. That is why current topology-unaware P2P file sharing applications are not expected to perform well when deployed over wireless ad hoc networks. Designing efficient file sharing solutions for such networks is an important area of research. Indeed, a P2P solution for file sharing has diverse advantages over other data dissemination techniques like multicast in general and this applies to wireless ad hoc networks in particular. For instance, in case of multicast, the construction and update of the virtual topology (tree or mesh) is costly in terms of bandwidth consumption namely in dynamic scenarios. Moreover, the data replication in multicast follows the virtual topology and so nodes like leaves of a tree only receive data and do not spend resources to provide it to other nodes. Thus, no fair cooperation is ensured when using multicast unless constructing a different virtual topology (or tree) per piece of data, which is technically unfeasible.

In this work, we investigate how well a P2P file sharing solution developed for the wired Internet performs over a wireless ad hoc network. Our aim is to come up with a solution that minimizes the content download time while at the same time improving collaboration by enforcing fair sharing among peers. As efficient and fair content sharing is targeted, we choose to adapt BitTorrent [1] as a file sharing protocol given its large usage and its known close to optimal performances in the wired Internet [15]. When data is distributed using BitTorrent, interested peers supply pieces of the data to other peers, reducing the burden on any individual peer, providing redundancy in the network, and reducing dependency on the original seed. In addition, BitTorrent implements incentives that encourage peers to collaborate in downloading the content, which is not the case of multicast-tree based solutions.

In a first effort to understand this problem, we consider the particular case when every ad hoc node is interested in downloading the content. In this case, the underlying topology has a big impact on the performance of BitTorrent. Indeed, any piece sent over a suboptimal route will cause resource consumption in

all intermediate nodes. When all nodes are peers, this will affect all peers located on these nodes by stealing bandwidth from them without being able to profit from this transmission since it happens at the routing layer. However, if intermediate nodes are not peers interested in the same content, this suboptimal piece transmission will have less impact on the torrent itself since it does not directly steal bandwidth from peers (it will steal bandwidth from other applications however). Add to this the fact that when all nodes are peers, the traffic generated by the torrent is maximal and an optimization is further required. We aim at well understanding this case and proposing an efficient solution for it before moving into less loaded scenarios in future work namely the scenario where only a part of the nodes are peers. The performance evaluation is done through extensive NS-2 simulations using regular modules for the ad hoc routing and wireless medium and our implementation of BitTorrent in NS-2¹. Our main contributions can be summarized as follows. Ordinary BitTorrent establishes TCP connections with neighbors independently of their location. This choice of neighbors can lead to slow TCP connections due to long multi-hop paths and routing overhead. Sharing can also be bad when using large pieces since complete pieces cannot be sent too far to be reused later by other peers. A first solution is to limit the scope of the neighborhood. In this case, we noticed shorter download times but poor sharing since there is no diversity of pieces in the network. To recover from these impairments, we propose an enhanced variant of BitTorrent, tuned to ad hoc networks, which considers a restricted neighborhood to diminish routing overhead and to improve throughput, while establishing few connections to remote peers to improve diversity of pieces. To implement this, we modify the choking algorithm and add a new piece selection strategy. The simulations show that these enhancements to BitTorrent considerably improve the file completion time while fully benefiting from the incentives implemented in BitTorrent to enforce fair sharing.

Section 2 of this paper presents an overview of the state of the art in deploying P2P solutions over wireless ad hoc networks. Section 3 describes briefly the BitTorrent protocol. The framework of the study is discussed in Section 4. Section 5 shows the importance of the piece size in determining the performance of BitTorrent. Section 6 studies the impact of the scope of the neighborhood. Section 7 presents our enhanced variant of BitTorrent. Section 8 summarizes the work and gives some ideas on our future work.

2 State of the art

In this section, we present an overview of the state of the art of P2P file sharing applications and their different implementations in wireless ad hoc networks.

P2P applications in the Internet: There are several design approaches for the construction of P2P overlays over the Internet. One can distinguish between structured and non-structured overlays. This classification is done from

¹ NS-2 code and scripts at: http://planete.inria.fr/personnel/Mohamed_Karim.Sbai/BitTorrent/AdaptedBitTorrent.htm

the standpoint of resources lookup. In non-structured overlays like Gnutella [6], there is no control on the structure of the overlay. Peers discover each other by flooding the network and by learning from previous sessions. The P2P application in this case is not conscious of the topological location of the other peers. In case of structured overlays, an overlay routing algorithm is introduced to locate the content in the network. Several structured overlay networks have been proposed like CAN [8], Chord [9], Pastry [11] and Tapestry [10]. All of them use Distributed Hash Tables (DHT) in their routing of lookup requests. Such tables allow the lookup to scale logarithmically with the number of nodes in the overlay. Again most of these structured overlays are topology independent. On the other hand, there is BitTorrent [1] that does not concentrate on the information lookup since it uses a centralized tracker to discover neighbors. However, it concentrates on optimal utilization of the network capacity when sharing the file between the different interested peers. Since we are mainly concerned in this work by the data transfer plane, we adopt BitTorrent and we extend it to wireless ad hoc networks. More details on BitTorrent are presented in Section 3.

P2P applications in Mobile Ad hoc NETWORKS: Both structured and non-structured overlays have been implemented in MANET. Since nodes are both end-users and routers, some cross-layer design approaches have been introduced. These approaches suppose that P2P applications operate both at the network layer and at the application layer. One can divide the design space into four subspaces:

- Non-structured and layered design: Oliviera et al. study in [12] the performance of Gnutella deployed over three ad hoc routing protocols DSR, AODV and DSDV. Their results show that the ratio of delivered packets is lower than those of unicast applications deployed over MANET. This is due to the fact that Gnutella chooses neighbors independently of their locations. The overlay construction is topology independent.
- Non-structured and cross-layer design: The work done by Klemm et al. in [13] proposes to integrate the peer lookup mechanism of a P2P application like Gnutella in the network layer and compares this design to the layered design proposed by Oliviera et Al. They propose ORION that establishes connections on demand through the routing mechanism. The cross-layer lookup implemented by ORION is shown to provide higher successful transfers ratio than in the layered scenario.
- Structured and layered design: A proximity-conscious DHT (Pastry) has been deployed over the DSR routing protocol in [14]. As it is a layered design, there is no interaction between the DHT and the routing protocol. This leads to an overhead in maintaining routes for both the application layer and the network routing layer.
- Structured and cross-layer design: This design is named Ekta by Das et al. in [14]. The functionalities of the Pastry DHT are integrated within the routing protocol. The main idea is the mapping of the peer identifiers in the same namespace than the IP addresses. Their results show that Ekta is

better than the layered design in terms of number of successfully delivered packets.

Former studies on BitTorrent over wireless ad hoc networks: Several works tried to adapt BitTorrent to wireless ad hoc networks (e.g. [16] and [17]). They only focus on the tuning of the peer discovery phase without addressing the efficiency of the content sharing itself. Michiardi et al. study in [5] the performance of a cooperative mechanism to distribute content from one source to a potentially large number of destinations. They propose to deploy BitTorrent with a minor change allowing neighbor discovery and traffic locality. This is done by selecting only near neighbors as effective neighbors. The result is a decrease in the total download time and energy consumption. Their work is relevant to ours; however we go beyond by focusing not only on the download time but also on the sharing among peers which we will show to further improve the download time as well.

3 BitTorrent: a content distribution protocol

BitTorrent (see e.g., [1], [15]) is a scalable P2P content distribution protocol. Each client shares some of its upload bandwidth with other peers interested in the same content in order to increase the global system capacity. Peers cooperating to download the same content form a *torrent*. A peer discovers other peers by contacting a central rendezvous node called *tracker*. The latter stores IP addresses of all peers in the torrent and maintains statistics on uploads and downloads per peer. To facilitate the replication of the content in the network and to ensure multi-sourcing, a file is subdivided into a set of pieces. Each piece is also subdivided into blocks. A peer which has all pieces of the file is called *seed*. When the peer is still downloading pieces, it is called *leecher*. Each peer maintains a peer list. Neighbors are those of this list with whom the peer can open a TCP-connection to exchange data and information. Only four simultaneous outgoing *active* TCP connections are allowed by the protocol. The corresponding neighbors are called *effective neighbors*. They are selected according to the *choking algorithm* of BitTorrent. This algorithm is executed periodically. Once the *choking period* expires, a peer chooses to unchoke the 3 peers uploading to him at the highest rate. It is a best slot unchoking. This strategy, called *tit-for-tat*, ensures reciprocity and enforces collaboration among peers. Now to discover new upload capacities, a peer chooses randomly a fourth peer to unchoke. This unchoking slot is called optimistic slot. All other neighbors are left choked. When unchoked, a peer selects a piece to download using a specific piece selection strategy. This strategy is called *local rarest first*. Indeed, each peer maintains an update-to-date list of pieces owned by all its neighbors. When selecting a piece, a peer chooses the piece with the least redundancy in its neighborhood. In case of equality, one of the rarest pieces is chosen randomly. Rarest first is supposed to increase the entropy of pieces in the network which enforces collaboration and hence improves global performance.

Here are the performance metrics relevant to BitTorrent that we will use in our study and that are calculated at the end of the experimentation:

- U_{ij} : Total bytes uploaded by peer i to peer j .
- D_{ij} : Total bytes downloaded by peer i from peer j . ($U_{ij} = D_{ji}$)
- R_{ij} : Ratio of sharing between peer i and peer j .

$$R_{ij} = \frac{\min(U_{ij}, D_{ij})}{\max(U_{ij}, D_{ij})} \quad (1)$$

- N_i : Number of neighbors j of peer i such that $U_{ij} \neq 0$ or $D_{ij} \neq 0$.
- R_i : Sharing ratio for node i .

$$R_i = \frac{1}{N_i} \cdot \sum_{j | U_{ij} \neq 0 \text{ or } D_{ij} \neq 0} R_{ij} \quad (2)$$

- F_i : The finish time of peer i . It is the time by which it receives all pieces of the file.

As we are studying BitTorrent over wireless ad hoc networks where topology matters, we consider some additional performance metrics related to topological positions of peers. The file is supposed to exist at one seed S at the beginning of the session. Our metrics quantify the quality of service perceived by peers as a function of their relative positions with respect to the seed.

- F_h : Average finish time of peers (or nodes) located at h hops from seed S .

$$F_h = \frac{1}{n_h} \cdot \sum_{i | H(i)=h} F_i \quad (3)$$

where n_h is the number of peers located at h hops from seed S and $H(i)$ a function that gives the number of hops between any node i and the seed S .

- R_h : Average sharing ratio of peers (or nodes) located at h hops from seed S .

$$R_h = \frac{1}{n_h} \cdot \sum_{i | H(i)=h} R_i \quad (4)$$

4 Framework of the study

We proceed with an experimental approach using the NS-2 simulator. In this section, we describe preliminary changes we made to BitTorrent to allow peer discovery and the exchange of signaling over wireless ad hoc networks. Then, we discuss the stack of protocols we use in our deployment of BitTorrent in NS-2. Finally, we introduce the scenario used in our evaluation.

4.1 Trackerless BitTorrent

Wireless ad hoc networks are infrastructureless. It is convenient that one does not rely on a centralized tracker when applying BitTorrent to such networks. So, we opt in our study for a trackerless approach. Since the most important role of a tracker in the Internet is to provide peers with the identifiers of other peers, we need to introduce a peer discovery mechanism. In our evaluation framework, to

discover new peers, a peer floods periodically the network with a HELLO message and waits for HELLO REPLY messages. HELLO messages are transmitted to wireless neighbors with some initial TTL (Time-To-Live) to control the scope of the flood and hence the visibility of a peer. This TTL is a parameter of our study. Receiving a HELLO message, a peer decrements the TTL and forwards it to its wireless neighbors, and so on. The message is not forwarded when its TTL reaches zero.

4.2 Stack of protocols and packets exchanged between peers

In BitTorrent, peers exchange two types of packets: Data packets and control packets. We choose in our NS-2 implementation to send data packets via TCP connections because reliability and congestion control are needed when transporting blocks of file. However, control packets as for peer discovery and piece updates contain small and urgent information that is better to transport using UDP. Here are the different control packets exchanged between peers:

- **HELLO**: see Section 4.1.
- **HELLO REPLY**: see Section 4.1.
- **UPDATE PIECE LIST**: when a peer receives a new piece, it sends an UPDATE PIECE LIST to all peers with whom it can exchange data.
- **PIECE OFFER REQUEST**: when a peer i unchokes a peer j , it sends a PIECE OFFER REQUEST packet to j . This packet contains the list of pieces that i has already downloaded.
- **PIECE OFFER REPLY**: receiving a PIECE OFFER REQUEST, a peer answers with a PIECE OFFER REPLY packet. After applying the piece selection strategy, it decides whether to accept or to reject the offer. A flag included in the PIECE OFFER REPLY packet indicates this decision (ACCEPT or REJECT). In the case the offer is accepted, the peer indicates the number of the requested piece. During the choking period, many PIECE OFFER REPLY packets can be sent to the offering peer in order to allow the transmission of several pieces.

4.3 The main scenario

We consider a network of N nodes ($N=40$ when not specified) distributed in a plane following a grid topology (10 nodes per row). The distance between two physical neighbors is set to 40 m for a range of wireless transmissions equal to 50m. This ensures connectivity while minimizing interference. At the beginning of each simulation, node 0 located at the top left is the seed and the other nodes are leechers. The file size is set equal to 10 Mbytes, which is large enough to ensure the convergence of the protocol to equilibrium. All peers start downloading the file at the same time $t=1500s$ by first looking for each other then sharing the pieces of the file according to the BitTorrent algorithms. This time interval skipped at the beginning gives the network enough time to stabilize and calculate its routing tables. The bitTorrent choking algorithm period is taken in our simulations equal to 40s. A piece is subdivided into blocks of size 1KB. Concerning the underlying layers, the nodes connect to each other using the 802.11 MAC Layer with the RTS/CTS-Data/ACK mechanism enabled. The data rate is set to 1 Mb/s. For ad hoc routing, we use the DSDV proactive protocol [2].




Fig. 1. Average finish time as a function of number of hops to seed




Fig. 2. Average sharing ratio as a function of number of hops to seed

5 Impact of piece size

We start by evaluating regular BitTorrent where the overlay is constructed without considering the underlying wireless topology. We give a particular attention to the piece size and to its impact on both the finish time of peers and their sharing ratios. The reason to consider the piece size is that it decides how far pieces can be sent over the network. The TTL of HELLO messages is set to its maximum value so that all peers are neighbors of each other. Two sizes of pieces are used while keeping constant the size of the file. We consider respectively the values 100 blocks and 1000 blocks for small and big size of pieces. Figure 1 plots the average finish time F_h as a function of the number of hops h to the seed for both small and large size of pieces. Each point in this figure is an average over multiple simulations and over all nodes located at the same number of hops to the seed. As expected, the finish time increases as far as we move away from the source. One can notice in the figure that for small pieces, remote peers have better finish time than for large pieces. This is because the range of transmission of small pieces is longer. A remote peer can then receive more pieces in the choking period and share them with others, which improves the reusability of pieces and network resources. This is confirmed in Figure 2 where we plot the average sharing ratio R_h as a function of the number of hops to the seed. It is clear that the sharing ratio in case of small pieces is more important because distant nodes (or peers) can now get quickly complete pieces and replicate them in their neighborhood. Unfortunately, this is not the case with large pieces. Large pieces cannot be sent far in the choking period so they propagate in the network as a wave resulting in an under-utilization of network capacity. One can see the case of large pieces as being the absence of sharing between distant nodes and the fact that nodes wait for pieces to arrive to their upstream nodes before obtaining them. The use of small pieces however make the pieces spread over the network, which reduces the finish time and makes the sharing incentives implemented by BitTorrent work better in wireless ad-hoc networks. Our solution supports this modification.

6 Impact of the scope of the neighborhood

Another important factor in BitTorrent over wireless ad hoc networks is the scope of the neighborhood. In this section, we study the impact of reducing




Fig. 3. Average finish time as a function of number of hops to seed for different flooding scope




Fig. 4. Average sharing ratio as a function of number of hops to seed for different flooding scope

this scope on both the finish time and the sharing ratio. We run several simulations on the topology described in 4.3 changing each time the flooding scope (TTL) of HELLO messages destined to peer discovery. Figure 3 compares the finish time for TTL=max, 5, 2 and 1. Interestingly, the finish time improves when the neighborhood scope is decreased. This is mainly due to better TCP performance over short paths and to smaller routing overhead. Control packets, namely PIECE UPDATE and HELLO packets, are sent only in the restricted neighborhood. The case TTL=2 is slightly better than the case TTL=1 because of the interference between physical neighbors. Figure 4 plots the average sharing ratio R_h as a function of the number of hops to the seed for the different values of TTL. Unfortunately, we can see that the improvement in finish time when reducing the neighborhood comes at the expense of a lower sharing ratio. The diversity of pieces in the network decreases and the file propagates more or less as a wave in a unique direction from the seed to the farthest nodes. Hence, distant peers can not participate in the replication of pieces, they only wait for pieces to arrive to their physical neighbors to obtain them. Clearly, this is bad for cooperation among peers. An optimal solution should improve the finish time while preserving large values for the sharing ratio.

7 BitTorrent adapted to wireless ad hoc networks

The main objective of our variant of BitTorrent is to profit from the advantages of the limited neighborhood, namely the good performance of TCP on short paths, the reduced routing overhead, and the reduced load of flooding control packets. At the same time, we aim at improving the sharing ratio and the reusability of network resources by creating diversity of pieces in the network. Our main idea is to create few TCP connections to distant peers in addition to those with close peers. Pieces can then spread over the network and propagate in different directions, which improves the sharing and the download completion time. With this modification, several zones of the network can be active simultaneously, which is not the case of the wave generated by regular BitTorrent with limited neighborhood. To implement this idea, we tune BitTorrent to support the distinction between remote and close peers. The new choking algorithm is aware of the location of peers by using routing information. It distributes optimistic

unchokes between remote and close peers and adds a specific neighbor selection mechanism to select a distant peer. It also applies a new piece selection strategy when the peer offering the piece is distant. Unlike BitTorrent with limited neighborhood, this modification requires a global knowledge about the identifiers of peers in the network. We propose that each peer maintains two neighbor tables: NEARBY NEIGHBORS TABLE (NNT) and FAR NEIGHBORS TABLE (FNT). When discovering new peers, neighbors whose number of hops is less than or equal to 2 are added to NNT. Other peers belong to FNT. The PIECE UPDATE packets are sent only to neighbors in NNT. Peers do not need to know about all pieces in the network as their piece selection strategy operates only on their NNTs. Indeed, in wireless networks, the replication of pieces is more efficient when it is based on statistics in the close neighborhood since this guarantees a faster local replication compared to when statistics are based on a large neighborhood. As in BitTorrent, when the choking algorithm is executed, three best uploaders are selected as effective neighbors. These three neighbors are chosen from both nearby and far neighbor tables. The peer then serves these three neighbors during the next choking period. But in addition to these effective neighbors, the peer selects a fourth random neighbor from one of the two tables (optimistic slot). The table from which it selects the neighbor is decided by a round robin policy that guarantees an optimal balance between the random unchokes locally and the transmission of pieces to distant neighbors in order to improve diversity. For a succession of optimistic unchokes, the peer selects a peer one time from FNT, q times from NNT and so on. In our protocol, the quantum q represents the ratio of the number of time slots spent on serving nearby neighbors and those for serving far neighbors. It is also the number of slots that a peer should wait before unchoking a distant neighbor again. Our simulations indicate that the choice of this quantum is fundamental in deciding the performance of our solution. Furthermore, the strategies of selecting pieces proposed by distant neighbors and selecting effective neighbors from FNT should differ from the ordinary strategies applied by BitTorrent because the objective of our version of BitTorrent in unchoking far peers is mainly to improve diversity. The next paragraphs explain the different selection strategies we implement in our solution. The following ones study the performance of the enhanced BitTorrent and discuss the choice of the quantum q .

7.1 Selecting a far neighbor at random

When a regular BitTorrent client decides to optimistically unchoke a peer, it selects it at random with a uniform probability. In wireless networks however, the gain we get from optimistic unchoking in terms of diversity increases with the number of hops. So a peer has more interest in unchoking a farther peer than another one closer to it. Thus, in our adapted version of BitTorrent, to select a far peer to unchoke from FNT, the peer starts by selecting the number of hops to that peer with a probability that increases linearly with the number of hops. Let h_m be the maximum number of hops seen by the peer. We suppose that FNT contains only peers at h_m and $h_m - 1$ hops. These are the farthest peers that if we send pieces to them, we are sure of having the largest gain in terms

of diversity and reutilization of network resources.² It follows that the number of hops is first selected using a probability function p given by this formula:

$$p(h) = \begin{cases} \frac{h}{h_m + (h_m - 1)} & \text{if } h \geq h_m - 1 \\ 0 & \text{else} \end{cases}$$

When the number of hops h is chosen, the peer then selects, in a uniform random way, a peer among those located at h hops from it as the node to optimistically unchoke.

7.2 Selecting a nearby neighbor at random

When the peer needs to select a nearby neighbor, it chooses a node from NNT in a uniform random way. A nearby neighbor is supposed to replicate the pieces it receives in its two-hop limited neighborhood³. This replication is fast since the TCP protocol has a good throughput over short paths.

7.3 Piece selection strategy when the offering neighbor is far

When receiving a piece offer from a P2P neighbor, the peer checks the number of hops to the offering neighbor. If it is greater than 2, it considers that it is an offer from a far node. In this case, a specific piece selection strategy is applied in order to select the best piece to download from this node. This strategy will be called the *absent piece strategy*. The peer first computes the redundancy of the offered pieces in its close neighbors table and in its piece pool. At the opposite of BitTorrent, the candidate pieces will be those with zero redundancy (no need to download a piece from a distant node if it exists at less than two hops). So a piece can be accepted only if neither the peer nor one of its near neighbors has downloaded it before. In case of multiple absent pieces, one piece among them is chosen in a uniform random way. The absent piece can then be replicated quickly in the near neighborhood. If no absent piece is noticed, the peer sends a REJECT in the piece offer reply packet. In summary, our solution supposes that it is better to download a piece existing in the nearby neighborhood from a nearby neighbor. Only absent pieces are taken from far neighbors so as to reduce the routing overhead. This strategy is fundamental for getting good performances with our variant of BitTorrent.

7.4 Piece selection strategy when the offering neighbor is near

Local rarest first is used when the peer receives a piece offer from one of its nearby neighbors. Pieces with the least number of copies in the close neighborhood are selected. This is the normal behavior of the standard version of BitTorrent but only applied in the two-hop neighborhood. Here the throughput of TCP is good and the routing overhead is almost inexistent so we can allow ourselves to apply the rarest first policy that guarantees the fast replication of pieces.

² We add peers at h_{m-1} hops to FNT in order to reduce the load on the one or few peers located at h_m hops, but while having the pieces sent to those former peers seen by the latter ones.

³ We form NNT using two-hop neighborhood because according to our results in Section 6, this leads to slightly better finish time and sharing than if limiting the neighborhood to only one hop.




Fig. 5. Average finish time for our enhanced BitTorrent compared to ordinary BitTorrent with limited neighborhood




Fig. 6. Average sharing for our enhanced BitTorrent compared to BitTorrent with limited neighborhood

7.5 Simulation results

To study the performance of our solution, we run several NS-2 simulations over the previously described topology. We vary the values of the quantum q and observe the behavior of the download finish times of peers and their sharing ratios. Figure 5 compares finish time of ordinary BitTorrent with limited neighborhood ($TTL = 2$) with our version of BitTorrent using different values of the quantum q ($q=3, 2$ and 1). Each curve presents the average finish time F_h as a function of the number of hops to the seed. Recall that the role of q is to balance optimistic unchokes between close and remote peers. The larger the q , the smaller the number of unchokes to remote peers. The finish time for our solution is better and more equally distributed since far nodes can receive pieces from the beginning of the session and can replicate them in their close neighborhoods. Our solution limits the number of pieces sent to far nodes in order to reduce the routing overhead. This creates parallel areas of activity in the network. Far nodes do not need to wait for pieces to arrive to their neighborhoods to download them. Hence, pieces propagate in the network in all directions. This observation is illustrated in Figure 6 which compares sharing ratios of ordinary BitTorrent with limited neighborhood ($TTL=2$) with our variant of BitTorrent using different values of the quantum ($q=3, 2$ and 1). Each curve presents the average sharing ratio R_h as a function of number of hops to the seed. Figure 6 shows that the strategies used in our solution increase considerably the sharing ratios of all peers. This is due to the diversity created by sending original pieces to distant nodes. So, sharing incentives work well in this context and the distribution is less vulnerable to the selfishness of some nodes. Our results also show that a quantum equal to 1 gives a better finish time and a better sharing ratio in our setting. Clearly, the performance of our solution depends on the choice of the quantum q . This choice is treated in the next section.

7.6 Optimal choice of the quantum q

In this paragraph, we establish an empirical formula for q and then validate it through simulations. Let h_m be the maximum length of a path between two nodes in the network. Let α_i be the number of pieces that can be sent during a choking slot to a node located at i hops. The objective of our balanced optimistic unchoking strategy is to send a copy of each piece to the end of the network and




Fig. 7. Average finish time as a function of the chosen quantum




Fig. 8. Best quantum as a function of the number of nodes

wait for it to return to the middle of the network. Forward and backward pieces meet then in the middle of the network, which guarantees the best gain. If there were only one piece in the file, only one seed and the content is sent to the farthest node, the piece will take $\frac{h_m}{2}$ slots to return to the middle of the network. Now when the file contains several pieces, the node should wait $\frac{\alpha_{h_m}}{\alpha_1} \cdot \frac{h_m}{2}$ before unchoking the farthest node again. It is the number of slots needed for the α_{h_m} pieces to return to the middle of the network hop by hop. Now, if all peers in the network are interested in the content and if we assume nodes to be uniformly distributed in the plane, $\frac{N}{2}$ nodes at maximum can participate in sending pieces to the farthest node. So one needs to increase the waiting time by a factor of $\frac{N}{2}$. So, the formula approximating q will be:

$$q = \frac{\alpha_{h_m}}{\alpha_1} \cdot \frac{h_m}{2} \cdot \frac{N}{2} \quad (5)$$

To validate this formula, we vary the number of nodes in the network and observe how this impacts the optimal choice of q . We plot optimal q as a function of the number of nodes N . Simulations are done on grid topologies with $N=20$ to 100. Figure 7 plots the average finish time over all nodes as a function of the chosen quantum q for 40 nodes and 80 nodes (curves for all values of N are not included for clarity of presentation). Figure 8 plots both the computed and simulation results for best q . The values of α_{h_m} and α_1 are taken from simulations in both curves. Even though our expression for q is simple and approximate; we can see a good match between the two curves. In the figure, simulation values of the best q are rounded integer values of theoretical ones. Thus, the above formula describes well the behavior of the optimal q when number of nodes varies. One of the conclusions we can come to is that this quantum increases with N , which means less pieces sent by each peer to remote peers for larger networks.

8 Conclusions and perspectives

P2P data sharing applications in wireless ad hoc networks should provide good quality of service to their users in terms of finish time and sharing. There is a high potential for these applications but unfortunately, the wireless nature of the network imposes many constraints to be taken into consideration before using regular applications tuned for the wired Internet. Solutions that reduce

neighborhood scope allow better finish time than those with random graphs of communications. Nevertheless, limiting the neighborhood is shown, in this paper, to be dangerous in terms of reducing sharing ratios between peers. The solution we propose in this paper finds a good management of neighbor and piece selection that reduces finish time and encourages sharing. A peer concentrates on its nearby peers with few connections to far ones. When far neighbors are selected, a special piece selection strategy named absent piece strategy comes into effect. Simulation results show a decrease in service time and a great improve in sharing ratios. Our future work will be on adapting our solution to mobile scenarios. High dynamicity of such networks will open the way to new interesting problems.

References

1. BitTorrent protocol. <http://wiki.theory.org/BitTorrentSpecification>.
2. C.Perkins, P.Bhagwat, Highly Dynamic Destination-Sequenced Distance-Vector Routing (DSDV) for Mobile Computers, ACM SIGCOMM'94.
3. NS: The Network Simulator, <http://www.isi.edu/nsnam/ns/>
4. G. Ding, B. Bhargava, Peer-to-Peer File-Sharing over Mobile Ad hoc Networks. In IEEE PERCOM-W, Orlando, USA, 2004.
5. Michiardi P., Urvoy-Keller G., Performance analysis of cooperative content distribution for wireless ad hoc networks, WONS 2007, Obergurgl.
6. The Gnutella specification, 2000, <http://dss.clip2.com/GnutellaProtocol04.pdf>
7. The free network project, <http://freenetproject.org/>
8. S.Ratnasamy, P.Francis, M.Handley, R. Karp and S.shenker. A scalable content-addressable networks, ACM SIGCOMM, 2001.
9. I.Stoica, R.Morris, D.Karger, M.F.Kaashoek and H.Balakrishnan. Chord: a scalable peer-to-peer lookup service for internet applications, ACM SIGCOMM, 2001
10. B.Y.Zhao, J.D.Kubiatowicz, and A.D.Joseph. Tapestry: an infrastructure for fault-resilient wide-area location and routing. T.R. UCB//CSD-01-1141, U.C.Berkeley,2001
11. A.Rowstron, P.Druschel, Pastry: scalable, diributed object location and routing for large -scale peer-to-peer systems, Middleware, 2001.
12. L.B. Oliviera, I.G.Siqueira, A.A.Loureiro, Evaluation of ad hoc routing protocols under a peer-to-peer application, WCNC, 2003.
13. A.Klemm, C.Lindermann, O.Waldhorst. A special-purpose peer-to-peer file sharing system for mobile ad hoc networks, VTC, 2003.
14. S.M. Das, H.Pucha, Y.C. Hu. Ekta: an efficient peer-to-peer substrate for distributed applications in mobile ad hoc networks. TR-ECE-04-04, Purdue University, 2004.
15. A. Legout, G. Urvoy-Keller, P. Michiardi, Rarest First and Choke Algorithms Are Enough, IMC'2006, 2006, Rio de Janeiro.
16. A. Nandan, S. Das, G. Pau, M. Gerla, Cooperative downloading in vehicular ad hoc networks, WONS, Washington, USA, 2005.
17. S. Rajagopalan, C-C. Shen, A cross-Layer Decentralized BitTorrent for Mobile Ad hoc Networks, MOBIQUITOUS, San Jose, USA, 2006.