The Effect of Heterogeneous Link Capacities in BitTorrent-like File Sharing Systems

Francesca Lo Piccolo, Giovanni Neglia
D.I.E., Università degli Studi di Palermo, Italy
{francesca.lopiccolo,giovanni.neglia}@tti.unipa.it

Giuseppe Bianchi
D.I.E., Università degli Studi di Roma - Tor Vergata, Italy
bianchi@elet.polimi.it

Abstract

Despite the recent research efforts on Peer-to-Peer networks modeling and performance evaluation, very few literature results have appeared on what concerns the effect of network characteristics on the overlay peer-to-peer services. Goal of this paper is to contribute in assessing the effect of different access link capacities on performance of a Peer-to-Peer file sharing system. In order to make the problem tractable, an analytical model referring to a simplified network model with access links belonging to two different capacity classes is developed. Our preliminary results seem to imply that, depending on evaluation metrics, bandwidth heterogeneity can have a positive effect on content propagation among peers.

1. Introduction

Peer-to-Peer (P2P) paradigm has recently emerged as new and fundamental model for distributed networked services and applications. P2P overlays have been deployed in many different application areas, such as distributed grid computing [2], storage [3], web cache [4], service directory [5, 6, 7]. But file sharing networks are perhaps the most popular P2P application: many different P2P file sharing systems, such as Gnutella, Kazaa, eDonkey, BitTorrent, exist and collect million of users everywhere in the world. So it's not casual that the dominant traffic type observed by Internet Service Providers (ISPs) is associated with P2P file sharing applications, as SD-NAP trace [1] reveals.

Most of P2P research on file sharing systems, perhaps the most popular P2P application, has so far emphasized system design and traffic measurement. Only a few works have tried to quantitatively evaluate performance by mathematical models. The work in [13] is maybe the first one to propose a mathematical abstraction for P2P systems: a general P2P file sharing system is modelled by means of a multiple class closed queueing network and the dependence of stationary performance on parameters like peer request rate and peer number in the system is analyzed. In [14], a branching process and a simple Markovian model are used to study service capacity of BitTorrent-like file sharing systems respectively in transient regime and in steady state regime. Our paper follows such research direction too. In particular it aims to study the effect of capacity heterogeneity on P2P file sharing.

Extensive measurement studies [8, 9, 10] have been performed on P2P file sharing systems and have highlighted a significant amount of heterogeneity in P2P infrastructure: as it is well documented in [10], bandwidth, latency, availability and degree of shared contents can vary from three to five orders of magnitude across the peers in Gnutella and Napster networks. As regards peer bandwidth, in [10] Saroiu et al. have reported that it is possible to distinguish between two different kinds of heterogeneity. On the one hand, there is asymmetry in upstream and downstream bandwidth of many peers. On average a peer tends to have downstream bandwidth higher than upstream bandwidth; this is due to large fraction of peers depending on asymmetric links, such as ADSL lines, cable modems or regular modems using V90 protocol. On the other hand, measures performed on Gnutella network have revealed 8% of users connected by means of modem (64 Kbps or less), 60% using broadband connections (DSL, cable, T1, T3), 30% having very high bandwidth (at least 3 Mbps) connections (data relates to downstream bandwidth).

On basis of such results we propose to characterize impact of heterogeneity of peer bandwidth on P2P file sharing systems; it is to be specified that we ignore the effect of the possible asymmetry in upstream and downstream bandwidth of a single peer and we assume every peer having symmetric connection link. In particular we analyze effects of bandwidth heterogeneity on file transfer dynamics and content diffusion process by developing a simple fluid model. In doing this, we make three assumptions: i) we consider a file sharing P2P architecture, which allows peers to download chunks of the same content from different origins, ii) there are only two possible bandwidth values in the networks, iii) we focus on the diffusion of a single file. Preliminary results so achieved seems to demonstrate that, depending on metrics used to compare bandwidth heterogeneous systems with bandwidth homogeneous systems, file diffusion dynamics can take naturally advantage of capacity heterogeneity.

1.1. Related work

Bandwidth heterogeneity issue has already been explored from different points of view. For example new solutions enabling to accommodate natural capacity heterogeneity existent in P2P file sharing systems are proposed in [11, 12] in order to gain better scaling properties. Differently from [12, 11], we base on existent P2P systems and we propose no new and smart solution improving performance; our intention is instead to point out bandwidth heterogeneity impact on P2P file sharing networks and to show if and under which conditions bandwidth heterogeneity can intrinsically improve P2P file sharing performance.

To our knowledge, capacity heterogeneity issue has been considered from a mathematical point of view only in [14]. From some results of branching process theory, the authors incidentally show that if file transfer times are random variables, the greater the variance of the transfer time, the faster the file diffusion process in P2P network. More precisely, if we consider two different transfer time distributions, i.e., $T^{(1)}$ and $T^{(2)}$ such that $E[T^{(1)}] = E[T^{(2)}]$ and there is an increasing convex ordering (I.C.X.) on $T^{(1)}$ and $T^{(2)}$, i.e., $T^{(1)} \leq^{icx} T^{(2)}$, then diffusion process is faster with $T^{(2)}$. Note that $T^{(1)} \leq^{icx} T^{(2)}$ and $E[T^{(1)}] = E[T^{(2)}]$ imply $Var[T^{(1)}] \leq Var[T^{(2)}]$. Our study differs from such approach for two reasons: firstly we consider homogeneous and heterogeneous networks under different equivalence criteria (not only for a given mean transfer time); secondly in branching process model proposed in [14] transfer times are assumed to be independent identically distributed random variables, disregarding correlation due to the access link bandwidths.

In order to study the effect of bandwidth heterogeneity on P2P file sharing systems we have developed a simple fluid model. This model is an extension of the model presented in [15], that is in its turn partially motivated by the Markovian model proposed in [14]. The new model takes into account two different classes of peers, associated to two different values of service capacity. Such extension has brought about an increase of model differential equation number (from two to four) and has made some equations non linear. Besides presence of two different classes of peers has required to explicit download and upload policies of peers. The model in [15] has been considered to obtain performance indexes for the homogeneous scenario, to be compared with those ones of the heterogeneous scenario.

The paper is organized as it follows. Section 2 describes the simple fluid model we propose, while the effect of heterogeneity is discussed in section 3 via both analytical and numerical analysis. Finally, conclusive remarks and further research issues are given in section 4.

2. Fluid Model Description

In this section, we propose a fluid model able to capture the effect of heterogeneous access link rates in BitTorrent-like P2P file sharing systems. To keep the model simple, just two different link rates are considered. Moreover, we assume symmetric link capacity, i.e. the link rate to be the same in both upstream and downstream directions. The model foundation is based on the fluid approach presented in [15] and briefly described in the following subsection. Basic assumption in [15] is that all peers are connected by means of access links of equal capacity; in reason of this, we denote this model as single capacity model.

2.1. Single Capacity Model

The idea is to model temporal evolution of the number of users that are downloading or have completed to download a given file. Following BitTorrent jargon [3], these users will be referred to as either downloaders or seeds, respectively. Number of seeds and downloaders in the system at time t are $x(t)$ and $y(t)$. Each peer comes into the system only once, tries to download the only present file (or the file we are focusing on) and leaves the system either since too much time has been elapsed from the connection beginning and it is not willing to wait anymore, or since download is completed. In particular, peer arrivals are modelled in according to a Poisson process with rate λ. Moreover θ represents the rate at which downloaders abort their download and μ is the normalized bandwidth, i.e. measured in files/seconds, of the peers (identical full duplex links with equal bandwidth in both the directions). If a peer succeeds in completing download before disconnection, that is if a peer succeeds in becoming seed, it leaves system after a certain time amount exponentially distributed with mean $1/\gamma$. The rate at which downloaders become seeds is r; we are going to specify r in what follows.

Hence evolution of the number of peers is given by

$$x'(t) \quad = \quad \lambda - \theta x(t) - r \qquad (1)$$

$$y'(t) = r - \gamma y(t)$$

In a BitTorrent-like P2P network, a downloader can upload data to other peers even though it may only have parts of a file (chunks). The parameter η is used to indicate the effectiveness of this file sharing. The number of peers that can contribute to file upload is equal to $\tilde{y} = y + \eta x$. By setting $\eta = 0$, we assume that a randomly chosen chunk will be never found on a downloader. Conversely, by setting $\eta = 1$ we assume that the contribution to the file sharing given by a downloader is equal to the contribution given by a seed. It has been shown in [15] that, with realistic file sizes and BitTorrent typical chunk size of 256 KB, chunk number per file is of the order of several hundreds and it assures η to be very close to 1.

As regards parameter r, the biggest advantage of this approach is the possibility of referring to aggregate downloading rate: r represents, in fact, aggregate rate at which downloaders become seeds. In particular if there is no constraint due to uploading bandwidth (case A), total download rate of the system can be expressed as $r = \mu x(t)$. Conversely when there is constraint due to uploading bandwidth (case B), total download rate will be $r = \mu\tilde{y} = \mu(y+\eta x)$. In general we can write $r = \min(\mu x(t), \mu(\eta x(t) + y(t)))$.

We find useful to refer to the maximum uploading bandwidth $u = \mu\tilde{y}$ and the maximum downloading bandwidth $d = \mu x$. Note that u and d are time dependent. Cases A and B are characterized respectively by $d < u$ and $d > u$.

2.2. Steady State Analysis

In order to study the system in steady state solutions we can let the derivatives be equal to zero:

- case A, when $d < u$

$$\bar{x} = \frac{\lambda}{(1+\alpha)\mu} \qquad \bar{y} = \frac{\lambda}{(1+\alpha)\gamma} \tag{2}$$

- case B, when $d > u$

$$\begin{aligned} \bar{x} &= \frac{\lambda(\gamma - \mu)}{\mu(\alpha\gamma + \gamma\eta - \alpha\mu)} \\ \bar{y} &= \frac{\lambda\eta}{\alpha\gamma + \gamma\eta - \alpha\mu} \end{aligned} \tag{3}$$

In the previous relations we have assumed $\theta = \alpha\mu$ with $0 < \alpha < 1$: θ is in fact reasonable to be lower than μ, otherwise on average downloaders would abandon their download attempts before completion.

Relation between two described cases can be advantageously shown in the plane (η,γ), as in figure 1. From simple calculations it follows that $d < u$, if only if $\gamma < \frac{\mu}{(1-\eta)}$. Hence the curve $\gamma = \frac{\mu}{(1-\eta)}$ is the boundary between the two cases.

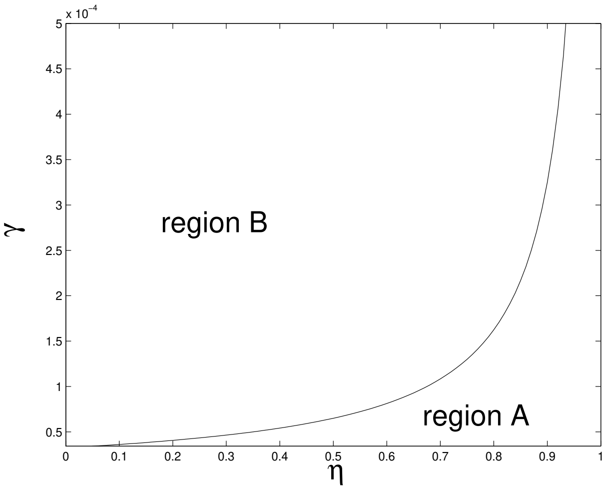

Figure 1. Transitions among cases A and B for the Single Capacity Scenario

2.3. Two Capacities Model

In this situation an user will connect to the system through an access link which may assume two possible capacity values. Hereafter, we will refer to these two link capacities as low-rate (l) and high-rate (h). Equations (1) can be easily extended as

$$\begin{aligned} x_l'(t) &= \lambda_l - \theta_l x_l(t) - r_l \\ x_h'(t) &= \lambda_h - \theta_h x_h(t) - r_h \\ y_l'(t) &= r_l - \gamma_l y_l(t) \\ y_h'(t) &= r_h - \gamma_h y_h(t) \end{aligned} \tag{4}$$

where all the variables have the same meaning as in (1), but they refer to the class of users with low-rate access links or high-rate access links.

In order to evaluate r_l and r_h, some considerations are required. We assume that P2P architecture is able to support *parallel download* of chunks of the same content from different peers, as BitTorrent does. We are going to discuss this assumption at the end of the section.

Figure 2 shows the network abstraction. Downloaders (x_h,x_l) are on the left, while uploaders ($\tilde{y}_h$,$\tilde{y}_l$) are on the right. Rate constraints are due to access links, not to the core network (the same assumption holds for the previous model).

We need to characterize how potential uploading resources are shared between slow and fast users. There is not such need for the single capacity model, because the users are indistinguishable and, whatever the resources assignment criteria are, aggregate download rate does not change if we assume that no resource is wasted. Resource waste can

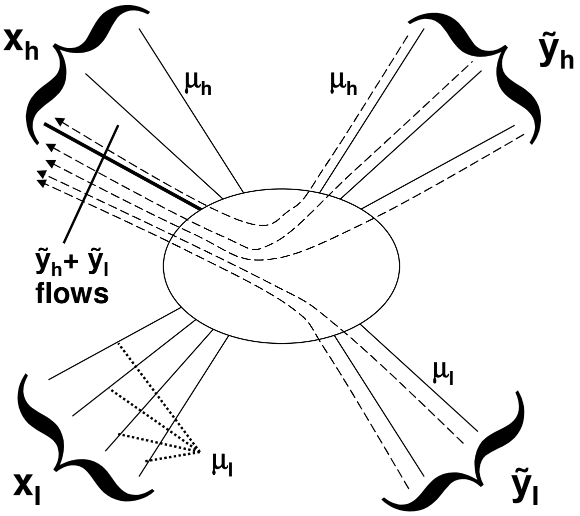

Figure 2. Network topology

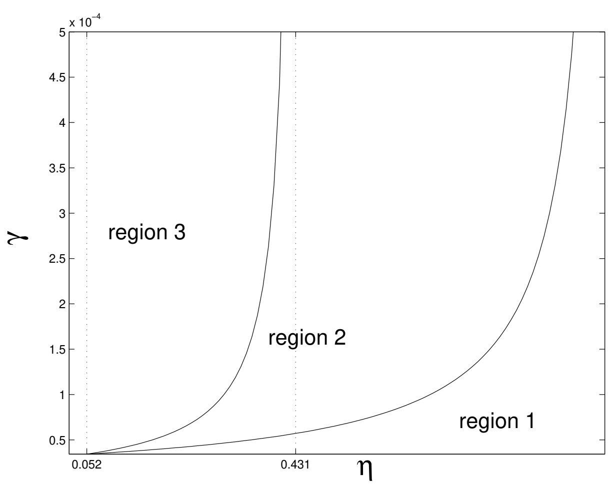

Figure 3. Transitions among cases 1, 2 and 3 for the Two Capacities Scenario

occur if peer knowledge of the resources available in the network is not complete due, for example, to some propagation delay of information or to limitations on signalling.

We implicitly derive no resource waste and reasonable resource sharing criteria for the two capacities scenario by the following assumptions:

- each peer establishes connections with all the $\tilde{y}_l + \tilde{y}_h$ uploaders;
- bandwidth of each connection is determined according to the max-min fair criteria.

Assumption on the number of connections simplifies the calculations. We are however interested into aggregate download rates for both the classes of users and these rates should not differ from those ones evaluated here, if the number of admissible parallel downloads is high enough to assure that no potential resources are wasted.

The particular symmetry of network in figure 2 and of its flows distribution leads to the following considerations: i) the links are divided in four groups (x_h links with capacity μ_h, $\tilde{y}_h$ links with capacity μ_h, x_l links with capacity μ_l and $\tilde{y}_l$ links with capacity μ_l) and rate across each link of a group is the same, in particular the links of the same group are all saturated or all unsaturated, ii) at least the two groups of downloaders or the two groups of uploaders have to be saturated, iii) if three groups are saturated, then both the groups corresponding to slow peers are saturated.

As consequence, because our interest is to evaluate aggregate rates r_l and r_h at which downloaders become seeds, only three cases are relevant: case 1, both the groups of downloaders are saturated; case 2, the group of slow downloaders is saturated, while the group of fast downloaders is not saturated; case 3, no dowloaders group is saturated. Note that, while in the first case the $\tilde{y}_l + \tilde{y}_h$ uploaders are not fully saturated, in the second and in the third case all the $\tilde{y}_l + \tilde{y}_h$ uploaders are saturated.

By generalization of the above notation, the maximum uploading bandwidth of the system is $u = \mu_l \tilde{y}_l + \mu_h \tilde{y}_h$ and the maximum downloading bandwidth is $d = \mu_l x_l + \mu_h x_h$.

It is possible at this point to evaluate rates r_l and r_h by considering the three cases previously described.

- case 1 when $u \geq d$
  Total downloading capacity is constraining bottleneck and total uploading capacity is sufficient to saturate both low rate links and high rate links. As consequence of this,
  $$r_l = \mu_l x_l(t) \qquad r_h = \mu_h x_h(t)$$
- case 2 when $\mu_l(x_l + x_h) < u < d$
  Total uploading capacity is constraining bottleneck, but it is sufficient to saturate at least low rate links. In this case
  $$r_l = \mu_l x_l(t) \qquad r_h = u - \mu_l x_l(t)$$
- case 3 when $u \leq \mu_l(x_l + x_h)$
  Total uploading capacity is constraining bottleneck and it is sufficient to saturate neither low rate links nor high rate links. In this case
  $$r_l = \frac{u}{x_l(t) + x_h(t)} \cdot x_l(t)$$
  $$r_h = \frac{u}{x_l(t) + x_h(t)} \cdot x_h(t)$$

As a final remark, note that *parallel download* capability does not necessarily correspond to capability to *upload* chunks of an uncomplete content. A system able to manage chunks will in general permit both the actions, but from the modelling point of view there are some differences. In particular in our model uploading of chunks of an uncomplete content can be disabled by simply considering $\eta = 0$ in equation system, while effect of parallel downloading is intrinsically considered in the expression of r_l and r_h.

2.4. Steady State Analysis

In order to simplify equation mathematical solution, it has been performed under the following assumptions:

- low-rate and high-rate links are uniformly distributed among peer population which downloaders come from; so it is possible to assume $\lambda_l = \lambda_h = \widehat{\lambda}$;
- as in the single capacity model, θ_l (θ_h) is reasonable to be lower than μ_l (μ_h); therefore we assume $\theta_l = \alpha\mu_l$ ($\theta_h = \alpha\mu_h$) with $0 < \alpha < 1$;
- peer service capacity should not affect seed departure rate; this involves $\gamma_l = \gamma_h = \gamma$.

Results of steady state study are reported by distinguishing among the three cases before singled out.

1. $u \geq d$
   Relate solutions are
$$\bar{x}_l = \frac{\hat{\lambda}}{(1+\alpha)\mu_l}, \qquad \bar{y}_l = \frac{\hat{\lambda}}{(1+\alpha)\gamma} \qquad (5)$$
$$\bar{x}_h = \frac{\hat{\lambda}}{(1+\alpha)\mu_h}, \qquad \bar{y}_h = \frac{\hat{\lambda}}{(1+\alpha)\gamma}$$
2. $\mu_l(x_l + x_h) < u < d$
   Corresponding solutions are
$$\begin{aligned} \bar{x_l} &= \frac{\hat{\lambda}}{(1+\alpha)\mu_l} && (6)\\ \bar{x_h} &= \hat{\lambda}\frac{(1+\alpha)\mu_h + \mu_l + (\eta - 2 - \alpha)\gamma}{(1+\alpha)\mu_h[\alpha(\mu_h - \gamma) - \gamma\eta]} \\ \bar{y_l} &= \frac{\hat{\lambda}}{(1+\alpha)\gamma} \\ \bar{y_h} &= \hat{\lambda}\frac{\gamma(\alpha - \eta - 2\alpha\eta) - \alpha\mu_l}{(1+\alpha)\gamma[\alpha(\mu_h - \gamma) - \gamma\eta]} \end{aligned}$$
3. $u \leq \mu_l(x_l + x_h)$
   This is the most complicated case from resolution point of view due to equation non-linearity. In what follows we show only numerical results for such case.

While in the previous case representation in the plane (η,γ) could appear trivial, here it can be usefully employed to understand the phenomenon, and to overcome analytical difficulties to manage and to solve system equations, in particular for case 3, where no closed formula has been found. Boundaries between the three regions -each corresponding to one of the above cases- can be easily found as follows.

- For boundary between region 1 and region 2 we can impose $d = u$ and we obtain the following analytical formula for the curve:
$$\gamma = \frac{\mu_l + \mu_h}{2(1-\eta)}$$
  similar to that one found for the single capacity model.
- For boundary between region 2 and region 3 we can impose $u = \mu_l(x_l + x_h)$, and we obtain the following analytical formula:
$$\gamma = \frac{\mu_l(\mu_h + 2\alpha\mu_h + \mu_l)}{2\mu_l + \alpha(\mu_h + \mu_l) - \eta(\mu_h + 2\alpha\mu_h + \mu_l)}$$
  for $\eta > \eta' = \frac{\alpha(\mu_h - \mu_l)}{\mu_h + 2\alpha\mu_h + \mu_l}$.

Figure 3 shows the above results for $\eta' < \eta < 1$. The range $0 < \eta < \eta'$ has not been explored, due to unavailability of closed formula solutions foe case 3, and we suspect that in this range region 1 borders on region 3. Anyway it can be neglected because η' is significantly less than 1 and in [15] η has been proven to be almost equal to 1. For example numbers in figure 3 have been obtained with the parameters specified in section 3.4. It appears that $\eta' \approx 0.052 \ll 1$.

The figure shows some qualitative differences in comparison to the single capacity scenario. First of all it is possible to note the presence of region 2 between regions 1 and 3 that correspond to regions A and B in figure 1. Secondly, referring to asymptote of boundary curve between region 2 and region 3 as $\eta = \eta''$, it appears that for $\eta' < \eta < \eta''$ increasing γ leads the system from case 1 to case 3 going through case 2. On the contrary for $\eta'' < \eta < 1$ case 3 is unreachable.

3. Performance evaluation

In this section we present some preliminary results about the impact of heterogeneous link capacities on the performance of a file sharing architecture. The outline of this section is the following. Firstly we investigate the key issue of equivalence between homogeneous and heterogeneous scenarios: when is the comparison meaningful? Secondly we present and justify the performance index we adopted in accordance to [15]. Then we analytically compare the performance of heterogeneous and homogeneous networks in a specific situation (case 1 for heterogeneous network and case A for homogenous one), showing that results depend on equivalence criteria chosen; finally we show by numerical evaluation that the same results hold also in other situations.

3.1. Equivalence Criteria

In order to compare single capacity model with two capacities model, we need to accurately choose corresponding parameters. Some choices appear quite reasonable. For example, we should compare:

- under the same arrival rate, i.e. $\lambda = \lambda_l + \lambda_h = 2\hat{\lambda}$;
- under the same or correspondent departure rate, i.e., $\gamma = \gamma_l = \gamma_h$, and equal α values.

Much more questionable is how μ, μ_h and μ_l should be related. In what follows we consider two different equivalence criteria, we will refer to in next sections.

Perhaps the most simple idea is to compare the two scenarios, under the same mean capacity value: $\lambda\mu = \lambda_l\mu_l + \lambda_h\mu_h$, i.e. under our assumptions $\mu = (\mu_l + \mu_h)/2$. We denote such equivalence with subscript *m*.

Another idea is to compare the scenarios, given the same average transfer time per unity of content, for example per chunk. The idea is that network offers transfer of a file (or a chunk) between two users (hence according to a client-server paradigm) as a basic service. According to this point of view, we consider two networks that offers identical average basic services, i.e. identical average transfer time, and we evaluate increase in transfer speed due to the P2P architecture. If we consider network in figure 2, where only access links can limit connection bandwidth, then if a pair of users is randomly selected, connection bandwidth will be equal to μ_h with probability $p = (\frac{\lambda_h}{\lambda_h+\lambda_l})^2$ and equal to μ_l with probability $1 - p$. Under assumption $\lambda_l = \lambda_h$ it results $p = 1/4$. So equivalence requires $1/\mu = 1/(4\mu_h) + 3/(4\mu_l)$. We denote such equivalence with subscript *t*. By the way the result in [14] about diffusion speed in branching assumed the same average value for the single transfer time.

It is evident that these different equivalence criteria lead to radically different parameters for the comparison. For example, given μ_l and μ_h, being μ_m and μ_t the capacity values for the homogeneous scenario according respectively to the first and the second criterion, it holds $\mu_t < \mu_m$.

3.2. Performance Index

As a performance index of P2P network good operation we consider the average number of downloaders in the network: the lower this number the better the network status. Intuitively, rate at which downloaders complete transfer and become seeds is equal -in steady state regime- to difference between arrival rate at the system (λ) and departure rate of unsatisfied downloaders (θx); hence the lower x the higher the number of users who complete file transfer in a certain time interval. From the dual user point of view, in [15] authors by the Little's law evaluate the average downloading time for a peer in steady state as $T = \frac{x}{\lambda}$. Again as number of users decreases, performance in terms of downloading time improves. As regards the two capacities scenario, it can be proven similarly that average downloading time is equal to $T_l = \frac{x_l}{\lambda}$ and $T_h = \frac{x_h}{\lambda}$ respectively for low-rate users and high-rate ones. Average downloading time could be evaluated as

$$T = \frac{\lambda_l}{\lambda_l + \lambda_h}\frac{\lambda_l - \theta x_l}{\lambda_l}T_l + \frac{\lambda_h}{\lambda_l + \lambda_h}\frac{\lambda_h - \theta x_h}{\lambda_h}T_h$$

where $\frac{\lambda_l}{\lambda_l+\lambda_h}$ ($\frac{\lambda_h}{\lambda_l+\lambda_h}$) is the probability that a new peer is a low-rate (high-rate) peer and $\frac{\lambda_l-\theta x_l}{\lambda_l}$ ($\frac{\lambda_h-\theta x_h}{\lambda_h}$) is the probability that such low-rate (high-rate) user completes download. Under our assumption, in case 1 it holds $T \propto (x_l + x_h)$. This relation is not in general true, anyway we have considered $x_l + x_h$ as the performance index in the two capacities scenario.

3.3. Analytical Results

Here we limit to case 1 for the two capacities model and case A for the single capacity one. Let μ and x be respectively bandwidth -according to the criterion chosen- and average number of downloaders of homogeneous system, and μ_l, μ_h, x_l, x_h the analogous quantities for heterogeneous network; it can be easily shown by the formula for the steady state that

$$x < x_l + x_h$$

i.e. P2P paradigm is more effective in a homogeneous scenario, if and only if the equivalent bandwidth is

$$\mu > \frac{2\mu_l\mu_h}{\mu_l + \mu_h}$$

With reference to the two equivalence criteria we have considered in section 3.1, it holds that

$$\mu_m > \frac{2\mu_l\mu_h}{\mu_l + \mu_h}$$

while on the contrary

$$\mu_t < \frac{2\mu_l\mu_h}{\mu_l + \mu_h}$$

We can read such results from two different points of view. The result on μ_m says that if we are dimensioning a network in order to achieve optimal performance for P2P transfer and we assume that we can distribute among links a fixed amount of bandwidth (or equivalently we have a certain amount of money and link prices increase linearly with the link bandwidth), then the best we can do is to choose all the links equal. The result on μ_t says that if we are considering two existing networks, where contents are spread

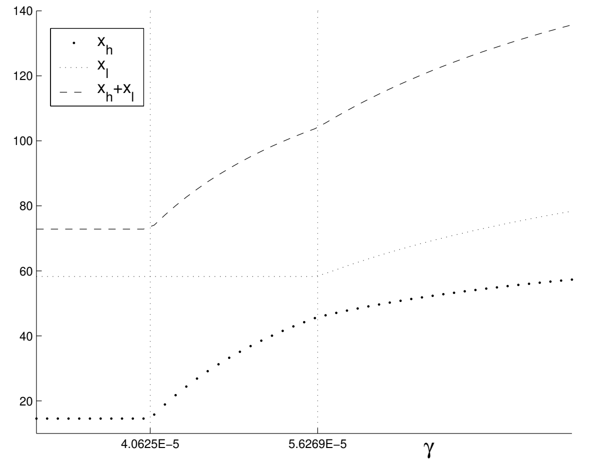

Figure 4. Average number of downloaders in heterogeneous scenario for $\eta = 0.2$

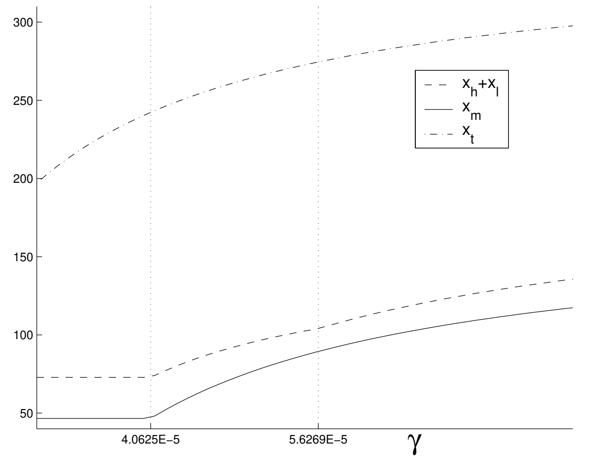

Figure 5. Comparison between heterogeneous and homogeneous networks for $\eta = 0.2$

according to a server/client paradigm, and the average completion time in the two networks is the same, network with the higher variability of link bandwidths will obtain the greater improvement from the adoption of a P2P file sharing paradigm.

Clearly results we have shown in this section are quite limited, because they refer to a two capacities scenario and a particular way to operate, essentially when constraint is the downloading bandwidth. As regards such constraint, it is what it happens usually when efficiency of data transfer from other downloaders is high ($\eta \approx 1$), as it appears from calculation in [15].

3.4. Numerical Results

In order to investigate effect of heterogeneity in different conditions, we compared, for typical values of the parameters, the performance of P2P networks in terms of chosen performance index. We consider the two equivalence criteria specified above.

As regards parameters values, we considered realistic values. Results shown have been obtained for $\mu_h = 5.2 * 10^{-5}s^{-1}$, $\mu_l = 1.3*10^{-5}s^{-1}$, $\alpha = 0.1$ and $\lambda = 1/600s^{-1}$.

Figure 4 shows x_l, x_h and $x_l + x_h$ versus γ in steady state for $\eta = 0.2$. For the particular values considered it holds $\eta' \approx 0.052$ and $\eta'' \approx 0.431$, hence η belongs to range where as γ increases system goes from region 1 to region 3 across region 2. The boundaries crossing is evidenced by the two vertical dotted lines, but it is also revealed by the change of curves trends. In particular during the first transition (from region 1 to region 2) high-rate links of downloaders are no more saturated, hence x_h starts increasing (because downloading time increases, see [15]), while x_l does not change. When γ exceeds a certain threshold, the uploaders are no more able to saturate low-rate downloaders and also x_l starts increasing.

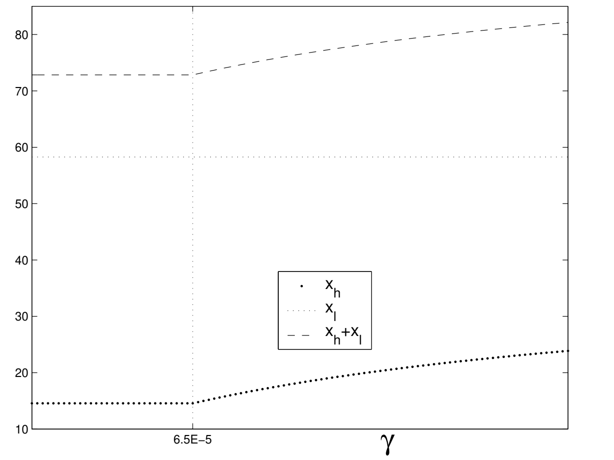

Figure 6. Average number of downloaders in heterogeneous scenario for $\eta = 0.5$

Figure 5 compares average number of downloaders of the heterogenous network ($x_l + x_h$) and of two homogeneous networks (x_m and x_t when the bandwidth are respectively μ_m and μ_t). Results confirm those analytically obtained in section 3.3: if bandwidth is equal to μ_m, then there is an improvement, while the opposite is true when μ_t is considered.

Figure 6 is analogous to figure 4 but here $\eta = 0.5$ is considered, hence $\eta > \eta''$ and system goes from region 1 to re-

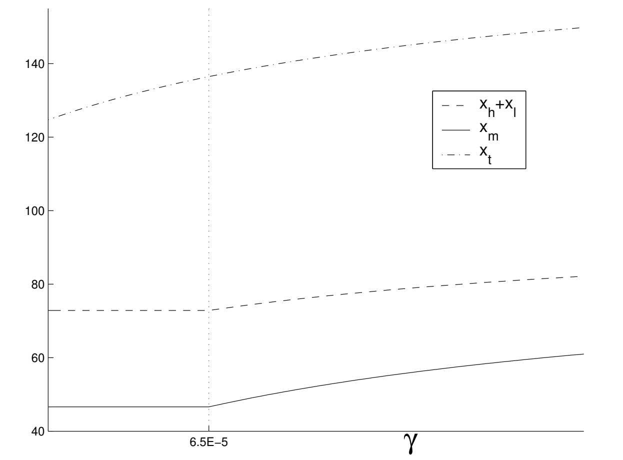

Figure 7. Comparison between heterogeneous and homogeneous networks for $\eta = 0.5$

gion 2, but it cannot reach region 3. In fact the curve shows only one change in its behavior and x_l is constant for all the values of γ, because low-rate links are always saturated.

Like figure 5, figure 7 compares average number of downloaders of the heterogenous network ($x_l + x_h$) and of two homogeneous networks (x_m and x_t when the bandwidth are respectively μ_m and μ_t), but it refers to $\eta = 0.5$. It confirms results previously described.

4. Conclusions and further research issues

This paper investigates the effect of different access link capacities on the performance of a Peer-to-Peer file sharing system. A simple fluid model with access links belonging to two different capacity classes has been developed. Preliminary results show that bandwidth heterogeneity can have a positive effect on content propagation among peers, and in general it depends on the equivalence criteria chosen. Future research will investigate more deeply these aspects, and the effect of other forms of heterogeneity, like file sizes and users behaviors.

References

[1] *Top applications (bytes) for subinterface: Sd-nap traffic*, in CA/DA workload analysis of SD-NAP data, http://www.caida.org/analysis/workload/byapplication/sdnap/index.xml, 2002

[2] Entropia, http://www.entropia.com

[3] B. Cohen, *Incentives build robustness in BitTorrent*, May 2003

[4] F. Dabek, M.F. Kaashoek, D. Karger,R. Morris, I. Stoica, *Wide-area cooperative storage with CFS*, in Proceedings of the 18th ACM Symposium on Operating Systems Principles (SOSP '01), Chateau Lake Lou, Banff, Canada, October 2001

[5] S. Iyer, A. Rowstron, P. Druschel, *Squirrel: a decentralized peer-to-peer web cache*, in Proceedings of ACM Symposium on Principles of Distributed Computing (PODC '02), Monterey, California, 2002

[6] S. Ratnasamy, P. Francis, P.,M. Handley, R. Karp, S. Shenker, *A scalable content addressable network*, in Proceedings of ACM SIGCOMM, 2001

[7] I. Stoica, R. Morris, D. Karger, M. Kaashoek, H. Balakrishman, *Chord: a scalable peer-to-peer lookup protocol for internet applications*, in Proceedings of ACM SIGCOMM, 2001

[8] E. Adar, B.A. Huberman, *Free Riding on Gnutella*, in First Monday, volume 5, number 10, October 2000

[9] M. Ripeanu, I. Foster, A. Iamnitchi, *Mapping the Gnutella Network: Properties of Large-Scale Peer-to-Peer Systems and Implications for System Design*, in IEEE Internet Computing Journal special issue on peer-to-peer networking, vol. 6(1), 2002

[10] S. Saroiu, P.K. Gummadi, S.D Gribble, *A Measurement Study of Peer-to-Peer File Sharing Systems*, in Proceedings of Multimedia Computing and Networking 2002 (MMCN'02), January 2002

[11] Y. Chawathe, S. Ratnasamy, L. Breslau, S. Shenker, *Making Gnutella-like P2P Systems Scalable*, in Proceedings of ACM SIGCOMM, 2003

[12] Q. Lv, S. Ratnasamy, S. Shenker, *Can heterogeneity make Gnutella scalable?*, in 1st Int. Workshop on Peer-to-Peer Systems (IPTPS), Cambridge, March 2002

[13] Z. Ge, D. Figueiredo, S. Jaiswal, J. Kurose, D. Towsley, *Modeling peer-peer file sharing systems*, in Proceedings of Infocom, 2003

[14] X. Yang, G. de Veciana, *Service Capacity of Peer to Peer Networks*, in Proceedings of IEEE INFOCOM, 2004

[15] D. Qiu, R. Srikant, *Modeling and Performance Analysis of BitTorrent-Like Peer-to-Peer Networks*, in Proceedings of ACM SIGCOMM, Portland, OR, September 2004

[16] D. Bertsekas, R. Gallager, *Data Networks*, Prentice Hall, Englewood Cliffs, NJ, 1987