Available online at www.sciencedirect.com
ScienceDirect
Performance Evaluation 64 (2007) 819–837
PERFORMANCE EVALUATION
An International Journal
www.elsevier.com/locate/peva

A modeling framework to understand the tussle between ISPs and peer-to-peer file-sharing users[☆]

M. Garetto[a], D.R. Figueiredo[b], R. Gaeta[a], M. Sereno[a,*]

[a] *Dipartimento di Informatica, Università di Torino, Torino, Italy*
[b] *PESC/COPPE Federal University of Rio de Janeiro, Rio de Janeiro, Brazil*

Available online 20 June 2007

Abstract

Recent measurement studies have shown that traffic generated by peer-to-peer (P2P) file-sharing applications has started to dominate the bandwidth consumption on Internet access links. The prevailing use of P2P applications carries with it significant implications for Internet Service Providers (ISPs): on the one hand increased levels of P2P traffic result in additional costs for an ISP, which has to provide a satisfactory service level to its subscribers. On the other hand, P2P applications are a major driving force for the adoption of broadband access, which is a significant source of revenue for the ISPs. A successful strategy to manage P2P traffic must address both the ISP perspective of costs and the subscriber perspective of quality of service. While several practical solutions have been identified to manage P2P traffic in a network, no analytical studies have been proposed so far to evaluate their effectiveness in specific contexts. In this paper we propose a modeling framework that allows the optimal strategy to be identified for an ISP as a function of the several factors that come into play. In particular, our model shows that P2P-friendly solutions become lucrative when the ISP can attract a sufficiently large number of subscribers. Our modeling framework also illustrates several other interesting phenomena that occur in the tussle between the ISP and its subscribers.

Keywords: Peer-to-peer sharing systems; Analytical modeling; Traffic management; Utility-based modeling

1. Introduction

Peer-to-Peer (P2P) file-sharing applications have reached a level of popularity among Internet users such that the volume of traffic generated by such applications nowadays represents a significant fraction of the overall Internet traffic [1]. Bandwidth consumption due to P2P file sharing applications continues to grow unabatedly, edging ahead of the traffic volume generated by traditional Internet applications (i.e. web browsing, email, etc.) [2,3].

As P2P file-sharing applications gain popularity, Internet service providers (ISPs) are increasingly faced with the problem of how to manage the vast amount of traffic generated by such applications. Indeed, P2P traffic represents

☆ This work has been partially supported by the Italian Ministry for University and Research (MIUR) within the frameworks of the FAMOUS (PRIN) and of the PROFILES (PRIN) projects.

* Corresponding author. Tel.: +39 011 6706718; fax: +39 011 751603.
E-mail address: matteo.sereno@di.unito.it (M. Sereno).

0166-5316/$ - see front matter
doi:10.1016/j.peva.2007.06.003

a huge problem for last-mile providers, where it makes up 80% or more of the traffic volume, according to some reports [1]. Such traffic levels can cause congestion and overall performance degradation in an ISP's network, ultimately leading to customer dissatisfaction. At the same time, P2P file-sharing applications are a major driving force for the adoption of residential broadband access, which is a significant source of revenue to ISPs. Thus, ISPs today face the following dilemma: Are P2P file-sharing applications friends or foes? In particular, should ISPs curb P2P file-sharing traffic in their network? What strategies should the ISP take in order to keep P2P file-sharing traffic under control, while continuing to maximize the subscribers' experience? Such a dilemma has been transforming broadband business models worldwide.

Many providers that have chosen completely to block P2P traffic (thus eliminating the additional costs generated by such applications) are now starting to realize that this approach, in the long term, does not pay off, since it dramatically increases customer churn in today's competitive broadband market. Moreover, many users of P2P file-sharing applications are now able to circumvent standard traffic-blocking techniques deployed by ISPs. As a result, P2P-friendly solutions are currently being sought and some approaches have recently appeared in the literature. For example, ISPs can try to exploit traffic locality or content-caching of P2P applications in order to reduce the transit costs [4,1]. Nevertheless, it is still unclear whether an ISP can indeed make any profit by embracing P2P file-sharing applications in an operational environment, with a given customer base.

The goal of this work is to shed some light on this dilemma by providing an analytical study of the impact of different strategies that might be available to an ISP. In particular, we propose an analytical framework that captures the economical and performance aspects of the problem and models the tussle between the ISP and its potential customers. Our approach is fairly general and captures many parameters of this complex interaction, such as the quality of service expected by users, the dynamics of content popularity and replication, the effects of limited bandwidth-sharing and user impatience, and, of course, the influence of the specific choices adopted by the ISP to manage P2P traffic.

The proposed framework can be applied to specific operational conditions to analyze the best strategy that an ISP should adopt to manage P2P traffic generated by its customers. For example, the framework can be used to determine the required effectiveness of mechanisms that reduce transit costs by exploiting traffic locality (e.g. caching content) in order for P2P file-sharing traffic to be a profitable business for both the ISP and its customers.

Interestingly, our analysis reveals that costs incurred by an ISP exhibit a counter-intuitive, non-monotonic behavior as the number of subscribers grows. This occurs because increasing levels of P2P usage also augment the probability that a given object request can be served from another peer within the ISP boundaries. Moreover, for a given number of subscribers, there exists an optimal subscription fee that the ISP should charge so as to maximize its profit. Our framework also quantitatively shows that the current practice of providing asymmetric access bandwidths to the users (i.e. download bandwidth greater than upload bandwidth) can be very detrimental for an ISP.

An important aspect of the problem is the temporal dynamics of the content available to users. In particular, new content is constantly being introduced into the system, replicated by the downloads of users, and suffering changes in its popularity (i.e. number of queries it receives). While developing our framework, we also introduce a novel and general approach to account for this temporal dynamics. This module of our framework could also be applied for other purposes.

The paper is organized as follows. Section 2 summarizes related work on the topic including a discussion of the measurement studies that we have used to set the parameters of our model. Section 3 describes the network scenario we address. Section 4 introduces a simple model that allows the derivation of basic insights into the system dynamics. Section 5 presents a selection of the most interesting results we obtained by exploring the parameter space of our simple model. Section 6 describes the approach that we have used in our framework to capture resource popularity and replication. Section 7 describes a refined model that allows some important limitations of the simple model previously introduced to be removed. Section 8 presents a few results obtained by solving our refined model. Finally, Section 9 draws conclusions and outlines directions of future work.

2. Related work

Peer-to-peer (P2P) file-sharing applications are a disruptive technology in many different ways. In particular, traffic generated by these applications can have a significant impact on networks, mainly due to its high volume. In fact, recent studies [3] have pointed out that, in contrast to some rumors spread by popular media, the traffic volume due to P2P file-sharing is still increasing. One pressing issue is how to identify accurately and measure P2P traffic, as more and more applications become available to users which deploy obfuscation techniques in order not to be identified (e.g. random port numbers, packet encryption). Recent efforts to cope with this problem require sophisticated traffic identification techniques at the transport [2] or even at the application layer [1].

Besides large volume, traffic generated by P2P file sharing also exhibits other peculiar characteristics that can impact network performance. For example, traffic patterns of these applications differ substantially from Web traffic, being mainly driven by a "fetch-at-most-one" behavior coupled with the birth of new objects and users [5]. Moreover, application-level topologies generated by P2P applications tend poorly to match the underlying Internet topology [6], leading to ineffective use of the network resources. Finally, as with several other applications, it has been observed that content popularity in P2P file-sharing applications is highly skewed. In [7] it is shown that 20% of the files account for more than 80% of the downloads. A comprehensive study of the characteristics of the P2P file-sharing workload is presented in [5].

Given the potential impact that P2P file-sharing applications have on networks, ISPs have started to search for mechanisms to control and cope with this type of traffic. Several alternatives are discussed in [4], among which: (1) acquire more bandwidth; (2) block P2P file-sharing traffic; (3) utilize caching; (4) implement pricing schemes based on traffic volume; (5) shape P2P file-sharing traffic by means of service differentiation policies; (6) introduce application-layer redirection of P2P file-sharing traffic to exploit traffic locality. The consequences of adopting one or more of these alternatives have been partially investigated in the literature, as we discuss in the following. However, the feasibility and effectiveness of practical implementations of the proposed alternatives are still an open issue.

The effectiveness of caching content generated by P2P file sharing applications has been discussed in [7–9]. Caching reduces the transit costs for downstream bandwidth by serving requests within the ISP network. This solution looks particularly attractive due to substantial skewness of content popularity [5]. In [7] the authors conclude that the disk space required for effective caching of P2P traffic is small enough to be practical. In [9] the authors propose a caching solution which leverages existing web cache proxies. Despite its clear advantages, P2P-caching remains a controversial approach. The foremost concern for ISPs is the legality of such a solution. The provider would no longer be merely providing network infrastructure but would also be storing and forwarding potentially copyrighted content, which could push the ISP into copyright violation.

The gains achievable by exploiting traffic locality have also been investigated in the literature, mainly by simulation [5,10]. In [5] the authors present a set of small-scale simulation experiments based on a P2P workload model derived from traces. They quantify the potential bandwidth saving that locality-aware P2P file-sharing applications would achieve in various operational conditions. Similarly, in [10] the authors explore what-if scenarios of locality-aware solutions using real BitTorrent trace logs. Both studies above suggest that bandwidth savings of the order of 60% are achievable by exploiting traffic locality.

In yet a different way, P2P file-sharing traffic may also disrupt economic agreements between ISPs. Since traffic can be routed at the application layer, transit traffic agreements between ISPs can be violated. In [11], the authors propose an analytical framework to model the dynamics between two competing ISPs and understand their transit agreements. Their model is based on the economics and performance aspects of the problem, although the end users' point of view is not considered.

In our work we consider a very different problem; namely, the tussle between an ISP and its customers. In particular, we propose an analytical framework to highlight fundamental insights on the system behavior, as well as a basic understanding of the impact of the various parameters lying inside or outside the sphere of influence of the ISP. Our model considers both the economics and the performance aspects of the problem and takes into consideration both the ISP and its customers. To the best of our knowledge, we are the first to propose an analytical study of the tussle between ISPs and P2P users based on a comprehensive and extensible model.

3. Network scenario and assumptions

Consider the network scenario illustrated in Fig. 1. An ISP provides Internet access to a set of n customers who are interested in running P2P file-sharing applications to download content from other peers. The total number of peers in the P2P file-sharing community as a whole is denoted by N. The network of the ISP is connected to the rest of the Internet by a single link of downstream[1] bandwidth B_d and upstream bandwidth B_u. Users who subscribe to the ISP

[1] In this paper *downstream* refers to the direction from the Internet to the user, whereas *upstream* refers to the direction from the user to the Internet.

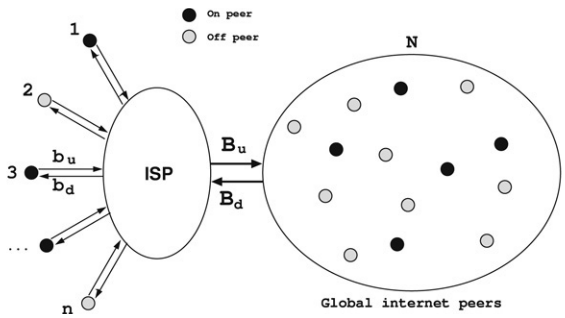

Fig. 1. Network scenario.

pay a fixed price c and get a downstream bandwidth b_d and an upstream bandwidth b_u, which are assumed to be the same for all of the ISP subscribers. We consider that traffic generated by P2P applications dominates the remaining traffic, and thus neglects the bandwidth consumption due to applications different from P2P. Alternatively, B_d could represent the bandwidth allocated to P2P traffic on a virtual link, in case the ISP is able to identify and segregate P2P traffic from the remaining traffic. The internal details of the ISP network are also neglected, as we assume that this part of the network is largely overprovisioned; hence, it does not affect the system performance. Moreover, we assume that the upstream link of capacity B_u connecting the ISP to the rest of the Internet does not constitute a bottleneck, and thus we neglect it in the remaining analysis.

The probability that a peer is on (i.e. running one or more P2P applications) is denoted by p_{on}. Let λ_q be the average demand of a peer for content, which can be interpreted as the mean rate at which a user generates queries using one or more P2P applications. We denote by σ the average probability that a query generated by a user results in a successful object retrieval. Notice that σ is not a system parameter, but a variable determined by the interaction between the ISP and its subscribers, as we soon describe.

We assume that the level of satisfaction of a user is just an increasing function of the successful object download probability σ. We chose this simple metric alone, instead of jointly considering the object transfer delay, because many P2P users are patient, and can tolerate long but reasonable download times; for example, because they keep the transfer active overnight. More specifically, we express the utility $U_i(\sigma, c)$ of user i as

$$U_i(\sigma, c) = \log(\alpha_i \sigma + 1) - c \tag{1}$$

where α_i is a shape parameter that can depend on the particular user. The log function has been chosen to model diminishing returns as σ increases, and in such a way that $\sigma = 0$ provides zero benefit to the user.

Users subscribe to the ISP only if they are minimally satisfied with the service provided. Thus, if $U_i(\sigma, c) \geq 0$ then user i subscribes to the ISP, otherwise the user prefers not to subscribe. Equivalently, user i subscribes to the ISP provided that $\sigma \geq \sigma_{\min_i}$ where $\sigma_{\min_i} = (\mathrm{e}^c - 1)/\alpha_i$.

The ISP receives revenue from its subscribers and must pay for the connection to a higher tier ISP. We assume the cost of the ISP is composed of a fixed component plus a cost proportional to the allocated bandwidth B_d. As mentioned above, subscribers only pay for subscription when they are minimally satisfied. Therefore, the ISP's utility function can be expressed as:

$$U_{\text{ISP}}(B_d, c) = \sum_{i=1}^{n} c\mathbb{I}_{(U_i(\sigma,c)\geq 0)} - (\beta_2 B_d + \beta_1) \tag{2}$$

where $\mathbb{I}_{(\cdot)}$ is the indicator function, β_1 is the fixed cost of providing bandwidth and β_2 is the cost for each unit of bandwidth. The objective of the ISP is to maximize its utility function over the control parameters B_d and c, for given values of β_2, β_1 and user population n.

4. Simple model

We now introduce a simple model to shed some light on the problem and acquire basic insights into the system dynamics. We consider, for simplicity, that all users are identical ($\alpha_i = \alpha$, $\forall i$). In this case, the ISP will convince all

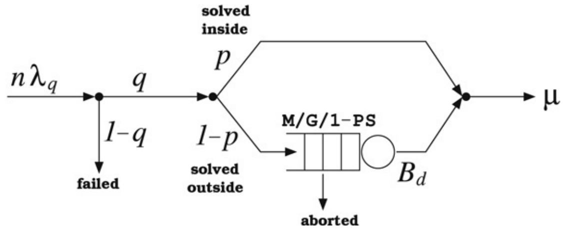

Fig. 2. Simple model for the outcome of queries generated by ISP subscribers.

of them to subscribe to the service, provided that $\sigma \geq \sigma_{\min} = (e^c - 1)/\alpha$, otherwise no user will decide to subscribe. Clearly, in order to minimize its costs, the ISP will allocate just the minimum amount of bandwidth B_d to guarantee that $\sigma \geq \sigma_{\min}$. Let us denote by $B_{\min}$ the minimum amount of bandwidth necessary to ensure that $\sigma = \sigma_{\min}$. We now describe a simple model that allows us to compute $B_{\min}$.

Fig. 2 illustrates the possible outcomes of queries generated by users subscribing to the ISP. Part of the queries generated by the users might fail to locate any copy of the desired content within the P2P community as a whole. This can happen because of several reasons, which are usually not under the control of the ISP. We thus treat the probability that a query is solved (i.e. at least one copy of the desired object is found) as a given parameter, and denote it by q. In general, a fraction of the queries that are solved can be served from within the ISP, while another fraction has to be served from peers residing outside the ISP. Let p be the probability that a solved query is served from within the ISP. To simplify the analysis further, we assume that internal downloads (i.e. data transfers between peers belonging to the ISP) always complete successfully. This corresponds to the best-case scenario of unlimited bandwidth within the ISP network. This assumption will be relaxed in Section 7.

If an object cannot be served from within the ISP, the downstream link of limited bandwidth B_d has to be used. Since this link is shared by all flows transferring objects from external peers, using elastic rate adaptation mechanisms (e.g. TCP), we adopt an M/G/1-Processor Sharing model to describe its dynamics. Insensitivity of this model with respect to the service time distribution permits to consider just the average object size, which is assumed to be known. We will actually normalize the object size to 1, and measure the bandwidth allocated on the link in objects per second.

We consider that users are impatient and tend to abort downloads that do not receive a minimum throughput. This is motivated by the fact that many P2P applications allow the user to inspect the current download rate of an object and abort the transfer. It is reasonable to assume that the abort decision is taken in the early phase of the transfer, relative to the total time required to download the object. Under these assumptions, the shared link is always stable and no bandwidth is wasted by partial downloads. Moreover, we will further assume that the shared link is always fully utilized (i.e. there is always at least one transfer in progress), yielding a throughput (expressed in objects per second) equal to B_d.

Let $\lambda = n\lambda_q$ be the aggregate generation rate of queries by the population of users belonging to the ISP. The aggregate system throughput μ, which is the rate at which objects are successfully retrieved by the users, can be expressed as

$$\mu = \lambda q p + B_d \tag{3}$$

where the first term of the sum corresponds to the throughput obtained by the queries that are served within the ISP and the second term corresponds to the throughput obtained by the queries served from outside the ISP. It follows that the probability that a query generated by a user results in a successful download is $\sigma = \mu/\lambda$, which is given by

$$\sigma = qp + \frac{B_d}{\lambda}. \tag{4}$$

Notice that, if $qp \geq \sigma_{\min}$, the ISP can allocate a bandwidth $B_d = 0$. This means that the system sustains itself even in isolation, i.e. without any connection to the rest of the P2P file-sharing community. Instead, if $qp < \sigma_{\min}$, the ISP has to allocate a minimum bandwidth $B_{\min} = \lambda(\sigma_{\min} - qp)$ so as to provide the minimum service for which users are willing to subscribe to the ISP. The ISP will actually do so (allocate $B_{\min}$) provided that this strategy leads to positive utility according to Eq. (2), i.e. $nc - \beta_1 - \beta_2 B_{\min} > 0$, otherwise the ISP will not allow its users to run P2P file-sharing applications.

Finally, we introduce a simple model to determine the probability p that a query generated by a user is served by some other peer belonging to the ISP. In general a solved query will find one or more replicas of an object, part of them stored inside the ISP and part of them stored outside the ISP. The probability that at least one replica of a desired object resides within the ISP can be expressed as $1-(1-f/F)^r$, where f denotes the total number of files stored by users belonging to the ISP, F denotes the total number of files stored by the P2P community, and r denotes the number of replicas of the object that are present in the P2P system at the time the object is requested.

Under the assumption that $f \ll F$, which naturally stems from the assumption that $n \ll N$, the probability that at least one replica of a requested object resides within the ISP can be approximated by rf/F. Let γ be the probability that an object is indeed downloaded from an internal peer, given that at least one replica exists within the ISP. This probability represents the efficacy of the system in exploiting traffic locality. We do not model the details of how γ depends on the optimizations already incorporated in the P2P applications (e.g. preference mechanisms based on RTT measurements), or how γ can be influenced by ISP choices (e.g. query filtering/redirection). We observe that $\gamma = 1$ represents a best-case scenario to save bandwidth B_d (although we will later show that $\gamma = 1$ is not always optimal).

We further assume that a copy of an object is cancelled/removed by a user after an average time τ. In addition, users introduce new objects in the shared folder of their computers at rate λ_o. Quantities such as τ, λ_o, and λ_q are assumed to be equal for users residing inside or outside the ISP. Instead, we consider the possibility that users belonging to the ISP are treated differently from the rest of the users belonging to the P2P community. Indeed, different ISPs around the world can offer different types of contract to their customers, who in turn expect different service levels (in terms of object throughput). We take this fact into account introducing the average probability $\hat{\sigma}$ that objects are successfully retrieved by a generic user belonging to the P2P community. The value of $\hat{\sigma}$ can, in general, be different from the value of σ resulting from the interaction of the given ISP with its customers.

Applying Little's law, we have $f = n(\lambda_o + \lambda_q\sigma)\tau$, $F = N(\lambda_o + \lambda_q\hat{\sigma})\tau$, and we obtain the approximation (valid for $n \ll N$),

$$p \approx \gamma r \frac{n}{N}\frac{\lambda_o + \lambda_q\sigma}{\lambda_o + \lambda_q\hat{\sigma}}. \tag{5}$$

Putting things together, we obtain that the minimum amount of bandwidth that the ISP has to allocate (provided that this leads to positive utility for the ISP), can be expressed as

$$B_{\min} = \max\left[0, n\lambda_q\left(\sigma_{\min} - q\gamma r\frac{n}{N}\frac{\lambda_o + \lambda_q\sigma_{\min}}{\lambda_o + \lambda_q\hat{\sigma}}\right)\right]. \tag{6}$$

We observe that $B_{\min} = 0$ (which means that the P2P community within the ISP network sustains itself even in isolation) can be a solution for sufficiently large n (but, recall, we assume $n \ll N$). In the special case, in which all users of the P2P community are treated the same ($\sigma_{\min} = \hat{\sigma}$), we obtain the simplified expression

$$B_{\min} = \max\left[0, n\lambda_q\left(\sigma_{\min} - q\gamma r\frac{n}{N}\right)\right] \tag{7}$$

which does not depend on λ_o. The native parameters of our model are summarized in Table 1.

5. Basic insights obtained with the simple model

In this section we present a selection of the most interesting results that we have obtained while exploring the parameter space of our simple model. Some of the parameters have been chosen according to measurement studies of P2P file-sharing traffic recently appearing in the literature. The values that will be kept fixed throughout the experiments of this section, unless otherwise specified, are the following:

- $N = 5 \cdot 10^7$. This number has been chosen because the number of online users of popular P2P applications is typically around a few million, while the percentage of online time is quite small (less than 10%) [12], hence the size of the P2P file-sharing community is estimated in the order of tens to hundreds of millions of users.
- $\lambda_o = \frac{((20*5475)/365)}{N}$ objects/day. The average number of new video objects (movies) introduced by users in one year is 5475, according to [5]. Although video objects account for about 65% of overall P2P traffic, they represent only about 5% of all file requests [5]. Using the same introduction rate for all types of file (a factor of 20) and

Table 1
Native parameters

N	Number of peers in the global P2P community
n	Number of potential subscribers of the ISP
p_{on}	Probability that a peer is online
λ_o	Average introduction rate of new objects by a peer
λ_q	Average object request rate by a peer
α	Shape factor of users' utility function
c	Subscription cost per user
B_d	Downstream capacity of ISP access link to the Internet
β_1	Fixed costs for the ISP
β_2	Cost per unit of external bandwidth contracted by ISP
q	Probability that a query is solved by the P2P application
r	Mean number of replicas available for a requested object
γ	Probability to download an internal copy (if any)
$\hat{\sigma}$	Successful object retrieval probability for all peers
τ	Average time before an object is removed by a user

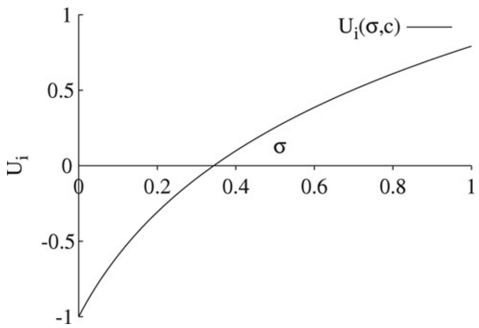

Fig. 3. The utility function U_i for $\alpha = 5$ and $c = 1$.

dividing by the number of days in one year (365), we obtain the above expression for the mean rate at which new objects are introduced per day by one user.

- $\lambda_q = 1$. This estimate of the mean object request rate by a user (per day) has been taken from the workload model proposed in [5].

At last, we set the parameters of the user utility function (1) as follows. We normalize the subscription cost c of a user to 1, and we choose the shape parameter $\alpha = 5$ so that the minimum value $\sigma_{\min} = (e^c - 1)/\alpha = 0.34$. This means that we consider a population of users who are minimally satisfied when they are able successfully to retrieve about one third of requested objects. The actual user utility function is depicted in Fig. 3.

5.1. Analyzing $B_{\min}$

We start considering how $B_{\min}$, the minimum amount of external bandwidth that the ISP must allocate to convince the users to subscribe, depends on various system parameters. Looking at Eq. (6), we notice that $B_{\min}$ is given by a second-order polynomial of the population size n. We observe that equation $B_{\min} = 0$ has always one root corresponding to $n = 0$. The other root, denoted by $n_{\max}$, represents the critical value of n above which no external bandwidth is needed[2]: this happens when the local P2P community is large enough to sustain itself. It follows that $B_{\min}$ exhibits a counter-intuitive, non monotonic behavior as a function of n (see plots in Fig. 4). In the following we explore the effect of several parameters.

The impact of r. In this set of experiments we take a scenario in which $\gamma = 1$, $q = 0.9$, $\widehat{\sigma} = \sigma_{\min}$, and we consider three different values of the average number of replicas r, 500, 1000, and 1500. Fig. 4 (left upper plot) depicts the

[2] Notice that when $B_{\min} = 0$ new contents are still introduced into the system by internal users.

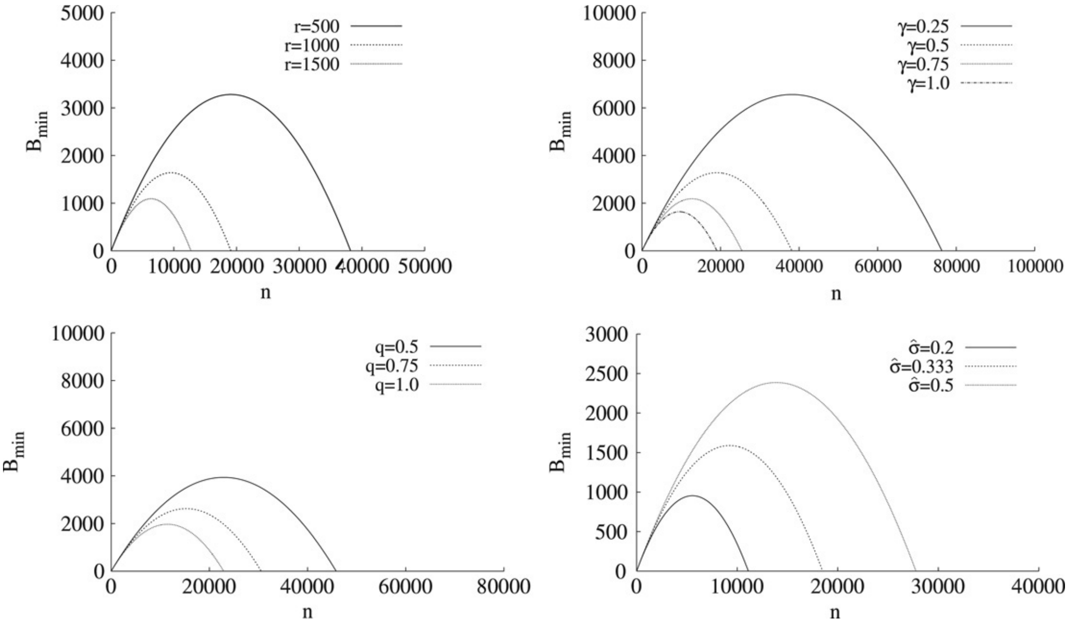

Fig. 4. $B_{\min}$ as a function n, for different values of r (upper left plot); $B_{\min}$ as a function of n, for different values of γ (upper right plot); $B_{\min}$ as a function of n, for different values of q (lower left plot); $B_{\min}$ as a function of n, for different values of $\widehat{\sigma}$ (lower right plot).

behavior of $B_{\min}$ as a function of n for the three considered values of r. Note that the larger the r, the smaller the value of $n_{\max}$ at which the ISP community could self-sustain.

The impact of γ. In the previous set of experiments we have considered the best-case scenario in which the efficacy of the system in selecting an internal copy of a desired object (provided that it exists) is maximum, i.e. $\gamma = 1$. Of course this is possible only in an ideal system. In practice, γ depends on native optimizing mechanisms implemented by P2P applications, but it can also be influenced by the ISP through application-layer redirection of queries, which, according to [4], is the preferred strategy for an ISP.

To evaluate the impact of γ, we fix $r = 1000$, $q = 0.9$, $\widehat{\sigma} = \sigma_{\min}$ and we let γ vary. Fig. 4 (upper-right plot) depicts the behavior of $B_{\min}$ as a function of n for four different values of γ, namely, 0.25, 0.5, 0.75, and 1.

We observe that, if the redirection mechanism to internal peers is not perfect ($\gamma < 1$), the value of $n_{\max}$ can increase considerably with respect to the ideal case ($\gamma = 1$), possibly making the solution $B_{\min} = 0$ impractical. In this case, for a given population size $n < n_{\max}$, curves such as those in Fig. 4 (upper-right plot) provide the additional external bandwidth that is needed for different values of γ.

The role played by γ is quite interesting because it could also be used to model the hit ratio of a content-caching solution.

The impact of q. Another important parameter that affects the system performance is q, the probability that a query is solved in the first place, i.e. the P2P application is able to locate at least one copy of the requested object at another peer of the P2P community. The value of q essentially depends on the effectiveness of the search mechanisms implemented by the P2P architecture, so it is only marginally modifiable by the ISP. In principle, an ISP could artificially drop queries (for example, in a probabilistic manner) to a priori limit the use of the external bandwidth. However, according to our model, this is not a sensible strategy. To prove this, we consider a scenario in which $\gamma = 0.75$, $\widehat{\sigma} = \sigma_{\min}$, $r = 1000$, and we vary q: 0.5, 0.75, and 1. Fig. 4 (lower-left plot) shows that, by reducing q, an ISP can only increase the amount of external bandwidth needed to fulfill the expectations of its customers. Conversely, if the ISP is able to increase the hit probability of queries, i.e. increase the value of q, then the value of $n_{\max}$ decreases because the need for the external bandwidth is lower.

The impact of $\widehat{\sigma}$. The quality of service enjoyed by P2P users throughout the world can be highly diverse, because users belong to ISPs that can offer very different contracts (e.g. in terms of subscription cost), for example because they operate in more or less competitive markets, with consequently different expectations by the customers. Therefore, the average capacity of an arbitrary peer of the community to using P2P applications (modelled by $\widehat{\sigma}$) can deviate substantially from the service level obtained by peers within the considered ISP.

To investigate the impact of $\widehat{\sigma}$, we consider a scenario in which $q = 0.9$, $\gamma = 1$, $r = 1000$. Recall that, in our basic parameters setting, the minimum level of satisfaction required by the ISP customers is $\sigma_{\min} \sim 0.34$. In Fig. 4 (lower-right plot) we have plotted the behavior of $B_{\min}$ as a function of n for three different values of $\widehat{\sigma}$, i.e. 0.2, 0.333, and 0.5. These values correspond to an average service level outside the reference ISP that is lower, similar, or higher, respectively, than the service level expected by peers belonging to the ISP. We observe that, if users outside the ISP enjoy (on average) a better (worse) service, the amount of external bandwidth that the ISP has to allocate increases (decreases). This is essentially due to the fact that the number of copies stored by users is proportional to their service level. The larger (smaller) the number of copies that exist within the ISP, the higher (lower) the probability that a query finds an object internally, thus avoiding the consumption of external bandwidth.

Curves such as those depicted in Fig. 4 can help an ISP to select a strategy to manage P2P traffic. If the population size is above $n_{\max}$, no external bandwidth is needed to support P2P traffic, provided that users are indeed able to locate copies of the desired objects within the ISP. This can be achieved by redirecting all queries towards internal peers, or by providing to the users a customized P2P application running within the ISP network boundaries. The feasibility of this solution, however, strongly depends on the degree of replication r of contents, confirming that this parameter is indeed crucial for the system performance. This actually motivates the in-depth analysis of content replication that will be presented in Section 6.

5.2. *Analysing* U_{ISP}

To determine whether it is really profitable for an ISP to start a service with given characteristics, we need to consider the costs associated with the required bandwidth $B_{\min}$, i.e. evaluate the utility function Eq. (2) for $B = B_{\min}$ and check if it is positive. Actually, it may well be that the total income obtained from the subscription fee of the customers does not cover the expenses due to fixed costs (β_1) and to the bandwidth contracted by the higher-tier ISP (proportional to β_2). It turns out that there exists a minimum value of n, denoted by $n_{\min}$, below which the ISP has no incentive to start the service, because by so doing, its utility would be negative. The value of $n_{\min}$ corresponds to the positive root of the second-order equation $U_{\mathrm{ISP}}(n) = 0$ in which $B = B_{\min}$.

The impact of c. From the ISP viewpoint it is especially interesting to evaluate how U_{ISP} (and threshold $n_{\min}$) depends on the subscription cost c. Indeed, on one side, larger values of c produce proportionally bigger incomes from the subscribers (equal to $n\,c$); on the other side, for larger values of c, users demand better and better service levels. So how to choose the value of c? To address this issue we consider a scenario in which $\gamma = 1.0$, $r = 1000$, $q = 1.0$, $\widehat{\sigma} = 0.6$, $\beta_1 = 500$ and $\beta_2 = 5$. We consider four different values of c, 0.6, 1.0, 1.4, 1.6, with corresponding minimum service levels $\sigma_{\min} = \frac{e^c - 1}{\alpha}$.

Fig. 5 (upper-left plot) depicts U_{ISP} as a function of n for the four considered values of c. We observe that the optimal value of c depends on the number of users. For large number of customers ($n > 25\,000$) the higher profit is obtained from the largest value of c, i.e. $c = 1.6$. However, the situation changes as the number of potential users decreases. Indeed, setting a smaller c eases the starting of the service ($n_{\min}$ reduces), and can provide positive utilities for values of n where larger c's are not feasible. Using the model, the optimum value of c for a given number of users can be found numerically, and the resulting values are shown in Fig. 5 (upper-right plot), together with the corresponding values of U_{ISP}. Notice that c cannot exceed the value $c_{\max} = \log(\alpha q + 1) \sim 1.8$, because it cannot provide a successful retrieval probability greater than the probability q that a query is solved by the P2P application. Moreover, even by tuning c, there exists a minimum value of n below which the ISP does not gain any profit by starting the service.

The impact of q. We consider a scenario in which $\gamma = 1.0$, $r = 1000$, $\widehat{\sigma} = \sigma_{\min}$, $\beta_1 = 1000$ and $\beta_2 = 10$, and vary the value of $q = 0.35, 0.5, 0.75, 1.0$.

Fig. 5 (lower-left plot) depicts U_{ISP} as a function of n for the considered values of q. We observe that the optimal utility for the ISP is always achieved for the largest q. Also, the threshold $n_{\min}$ is very sensitive to q.

The impact of r *and* β_2. At last, we investigate the dependency of $n_{\min}$ on the average number of replicas r, while also varying the cost β_2 per bandwidth unit incurred by the ISP. We considered a scenario in which $c = 1$, $\alpha = 5$, $\gamma = 1.0$, $q = 1.0$, $\widehat{\sigma} = 0.6$, $\beta_1 = 500$.

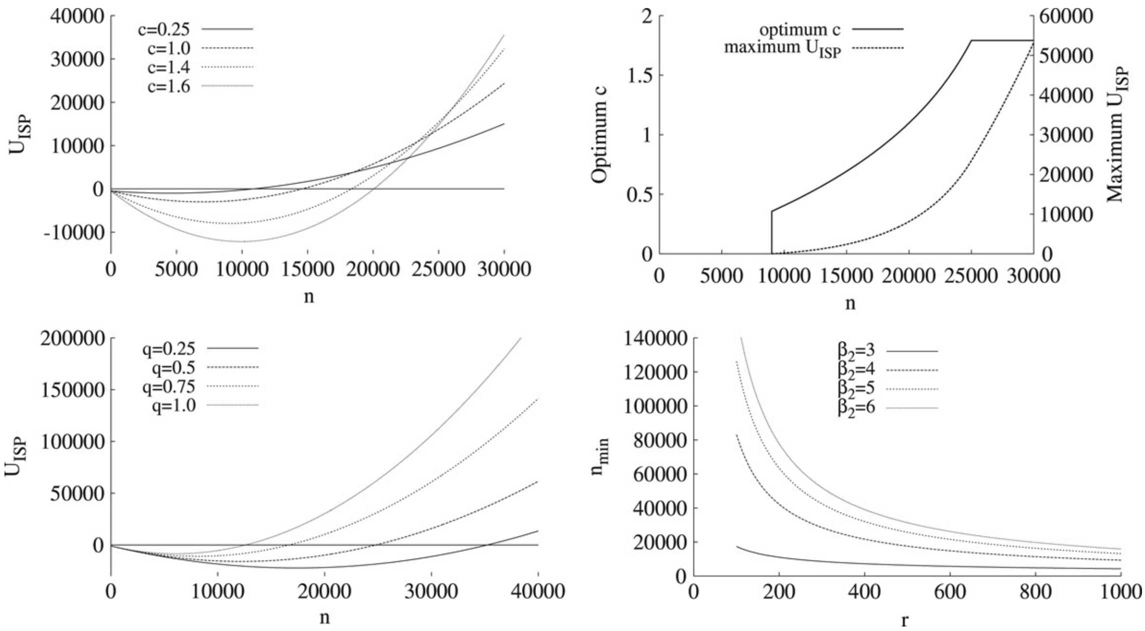

Fig. 5. U_{ISP} as a function of n, for different values of c (upper-left plot); optimum c as function of n, and corresponding maximum U_{ISP} (upper-right plot); U_{ISP} as a function of n, for different values of q (lower-left plot); $n_{\min}$ as a function of r, for different values of β_2 (lower-right plot).

Fig. 5 (lower-right plot) depicts the threshold $n_{\min}$ as a function of r for $\beta_2 = 3, 4, 5, 6$. We observe again the crucial role played by r in making the service lucrative to the ISP for a given customer base. The impact of β_2 on $n_{\min}$ is especially significant when r is small. This can be explained by recalling that the higher the value of r, the lower the need for an external bandwidth, hence β_2 is less important than fixed costs represented by β_1.

6. Modeling object popularity and replication

As we have seen, one fundamental parameter that affects the system performance is r, which has been introduced in the simple model to compute the probability that at least one replica of a requested content exists within the ISP network. Parameter r represents the average number of replicas of an arbitrary object that are available in the global P2P system *at the time the object is requested*. This number depends on both the level of popularity and the extent of replication of objects within the P2P community. Notice that these two metrics, for a given object, evolve over time, and their instantaneous values are not necessarily correlated: when a content reaches its highest popularity, there might be just a few copies of it in the system; similarly, there can be a large number of replicas of an object at the time the object has lost much of its popularity. The analysis is made complex by the fact that objects are heterogeneous and highly dynamic: new contents are constantly added into the system, and their popularity can change over time in very different ways. Some contents lose popularity very soon after they are introduced (e.g. videos related to sport events), whereas others can keep their popularity almost unchanged for many years (e.g. popular songs).

The usual approach taken in the literature to model the popularity of objects is to consider their rank distribution, which is typically assumed to be Zipf. New objects are then inserted at some popularity rank index according to the same distribution; pre-existing objects of equal or lesser popularity are pushed down, and the resulting distribution is re-normalized to keep the total probability equal to 1 [5]. Unfortunately, this model does not permit the temporal evolution of contents popularity and replication to be computed in an easy way.

Since existing techniques do not allow r to be evaluated, we have developed a new approach, based on native system parameters, that is able to compute the entire distribution of the number of replicas of an object that are present in the system at the time the object is requested. We believe the proposed methodology is an independent contribution of this paper.

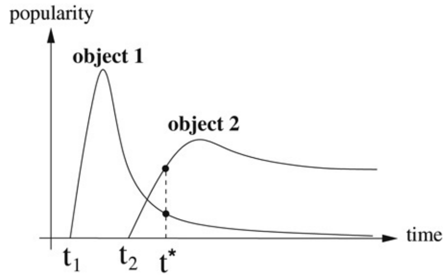

Fig. 6. Example of evolution of popularity for two different objects.

We consider that the popularity of an object can be quantified by some metric that evolves according to a generic function of time, independently for each object. The absolute value of this metric is not important, since what matters is just the relative popularity value of an object with respect to other objects present in the system at a given point in time.

We assume that new objects are inserted into the system according to a Poisson process of rate $N\lambda_o$. Upon arrival of object i, a popularity curve $p_i(t)$ belonging to a common family of curves $\mathcal{P}$ is instantiated and associated with the new file. The set $\mathcal{P}$ comprises, in general, all continuous, bounded functions defined over $[0, \infty]$. We assume that the probability that an object is requested at a given point in time is equal to its current popularity value, divided by the sum of the popularity of all available objects. As we will see, the number of objects in the whole P2P network is large enough for us to assume that the sum of the popularity of all objects is constant. It follows that the number of copies of an object, inserted at time t_i, that have been downloaded up to time t is proportional to $\int_{t_i}^{t} p_i(\tau - t_i)\mathrm{d}\tau$.

Since we are interested not in the number of times an object has been downloaded, but in the number of copies still available in the system, we need to account for the fact that users tend to cancel files from the hard disk of their computers, for instance to make room to new objects. We do so by allotting to each new object i a cancelation rate δ_i (possibly dependent on its particular popularity curve p_i) and assuming that copies are removed by users after an exponentially distributed amount of time of mean $\tau_i = 1/\delta_i$. We obtain the number of available copies $A_i(t)$ of object i at time t as being by

$$A_i(t) = G \int_{t_i}^{t} p_i(\tau - t_i)\mathrm{e}^{-\delta_i(t-\tau)}\mathrm{d}\tau \tag{8}$$

where G is a constant, and $A_i(t) = 0$ for $t < t_i$.

As an example, Fig. 6 depicts two objects whose popularity evolves in very different ways over time: object 1, inserted at time t_1 soon reaches a high value, but its popularity decays quickly, for instance because it refers to an event of the news that is no longer interesting a few days after its appearance. Object 2, inserted at time t_2, gains popularity more slowly, but maintains its popularity for a much longer time; it could represent, for example, a popular song. Now suppose that a peer issues a request at time t^*. At this time, object 2 has great popularity, thus it will be more likely to be requested than object 1. However, by this time, object 1 will have been downloaded many more times (in a number proportional to the integral of its popularity curve up to time t^*); therefore, it will be more likely to be found (assuming that the cancelation rate is the same for both objects).

Our model can be regarded as a *Poisson shot-noise process* [13], where a *shot* represents the evolution over time of the number of available copies of a given object inserted into the system, i.e. $A_i(t)$. We can thus exploit well-known results on such stochastic processes to study the dynamics of objects in the system. In particular, we can characterize the distribution of the total number $Z(t)$ of copies of all objects present in the system at a given point in time:

$$Z(t) = \sum_{i \in \mathbb{Z}} A_i(t).$$

Let $S_i = \int_{t_i}^{\infty} A_i(t)\mathrm{d}t$. For simplicity, we assume S_i to have finite expectations, so that process $Z(t)$ is stationary.[3] Under this assumption we have $\mathbb{E}[Z] = N\lambda_o\mathbb{E}[S]$. Since the number of different objects (and copies) available in the

[3] In reality $Z(t)$ is likely to be non-stationary, as we can argue that the total number of objects (and copies) shared by the P2P community increases over time. We neglect such long time-scale dynamics in this paper. The impact of the non-stationarity of the process is subject to future work.

system is extremely large, the coefficient of variation of $Z(t)$ is very small [13], and we can readily assume that $Z(t)$ is constant, i.e. $Z(t) = Z$. The constant G to be used in (8) is then given by:

$$G = \frac{Z}{N\lambda_o}\left(\mathbb{E}\left[\int_0^\infty \int_0^t p_i(\tau)\mathrm{e}^{-\delta_i(t-\tau)}\mathrm{d}\tau\,\mathrm{d}t\right]\right)^{-1}.$$

Now, let $f(\boldsymbol{\omega}, \delta)$ be the joint pdf of the parameters $\boldsymbol{\omega}$ that specify a generic popularity curve $p_\omega(t)$ and the associated cancelation rate δ.

Let $A_{\omega,\delta}(t) = G\int_0^t p_\omega(\tau)\mathrm{e}^{-\delta(t-\tau)}\mathrm{d}\tau$. The probability $\boldsymbol{P}(A > a)$ that a query is directed to an object available in a number of copies greater than a can be computed as

$$\boldsymbol{P}(A > a) = \frac{1}{K}\int_{\boldsymbol{\omega},\delta} f(\boldsymbol{\omega}, \delta)\int_{t:A_{\omega,\delta}(t)>a} p_\omega(t)\mathrm{d}t\,\mathrm{d}\boldsymbol{\omega}\mathrm{d}\delta. \tag{9}$$

where K is a normalization constant

$$K = \int_{\boldsymbol{\omega},\delta} f(\boldsymbol{\omega}, \delta)\int_0^\infty p_\omega(t)\mathrm{d}t\,\mathrm{d}\boldsymbol{\omega}\mathrm{d}\delta.$$

From (9) one can compute the pdf $f_A(a)$ of the number of available copies of a requested object. We remark that our model is intended to capture the macroscopic dynamics of objects available in the P2P community, and provides great flexibility in coping with the heterogeneous nature of contents, by considering different families of popularity curves and cancelation rates.

7. Model refinements

The simple model described in Section 4 could be refined in many different ways. Here, we remove one fundamental limitation of the simple model, that is, the assumption that the finite bandwidth on the access links of the users does not affect the system performance. This is, indeed, a crucial point, since the exploitation of traffic locality is indeed effective only when the upload bandwidth available on the access link of internal peers can be used in place of the external bandwidth B_d.

To account for the finite downstream bandwidth b_d of the users (see Fig. 1), we adopt an M/G/1-Processor Sharing model with rate limit b_d to describe the sharing of the external bandwidth B_d. This approach has also been used in [14] to model elastic flows through a bottleneck link. To maintain the queue stable under any load condition, we explicitly model user impatience as follows: when an object starts to be downloaded, we assume that the user inspects the rate achieved by the flow at the very beginning of the data transfer (let this rate be equal to b), and decides to keep it going with a probability p_g that depends on the ratio between the achieved rate and the access bandwidth b_d. With probability $1 - p_g$ the transfer is instead prematurely aborted. Probability p_g can be modelled as an arbitrary function of the ratio b/b_d. We have chosen the following dependency of p_g as a function of b/b_d:

$$p_g(b) = \left(\frac{b}{b_d}\right)^\eta$$

where $\eta \geq 0$ is a tunable parameter accounting for different degrees of user impatience. Fig. 7 depicts how probability p_g depends on η according to the chosen formula. The case $\eta = 0$ corresponds to the case of infinite patience ($p_g = 1$). As $\eta \to \infty$, the user becomes more and more intolerant, claiming always to obtain a rate close to the bandwidth he/she has paid for.

This model of user impatience can be easily incorporated into our bandwidth-sharing model by introducing arrival rates of flows that depend on the number of flows in progress. Using the same technique described in [15], it can be shown that the resulting model is still insensitive to the object size distribution. Let $\lambda_e = n\lambda_q q(1 - p)$ be the arrival rate of flows at the processor-sharing queue representing link B_d (before drops). The probability $\pi(i)$ that i flows are in progress at an arbitrary time instant can be computed as

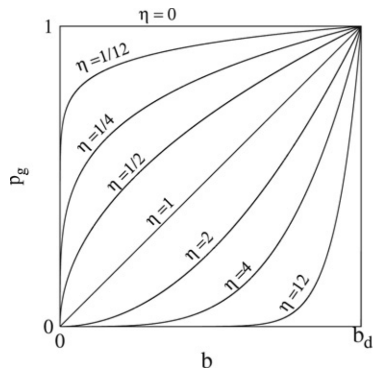

Fig. 7. Probability p_g that a user keeps doing a data transfer as a function of the initial achieved rate b.

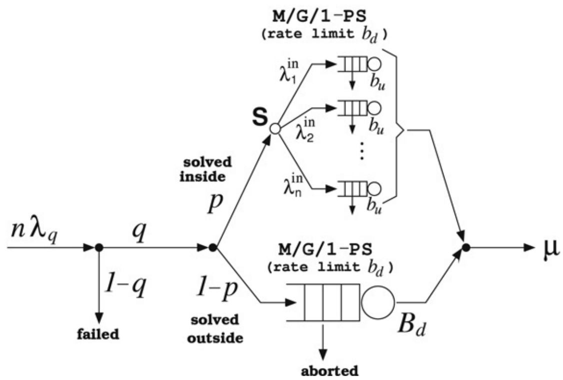

Fig. 8. Refined model for the outcome of queries generated by ISP subscribers.

$$\pi(i) = \begin{cases} \pi(0)\left(\dfrac{\lambda_e}{b_d}\right)^i \dfrac{1}{i!} & \text{if } i \le m, \\ \pi(0)\left(\dfrac{h^m}{m!}\right)^{1-\eta}\left(\dfrac{\lambda_e h^\eta}{B_d}\right)^i \dfrac{1}{i!^\eta} & \text{if } i > m \end{cases} \tag{10}$$

where $h = B_d/b_d$, and m denotes the integer part of h, i.e., $m = \lfloor h \rfloor$; probability $\pi(0)$ is obtained imposing that $\sum_{i=0}^{\infty} \pi(i) = 1$. We observe that the queue is always stable, provided that $\eta > 0$. In the case of $\eta = 0$, the queue is stable if $\lambda_e/B_d < 1$. Finally, the rate at which flows using the external bandwidth B_d end successfully, denoted by μ_e, is given by

$$\mu_e = [1 - \pi(0)]B_d - \sum_{i=1}^{m}[\pi(i)(B_d - ib_d)]. \tag{11}$$

To account for the finite upstream bandwidth b_u of the users, we model each internal peer j as an M/G/1-processor-sharing queue with service capacity b_u, rate limit b_d, and flow arrival rate λ_j^{in}, $1 \le j \le n$. A schematic representation of our refined model is shown in Fig. 8. Notice that the aggregate internal arrival rate of flows $\lambda^{\text{in}} = \sum_{j=1}^{n} \lambda_j^{\text{in}}$ is equal to $n\lambda_q qp$. For each internal queue we can apply the same formulas derived for the bandwidth-sharing model of link B_d, having substituted b_u for B_d, and λ_j^{in} for λ_e. Each peer provides a throughput μ_j^{in} computed by a formula analogous to (11). The overall system throughput is then computed as $\mu = \mu_e + \sum_{j=1}^{n} \mu_j^{\text{in}}$.

The only remaining difficulty is how to distribute the internal load among the peers, as represented in Fig. 8 by the splitting node S. A simple approach would be to consider all nodes identical. In this case we would have $\lambda_j^{\text{in}} = \lambda^{\text{in}}/n$, $\forall j$. However, peers are very heterogeneous, and, in particular, the fraction of time during which a peer is online, as well as the number of files made available by a peer, is expected to be highly diverse within the user population.

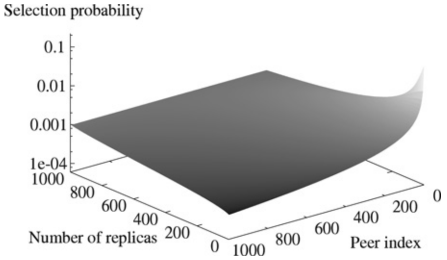

Fig. 9. Peer selection probability as a function of number of replicas, for $\nu = 0.8$.

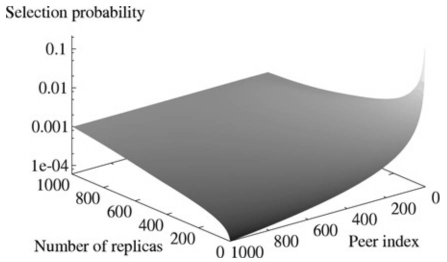

Fig. 10. Peer selection probability as a function of number of replicas, for $\nu = 1.2$.

This fact produces an uneven load distribution among the peers that can saturate the upload bandwidth on some of the most available peers, thus giving rise to internal bottlenecks that can disrupt the exploitation of traffic locality. Due to the relevance of this issue, we propose here an analytical approach to compute the load distribution among the peers, which makes use of the distribution of replicas for a requested object computed in Section 6, thus justifying and strengthening the proposed technique to modeling object popularity and replication.

We assume that the availability of resources at peers is described by a generic weight function $w(j)$; we consider the normalized weight function $w_n(j) = \frac{w(j)}{\sum_{k=1}^{n} w(k)}$. The quantity $w_n(j)$ represents the probability that peer j is selected as an uploader in the case in which there exists exactly one replica of the requested content within the ISP network. The problem is to evaluate the selection probability of each peer in case there are multiple replicas of a content. Since each peer is supposed to have only one copy of an object, our analysis maps onto the known problem of computing the inclusion probability of a finite set of items in a sampling procedure without replacement and varying draw probabilities [16]. This problem is, in general, very difficult to solve exactly. Therefore we have followed the approach proposed in [16], that provides an asymptotic analysis for large population sizes. We have verified that existing asymptotic formulas provide a very good approximation for the typical values of n that we consider, by comparing results against detailed simulations of successive sampling.

Following [16] we define the implicit function $t(y)$ by the relation

$$n - y = \sum_{s=1}^{n} \mathrm{e}^{-w_n(s)t(y)}, \quad 0 \le y < n. \tag{12}$$

The probability that peer j has a replica, given that the object is replicated a_i times, is given by

$$u(j) = \frac{1 - \mathrm{e}^{-w_n(j)t(a_i)}}{a_i}. \tag{13}$$

The peer selection probability is then obtained by normalizing probabilities $u(j)$ to sum up to one.

As an example, consider the case of 1000 peers, characterized by a weight function $w(j)$ which behaves as a power-law, i.e. $w(j) \propto j^{-\nu}$. This is a realistic weight function, because it has been observed that availability of resources at peers is highly skewed, with a small fraction of peers providing a significant percentage of the total resources available in the P2P community. Figs. 9 and 10 show the probability that each of the 1000 peers is selected as an uploader, as a function of the number of replicas of a content available within the ISP network, for $\nu = 0.8$

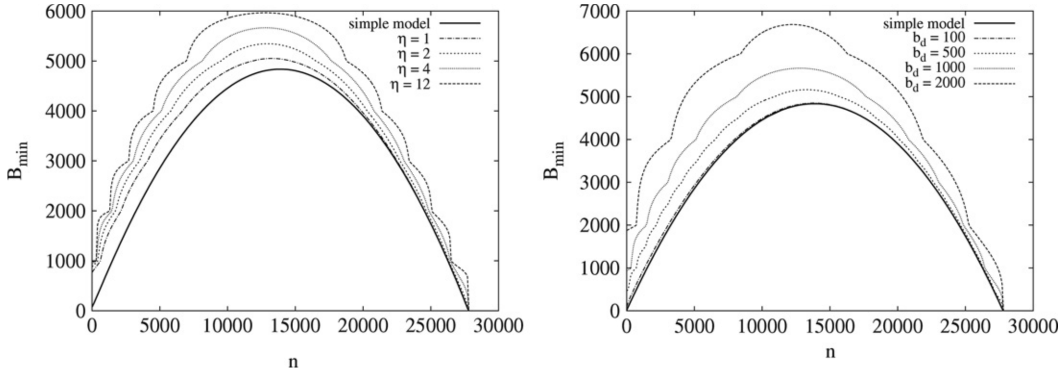

Fig. 11. $B_{\min}$ as a function n, for different values of η, $b_d = 1000$ (left plot); $B_{\min}$ as a function of n, for different values of b_d, $\eta = 4$ (right plot).

and $\nu = 1.2$, respectively. Notice that, as the number of replicas increases, the distribution approaches a uniform distribution (i.e. $p(j) = 1/1000, \forall j$); however, this happens more slowly when ν is larger.

To obtain the final load splitting distribution in the ISP network, we need to make an average of the peer selection probabilities weighted by the conditional probabilities $P(a_i)$ that there are a_i replicas available in the ISP network, given that there exists at least one replica (since we already know that the query has been solved internally). To compute $P(a_i)$, we start from the pdf $f_A(a)$ of the number of copies of a requested object that are present in the whole P2P community, as derived in Section 6. The probability that a copy in available at an online internal peer is $P_{oi} = p_{\text{on}} \frac{n}{N} \frac{\lambda_o + \lambda_q \sigma}{\lambda_o + \lambda_q \hat{\sigma}}$. The probability that there are m such replicas is given by $Q(m) = \sum_{i \geq m} f_A(i) \binom{i}{m} P_{oi}^m (1 - P_{oi})^{i-m}$. Finally, $P(a_i)$ is obtained as $Q(a_i)/\sum_{j \geq 1} Q(j)$.

8. Results with finite access bandwidth

In this section we present a few results obtained by solving the refined model that accounts for finite user access bandwidth, as described in Section 7. We will first isolate the impact of finite download bandwidth b_d, keeping the assumption that the upload bandwidth b_u of each user is infinite. Then, we will consider the more complex case in which both b_d and b_u are finite.

The impact of b_d. We consider a scenario in which $\gamma = 1$, $q = 0.9$, $\sigma_{\min} = \widehat{\sigma} = 0.5$. The impact of b_d can be very different, depending on the degree of user impatience, which has been modelled by parameter η (see Fig. 7). To explore the joint impact of b_d and η, we conduct two different experiments. In the first one we fix $b_d = 1000$, and let η vary. The results of these experiments are shown in the left part of Fig. 11, in which we plot $B_{\min}$ as function of n. In the second one we fix $\eta = 4$, and let b_d vary. The results of these experiments are shown in the right part of Fig. 11. Notice that, although we vary b_d, the content demand generated by the users is kept constant. When either $b_d \ll B_d$, or $\eta \sim 0$, the results are almost indistinguishable from those obtained by the simple model introduced in Section 4. The reason is that in these conditions the shared link is fully utilized almost all the time; thus, we do not make any significant error by assuming that it contributes a constant throughput equal to B_d. The behavior deviates from that predicted by the simple model when b_d is comparable to B_d, and, at the same time, η is not too small. In this case, the number of active flows is small enough for the link to fail to be fully utilized with non-negligible probability. For large values of η (e.g. 12), note that the required values of $B_{\min}$ concentrate at multiples of b_d. This can be explained by the fact that, when users are extremely impatient, they tend to abort a transfer whenever they do not obtain a rate equal to b_d. As a consequence, values of $B_{\min}$ in between two consecutive multiples of b_d are not useful, and $B_{\min}$ has to be varied (in the limit of $\eta \to \infty$) with granularity equal to b_d in order to have any impact on system performance; hence the 'cloud' shape of the curves in Fig. 11.

The impact of b_u. We now show the most interesting phenomena that we have observed while considering the joint impact of finite b_d and finite b_u. In this case, the distribution of the number of replicas of a requested content is important; thus, we can take full advantage of the popularity and replication model developed in Section 6. We consider a family of popularity curves $p_w(t)$ which decay exponentially over time, i.e. $p_w(t) = V\mathrm{e}^{-\theta t}$. We assume that both the initial popularity value V and the decay parameter θ are random variables. The cancelation rate δ is assumed to be the

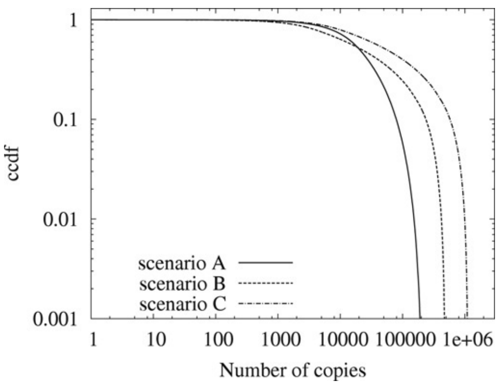

Fig. 12. Ccdf of the number of replicas of a requested object in the whole P2P community.

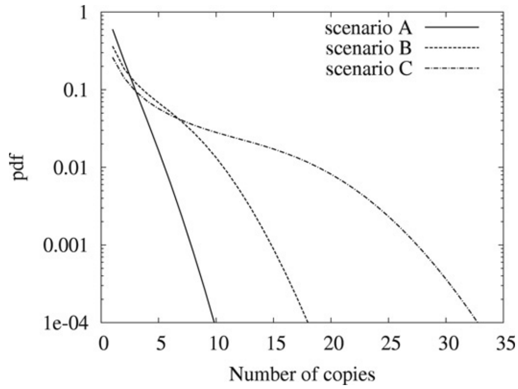

Fig. 13. Conditional pdf of the number of online replicas within the ISP network.

same for all objects. More specifically, we consider the following three scenarios: (A) V uniform in [1, 10], θ uniform in [0.001, 1], $\delta = 0.001$; (B) V uniform in [1, 10], θ uniform in [0.001, 1], $\delta = 0.02$; (C) V uniform in [1, 10], θ uniform in [0.001, 3], $\delta = 0.02$. Notice that the unit of measure of θ and δ is days^{-1}. We assume that the total number of copies in the P2P community is $Z = 4 \cdot 10^8$ (an average of eight objects per peer). This is consistent with the fact that the total amount of data stored by peers, according to online statistics reported by popular P2P file-sharing clients, is of the order of several thousand TBytes, while the average object size is around 30 MB [5].

We fix the number of ISP users to $n = 20\,000$, and consider an online probability $p_{\text{on}} = 0.05$ such that, on average, we have 1000 online peers in the ISP network. The complementary cumulative distribution function of the number of copies of a requested object in the whole P2P system is depicted in Fig. 12 for the three considered scenarios. In Fig. 13 we have plotted the resulting conditional pdf $P(a_i)$ of the number of online replicas within the ISP network, computed as explained at the end of Section 7. We observe that the mean number of replicas of a requested object increases when we increase either the cancelation rate δ or the popularity decay parameter θ (notice that the total number of copies in the system is kept constant).

In the following experiments we fix $q = 0.9$, $\widehat{\sigma} = \sigma_{\min} = 0.5$. We first take scenario (A), with $\gamma = 0.9$, and assume that $b_u = b_d$, i.e. the upload bandwidth of the users is equal to their download bandwidth. To account for user diversity, we consider a power-law user weight function with $\nu = 1.2$, the same as that considered in Fig. 10. In Fig. 14 we show the minimum required bandwidth $B_{\min}$ for different values of the user impatience parameter η. Quite surprisingly, $B_{\min}$ can exhibit a non-monotonic behavior as we increase the value of $b_u = b_d$, especially when users are very impatience. This can happen when $B_{\min}$ is comparable with b_d, and can be explained by the fact that the benefit resulting from increased upload bandwidth b_u is offset by the effect of transfer aborts by users who do not get a correspondingly high download rate when they have to use the external link bandwidth.

In general, to exploit traffic locality better, it is preferable to provide to the users a larger upload bandwidth b_u, rather than a larger download bandwidth b_d, contrary to common practice (e.g. ADSL lines). This is illustrated in

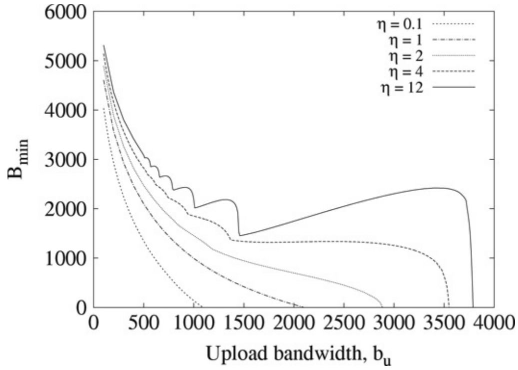

Fig. 14. $B_{\min}$ as a function $b_u = b_d$, for different values of the impatience parameter η.

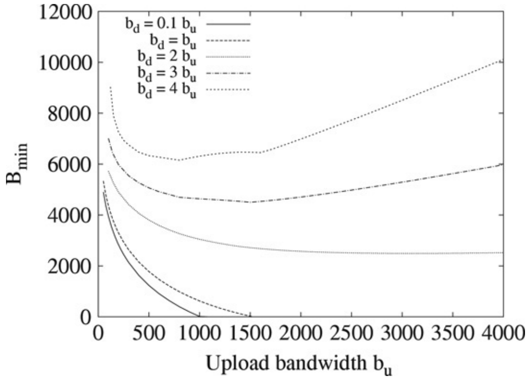

Fig. 15. $B_{\min}$ as a function of b_u for $\eta = 0.5$ and different values of b_d.

Fig. 15, where we fix $\eta = 0.5$ and vary the ratio between b_u and b_d. Even if users are quite patient ($\eta = 0.5$), it can be very detrimental to the ISP to provide values of b_d much larger than b_u. This choice can shift the point at which the external bandwidth is no longer necessary ($B_{\min} = 0$), or even result in increasing values of $B_{\min}$ for increasing user access rates. We conclude that asymmetric access rates are ill-suited to P2P traffic, and this should be taken into account by ISPs that want to be friendly towards P2P usage.

Next, we consider scenario (B), in which the number of replicas of an object is large enough that $p \sim 1$, i.e. it is very likely that a copy of a requested object is available at an internal peer. Here, we show another interesting phenomenon, that is, $\gamma = 1$ (perfect exploitation of traffic locality) is not always optimal, especially when the ISP provides an excessively small upload bandwidth to the users. This is illustrated in Fig. 16, in which we plot $B_{\min}$ as a function $b_u = b_d$, with $\eta = 1$, for different values of γ. When b_u is larger than 250, $\gamma = 1$ allows the ISP not to provide any external bandwidth ($B_{\min} = 0$). However, when $b_u < 250$, $\gamma = 1$ is not feasible, i.e. the bandwidth available within the ISP network is not enough to guarantee the throughput demanded by subscribers. In this case, the ISP has artificially to redirect queries out of its network ($\gamma < 1$), forcing the downloads to use the external bandwidth. Actually, for each value of $b_u < 250$ there exists a critical γ, illustrated in Fig. 16 by a thick solid line, such that values of γ larger than this critical value are not feasible (i.e. users are not satisfied by the service). However, unnecessarily reducing γ below the critical value results in increasingly higher values of $B_{\min}$ (see Fig. 16).

At last, we show the impact of different levels of user heterogeneity, by considering user weight functions $w(j) = j^{-\nu}$ with different exponents ν. We consider scenario (C), with $\gamma = 0.8$, $\eta = 1$. This time, we compute the maximum service level σ that can be provided to the users, assuming that the average service level obtained by other peers in the community is constant, $\widehat{\sigma} = 0.5$. The results are shown in Fig. 17, as a function of $b_u = b_d$. We observe that, the more even the internal load distribution (small ν), the bigger the effectiveness of traffic locality and thus the quality of service that can be offered to the subscribers. Notice that at all considered points $B_{\min} = 0$.

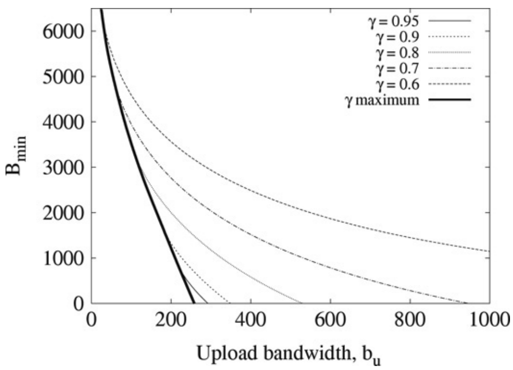

Fig. 16. $B_{\min}$ as a function of $b_u = b_d$, for different values of γ.

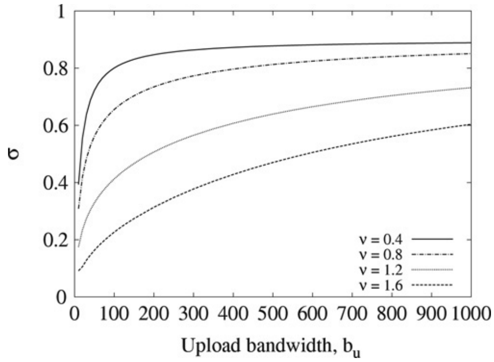

Fig. 17. Maximum σ as a function of $b_u = b_d$, for different values of ν.

Conversely, for a given desired service level, curves such as those in Fig. 17 can be used to compute the minimum upload bandwidth that has to be provided to the users.

9. Conclusions and further developments

The shift from traditional Internet applications to P2P applications is a rather new phenomenon with significant implications for ISPs, which are still not well understood. Although different strategies have been proposed to manage P2P traffic in a network, there is an increasing need of modeling techniques to evaluate their applicability and effectiveness in specific network contexts. This paper represents a first step towards an analytical understanding of the tussle between ISPs and P2P users. Our modeling framework takes into account a variety of factors that affect the system performance, and in particular the ability to exploit traffic locality to reduce transit costs. Our analysis has revealed many interesting, sometimes counter-intuitive, phenomena occurring while varying the system parameters. We are planning to extend our analysis along different directions. An important refinement of the model would be to consider users with different service level expectation for the same subscription fee, so that the ISP has to choose how many of them to attract. Another interesting development is the study of the interaction among multiple ISPs competing with each other for the same population of users.

References

[1] Cachelogic Research, Peer-to-Peer in 2005. URL: http://www.cachelogic.com/research/.

[2] T. Karagiannis, A. Broido, M. Faloutsos, K. Claffy, Transport layer identification of P2P traffic, in: Proc. IMC ’04, Taormina, Italy, 2004.

[3] T. Karagiannis, A. Broido, N. Brownlee, K. Claffy, M. Faloutsos, Is P2P dying or just hiding? in: IEEE Globecom 2004, Global Internet and Next Generation Networks, Dallas, Texas, USA, 29 November–3 December 2004.

[4] Meeting the Challenge of Today’s Evasive P2P Traffic, Sandvine co., Waterloo, Canada, 2004, An Industry White Paper.

[5] P.K. Gummadi, R.J. Dunn, S. Saroiu, S.D. Gribble, H.M. Levy, J. Zahorjan, Measurement, modeling, and analysis of a peer-to-peer file-sharing workload, in: Proceedings of the 19th ACM Symposium on Operating Systems Principles, SOSP-19, Bolton Landing, NY, USA, October 2003.

[6] M. Ripeanu, I. Foster, Mapping Gnutella network, in: Proc. of the First International Workshop on Peer-to-Peer Systems, IPTPS 2002, Cambridge, MA, USA, in: LNCS, vol. 2429, Springer, 2002.
[7] N. Lleibowitz, A. Bergman, R. Ben Shaul, A. Shavit, Are file swapping networks cacheable? Characterizing P2P traffic, in: Proceeding of 7th Int. WWW Caching Workshop, August 2002.
[8] O. Saleh, M. Hefeeda, Modeling and caching of peer-to-peer traffic, in: Proc. of ICPN '06, Santa Barbara, CA, US, November 2006.
[9] G. Shen, Y. Wang, Y. Xiong, B.Y. Zhao, Z.-L. Zhang, Hptp: Relieving the tension between isps and p2p, in: Proc. of IPTPS '07, Bellevue, WA, US, February 2007.
[10] T. Karagiannis, P. Rodriguez, D. Papagiannaki, Should ISPs fear peer-assisted content distribution? in: Prof. ACM USENIX IMC, Berkeley, CA, USA, 2005.
[11] J.H. Wang, D.M. Chiu, J.C. Lui, Modeling the peering and routing tussle between isps and p2p applications, in: Proc. of IEEE International Workshop on Quality of Service, IWQoS, Yale University, New Haven, CT, June 2006, pp. 51–59.
[12] D. Stutzbach, R. Rejaie, Understanding churn in peer-to-peer networks, in: Proc. of IMC '06, Rio de Janeiro, Brazil, 2006.
[13] D. Daley, D. Vere-Jones, An introduction to the theory of point processes, Springer-Verlag, 1988.
[14] S.B. Fred, T. Bonald, A. Proutière, G. Régnié, J.W. Roberts, Statistical bandwidth sharing: A study of congestion at flow level, in: Proc. of ACM SIGCOMM '01, San Diego, California, August 2001, pp. 111–122.
[15] J.W. Cohen, The multiple phase service network with generalized processor sharing, Acta Inform. 12 (1979) 245–284.
[16] B. Rosen, Asymptotic theory for successive sampling with varying probabilities without replacement, Ann. Math. Statist. 43 (2) (1972).

Michele Garetto (S'01–M'04) received the Dr.Ing. degree in Telecommunication Engineering and the Ph.D. degree in Electronic and Telecommunication Engineering, both from Politecnico di Torino, Italy, in 2000 and 2004, respectively. In 2002, he was a visiting scholar with the Networks Group of the University of Massachusetts, Amherst, and in 2004 he held a postdoctoral position at the ECE department of Rice University, Houston. He is currently assistant professor at the Computer Science Department of the University of Torino, Italy. His research interests are in the field of performance evaluation of wired and wireless communication networks.

Daniel Ratton Figueiredo received the B.S. degree and the M.Sc. degree in Computer Science from the Federal University of Rio de Janeiro, Brazil in 1996 and 1999, respectively. He recently received the M.Sc. degree and the Ph.D. degree from the Department of Computer Science at the University of Massachusetts, Amherst (in 2005). Following graduation he spent almost two years working as a post-doc researcher at the Swiss Federal Institute of Technology Lausanne (EPFL) in the Laboratory for Computer Communications and Applications (LCA). He was recently appointed assistant professor in COPPE/PESC Department at the Federal University of Rio de Janeiro, Brazil. His research interests include modeling and analysis of computer networks, and in particular, mechanisms used in the Internet.

Rossano Gaeta was born in Torino, Italy in 1968. He received his Laurea and Ph.D degrees in computer science from the University of Torino in 1992 and 1997, respectively. He is currently associate professor at the Computer Science Department of the University of Torino. His research interests include the use of formal models for dependability analysis of time critical systems and the performance evaluation of peer-to-peer computing systems and large scale sensor networks.

Matteo Sereno was born in Nocera Inferiore, Italy. He received his Laurea degree in Computer Science from the University of Salerno, Italy, in 1987, and his Ph.D. degree in Computer Science from the University of Torino, Italy, in 1992. He is currently full professor at the Computer Science Department of the University of Torino. His current research interests are in the area of performance evaluation of computer systems, communication networks, peer-to-peer systems, sensor networks, queueing networks and stochastic Petri net models.