Neighborhood selection in mobile P2P networks

Emna Salhi, Mohamed Karim Sbai, Chadi Barakat

Neighborhood selection in mobile P2P networks

Emna Salhi[1 2], Mohamed Karim Sbai[1] and Chadi Barakat[1]

[1] *EPI Planète, INRIA, BP 93, F-06902 Sophia Antipolis, France*
[2] *Ecole Nationale des Sciences de l'Informatique (ENSI), Tunisie*

La politique de choix des voisins joue un rôle capital dans la détermination de l'efficacité d'une application pair-à-pair pour le partage de fichiers. L'idéal serait d'obtenir des temps de téléchargement minimaux avec le maximum d'opportunittés de partage entre les utilisateurs. Ces applications ont été largement étudiées dans l'Internet, cependant leurs performances dans les réseaux ad hoc mobiles restent à évaluer. En particulier, toute politique de choix des voisins devrait tenir compte de plusieurs contraintes telles que la rareté des ressources radio et leur nature partagée ainsi que la mobilité des nœuds. Dans ce travail, et par l'intermédiaire du fameux protocole BitTorrent, nous étudions l'impact de la mobilité des nœuds sur les performances d'une application de partage de fichiers, et montrons à travers des simulations qu'il est, dans ce cas, inutile de communiquer avec des nœuds lointains car le gain en diversité de piéces de données assuré par cette communication est négligeable devant celui obtenu par la mobilité. En plus, le débit d'envoi de piéces est trés faible sur des chemins longs. Par contre, dans le cas d'un réseau fixe, un envoi sélectif à des nœuds lointains est nécessaire pour l'obtention de bonnes opportunités de partage et une distribution égale des efforts.

Mot clés: P2P, MANET, BitTorrent

1 Introduction

A MANET is a set of mobile devices that communicate directly with each other via multi-hop wireless links. They dynamically and arbitrary form a transient network without requiring any centralized administration or infrastructure. With the proliferation of small mobile devices (laptops, smartphones, PDAs, etc.), these networks gained lot of importance and became a good tool to share content and experiences among users. Given the absence of a centralized server infrastructure in MANETs, the peer-to-peer paradigm is the most convenient to enable content sharing. This has other advantages. MANETs are infrastructureless networks of independent and equal users, having usually modest resources. Hence, applications designed for these networks must distribute fairly their loads among all the nodes and optimize the consumption of scarce resources. Peer-to-peer applications are important to make use of users' individual offerings in order to provide large-scale distributed services. For instance, a peer-to- peer file sharing application like BitTorrent [1] profits from the upload capacities of peers willing to download the file to quickly replicate it in the network. The idea is to split the file into a set of pieces so that each peer can exchange pieces it holds with other peers. Moreover, BitTorrent provides efficient incentives that encourage peers to upload their pieces to other peers. This way, the load of serving the file is fairly distributed among all the peers.

Whereas efficient content localization in wireless ad hoc networks has attracted considerable research interest, the content replication problem is still in its first steps [2] [3]. In a previous work [3], we have studied the challenges of deploying BitTorrent over fixed wireless ad hoc networks. We found that the neighborhood selection policy is the key of the efficiency of content sharing applications over such networks. In fact, in wireless ad hoc networks, where resources are shared and scarce and where distant multi-hop communications are often subject to bad performances and failures, it is mandatory to render any content sharing

* This work has been supported by the ITEA European Project on Experience Sharing in Mobile Peer Communities (Expeshare)

application aware of the location of its members. Our main observation, in line with what was observed in [2] on a BitTorrent-like application, was that peers should focus on their physical neighbors instead of exchanging pieces with a random set of peers. We have shown that limiting the neighborhood decreases the routing overhead while sending some pieces to faraway peers improves the piece diversity. This increases the sharing opportunities and subsequently ameliorates the completion time. Moreover, by behaving in this way, the load is better distributed among nodes. In this paper, we investigate the impact of the mobility of nodes on the performance of BitTorrent over MANETs. Our goal is to come up with a mobility-aware neighborhood selection policy. In this context, the mobility has a beast and a beauty. The beast is that it increases packet losses over long multi-hop paths, thus it makes it almost inefficient to send pieces to faraway peers. The beauty is that the mobility improves piece diversity since peers are continuously exchanging new pieces with the new peers they meet while moving across the network. To better understand this tradeoff brought by mobility, we add a BitTorrent module into the NS-2 network simulator [4] and simulate several scenarios where a set of mobile wireless nodes forming a MANET are interested in sharing some data file. We study how the download time and reciprocity of data evolve for different neighborhood selection policies. The results confirm that mobility can solve the piece diversity issue while the focus should be on reducing the routing overhead by limiting the scope of peers.

The remainder of this paper is organized as follows. Section 2 quickly overviews the BitTorrent protocol and describes our methodology. Section 3 is dedicated to the analysis of the simulation results. We summarize the paper and the future work in Section 4.

2 Background, context and methodology

We start from the classical Internet version of BitTorrent [1] and investigate, in MANETs, on the best neighborhood selection strategy. We want to understand the impact of the mobility of nodes on both the download time and the reciprocity of data between peers. In this paper and for lack of space, we focus on the challenging case where all nodes are peers. Indeed, in this case, the routing overhead is maximum since on one hand the volume of exchanged data is large, and on the other hand any packet sent over multiple hops will steal bandwidth from all intermediate nodes which are also peers. In the following, we start by giving an overview of the Internet version of the BitTorrent protocol, then we present our methodology.

Peers cooperating to download a file using BitTorrent form a sharing session called *torrent*. In order to allow a fast replication in the network, the file is split into several pieces and pieces are split into several blocks. The transmission unit is the block and the exchange unit is the piece: a peer cannot start uploading a piece to another peer until it holds all its blocks. A peer that owns a complete copy of the file is called *seed* and peers that are still downloading it are called *leechers*. Peers contact periodically a central server called *tracker* to get an up-to-date set of randomly selected peers that may include both leechers and seeds. This peer set, called neighborhood, consists of the other peers with whom a peer can exchange pieces. BitTorrent is basically composed of two core algorithms: the *choking algorithm*, which describes the choice of the effective neighbors with whom to exchange pieces, and the *rarest first algorithm*, which describes the piece selection strategy within the neighborhood. According to the periodic choking algorithm, a peer can establish simultaneously, at most, four outgoing TCP connections with peers of its neighborhood. It establishes three connections with the peers that upload to it at the highest rates, and the fourth connection is established with a randomly chosen peer. This last connection is dedicated to the discovery of new upload capacities, and the bootstrapping of new peers that have no pieces, while the other connections are chosen in a way that guarantees reciprocity and encourages collaboration among peers. On the other side, the rarest first algorithm aims to increase the diversity of pieces in the network. Indeed, a peer chooses to download, in priority, pieces that are less redundant in its neighborhood.

We note that, in the Internet version of BitTorrent, the neighborhood is chosen randomly. This was proven to work efficiently in the Internet since the bandwidth is abundant and communications between faraway nodes is not constrained by the scarcity of physical resources [5]. In contrast, and as has been shown in previous works [3] [2], this random neighborhood selection policy is not suitable for wireless ad hoc networks where communications between faraway nodes incur a high routing overhead, which impairs the performance of BitTorrent both in terms of download time and reciprocity of data. This result was observed for fixed wireless networks and the performance was shown to increase when the neighborhood of a peer is restricted to its physical neighbors, with few pieces sent from time to time to faraway nodes to boost diversity and improve sharing. In this paper, we go one step further by evaluating the impact of the mobility of nodes. The mobility is known to increase the capacity of wireless networks [6], we want to check how the neighborhood of file sharing applications needs to be defined to fully profit from the mobility.

Our study is performed through extensive simulations run with our implementation of the different policies of BitTorrent in the NS-2 network simulator [4]. Our scenario consists of a MANET of 50 nodes moving in a 500m x 80m bounded area. The node mobility patterns are generated using the random waypoint model. In this model, each node moves along a zigzag line from one waypoint to the next. The speed of nodes is taken equal to 2m/s, a typical pedestrian speed. For the pause time, different values were considered, we show here the results for 2s and 10s. The nodes connect to each other using the 802.11 MAC layer with the RTS/CTS-DATA/ACK mechanism enabled. The transmission range is set to 50m and the data rate to 1Mb/s. The routing operations are handled by the DSDV ad hoc routing protocol. In our simulations, we consider a flash crowd scenario where all the nodes start the download at the same time. Initially a node is randomly designated as a seed, the other nodes are leechers. We also implemented a distributed peer discovery mechanism that emulates the functionalities of the centralized tracker in the Internet version of BitTorrent. This way we are sure that the sharing continues independently of the tracker location. The shared file size is equal to 10Mbytes, and is subdivided into 100 pieces of equal sizes. The choking period is set to 40s.

3 Impact of the mobility: Results and analysis

In our study, we define two performance metrics: the average download time and the average sharing ratio. The first one measures the average download time per node over all the performed simulations. This metric reflects how fast the file is replicated in the network. The second metric measures the average sharing ratio per pair of peers over all the simulations. The sharing ratio between a couple of peers $(i; j)$ is defined as $R_{ij} = \frac{\min(D_{ij}, D_{ji})}{\max(D_{ij}, D_{ji})}$, where D_{ij} is the total number of bytes downloaded by peer i from peer j during the torrent lifetime. This ratio reflects how well peers reciprocate data with each other. A value near zero means an almost one-way propagation of the file inside the network. The ideal would be to get a value of this ratio equal to 1, which means a fair sharing between any neighboring couple of peers.

A peer is authorized to exchange pieces with other peers located within its neighborhood scope. In the Internet version of BitTorrent, the neighborhood scope is set to a large value (often 80 peers) and peers in this neighborhood are selected independently of the underlying topology (random selection from the entire list). In our scenario, this amounts to connecting to all network nodes and exchanging pieces with them on a multi-hop basis. To reduce the overhead caused by this standard configuration, we propose two topology-aware policies for neighborhood selection and compare them to the standard one. In the first one, the neighborhood scope is limited to some number of IP hops around each peer. A peer exchanges data with other peers within this scope but it does not provide any explicit effort for piece diversification beyond this scope. The second policy is similar to the first one with the only difference that the fourth outgoing TCP connection of each peer is dedicated to the exchange of data with peers beyond the limited neighborhood scope. The purpose is to improve piece diversification while limiting routing overhead. We experiment these policies in both fixed and mobile wireless ad hoc networks. The results are presented in Figure 1 where we plot the average download time versus the neighborhood scope (in number of IP hops), and in Figure 2 where we plot the average sharing ratio versus the neighborhood scope.

One can notice that in both fixed and mobile wireless ad hoc networks, large values of the neighborhood scope incur large download times (Figure 1) and small sharing ratios (Figure 2). We explain these results by the fact that in wireless environments, multi-hop TCP communications cause a considerable routing overhead and are likely to fail. The performances achieved in mobile networks for large scopes are even worse than those in fixed networks. In fact, multi-hop paths are more subject to failures because of the mobility. Hence, sending blocks to distant peers steals time and bandwidth from close peers without considerably improving sharing opportunities. We also notice that the more the value of the pause time is low, the more the performance is bad because the network becomes more dynamic and hence is more prone to path failures.

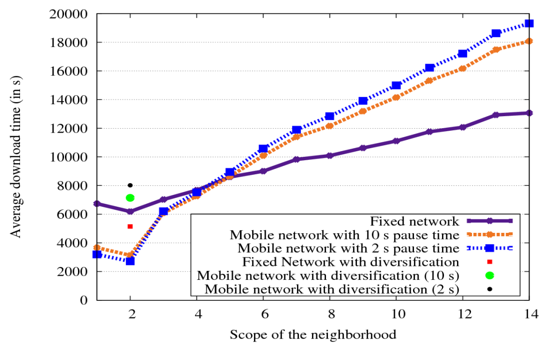

Fig. 1: Average download time Vs. Neighborhood scope

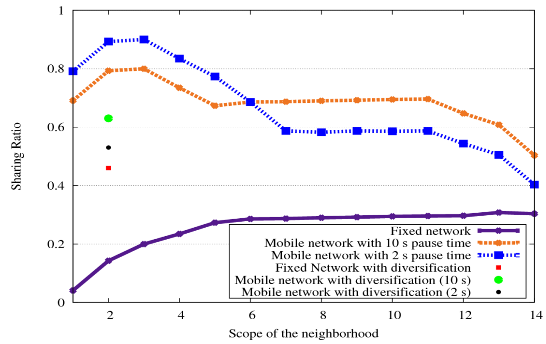

Fig. 2: Average sharing ratio Vs. Neighborhood scope

We attempt to reduce the network overhead by having low values for the neighborhood scope. Figures 1 and 2 show that, in fixed networks, low values of the scope improve the download time indeed, but yield the worst values for the sharing ratio. In fact, when the value of the neighborhood scope is small, pieces propagate in almost a unique direction from the initial seed to the farthest peers, which limits the reciprocation opportunities. It is for this main reason that the diversification policy was proposed for fixed networks [3]. The Figures show that limiting the neighborhood while sending some pieces to faraway nodes results in the best performances in fixed networks in terms of download time and sharing ratio. Surprisingly, when nodes are mobile, the best performances are achieved by only limiting the neighborhood scope to small values without any diversification effort. The sharing ratio has the highest values (approaching 0.8) and becomes even better when the pause time of nodes is reduced. These results can be easily linked to the diversity of pieces introduced by the mobility of nodes and the sharing opportunities it creates. One can conclude that in mobile environments only the scope of the neighborhood is to be limited. Furthermore, it is useless to provide any supplementary effort for the diversification because with the instability of paths, sending pieces to faraway peers in MANETs only increases the overhead and worsens the performances.

4 Conclusions and perspectives

The neighborhood selection policy is a fundamental building block in peer-to-peer content sharing applications. We believe it is the key for the success of such applications in challenging networks as MANETs. In this work, we investigate the impact of nodes mobility on the performance of BitTorrent when deployed over MANETs from the viewpoint of neighborhood selection. We show that sending pieces to remote peers is unnecessary because of the high loss rate caused by nodes' mobility and the important routing overhead. It is enough to exchange pieces in the nearby neighborhood and leave the mobility of nodes boost the diversity of pieces. Our future work will focus on validating this result by real experiments and analytical modeling, and on its extension to more realistic scenarios including cross traffic and non interested nodes.

References

[1] "BitTorrent protocol, http://wiki.theory.org/bittorrentspecification,"

[2] P.Michiardi and G.U.Keller, "Performance analysis of cooperative content distribution for wireless ad hoc networks, WONS," 2007.

[3] M.K.Sbai, C.Barakat, J.Choi, A.AlHamra, and T.Turletti, "Adapting BitTorrent to wireless ad-hoc networks, in Ad-Hoc Now, Sophia Antipolis," 2008.

[4] "NS: The Network Simulator, http://www.isi.edu/nsnam/ns/,"

[5] A.Legout, G.U.Keller, and P.Michiardi, "Rarest fisrt and choke algorithms are enough, IMC," 2006.

[6] M.Grossglauser and D.Tse, "Mobility increases the capacity of ad-hoc wireless networks, INFOCOM," 2009.