Multiclass P2P Networks: Static Resource Allocation for Service Differentiation and Bandwidth Diversity

Florence Clévenot-Perronnin* Philippe Nain* Keith W. Ross†

Abstract

We propose a multiclass fluid model for BitTorrent-like content distribution systems. The new model can model heterogeneous peers, in which peers have different access bandwidths. The model can also model BitTorrent-like systems which provide differential service (for example, first class and second class service) to the participating peers. The fluid model leads to a non-linear system of differential equations with special structure. For the service differentiation problem, we prove that the system of differential equations admits a unique stable equilibrium, that we compute in closed-form. We also give the average download times for both classes. For the bandwidth diversity problem, we show that the system of differential equations has a stable state that may depend on the initial conditions. We give the average download time of both classes for each reachable steady-state.

Keywords: P2P networks, fluid models, service differentiation

1 Introduction

In a traditional client-server content distribution system, such as distribution from an ordinary Web server, a large number of clients download content from a single server. If the single server cannot keep up with the demand from all the clients, the load can potentially

*INRIA, BP 93, 06902 Sophia Antipolis, France. E-mail: {Florence.Clevenot, Philippe.Nain}@sophia.inria.fr.
†Polytechnic University, Six MetroTech Center, Brooklyn, NY 11201, USA. E-mail: ross@poly.edu.

be handled by replacing the server with a server farm and increasing the access bandwidth from the server farm. Although it is possible in theory to match any demand with a sufficient number of servers and sufficiently wide access pipes, the cost can easily become prohibitive.

BitTorrent is a content-distribution booster which enables a content provider to distribute popular content to large number of clients without the need of large server farms and expensive high-speed Internet connections. The idea in essence is to split the file into small chunks, distribute different chunks to different downloading peers, and then have the different downloading peers obtain their missing chunks from each other. In this manner, the clients become servers, contributing bandwidth to the content-distribution system. This approach has proved to be a highly successful mechanism to distribute popular content at low cost. In BitTorrent terminology, the servers that make available the entire file are called "seeds". The clients that are collecting and sharing chunks are called "leechers". Once a leecher has downloaded the entire file, it becomes a seed for as long as it continues to distribute chunks to other clients. The BitTorrent protocol includes a "tit-for-tat" mechanism to ensure that leechers not only download content but also upload content [1]. BitTorrent is a peer-to-peer system since clients (peers) upload chunks directly to each other.

Qiu and Srikant [7] developed a tractable fluid model for BitTorrent-like content distribution systems. The model sheds insight on throughput, average download times, and stability of BitTorrent-like systems. Although the model is elegant and tractable, it has limited applicability. First, the model assumes that all peers are homogeneous, with all peers having the same upload and download capacity. In actuality, peers have diverse bandwidth characteristics, including dial-up modem access, broadband access (cable and ADSL), and high-speed Ethernet access. Second, the model does not allow for the exploration of distribution systems that provide application-layer differentiated services. Indeed, it is natural to conceive of a BitTorrent-like system in which there are, say, first-class peers and second-class peers. The first-class peers pay more (in some sense) and should receive better service – that is, shorter average download times – than the second-class peers. This is a form of "application-layer differentiated-service" as the service differentiation would be provided by the BitTorrent-like application rather than by the core of the Internet. Intuitively, BitTorrent-like systems could provide differentiated service by having the seeds and leechers allocate more of their upload bandwidth to first-class peers.

In this paper we propose a multiclass fluid model for BitTorrent-like content distribution systems. The new fluid model can model both heterogeneous peer access and multiple differentiated service classes. Our multiclass fluid model results in a system of differential equations which generalize the single-class equation in [7]. The equations are significantly more complex and difficult to resolve, as they explicitly distinguish between the various classes. The system of differential equations are so-called “linear switched systems” which are nonlinear differential equations with special structure. Nevertheless, for a number of important special cases, we explicitly resolve the equations, obtaining closed-form solutions for average download times for each of the classes.

In particular, we consider the special case where downloaders leave the system immediately after completing their download. This is a worst-case scenario since altruistic seeds could instead stay in the system when they have completed their download, contributing bandwidth and providing any missing chunk to other peers. For the service differentiation problem we prove that the system of differential equations governing the system dynamics admits a unique stable equilibrium, that we compute in closed-form. From this result, we find the average download time for each class of peers and show how this result can be used to achieve service differentiation among the peers. We also indicate to what extent our results remain valid when seeds stay in the system for a non-negligible amount of time.

In the second part of the paper, we address the bandwidth diversity problem. We show that the system of differential has a stable stationary state that may depend on the initial conditions. We identify all stationary solutions and compute the average download time associated with each solution. Last, we minimize the maximum average download time of both classes, regardless of the initial conditions.

The paper is organized as follows. In Section 2 we introduce the multiclass model and derive the equations governing the system dynamics. Section 3 and Section 4 provide results for the service differentiation problem and bandwidth diversity problem, respectively. Section 5 concludes the paper.

2 Multiclass Model

In this paper we consider a BitTorrent-like system with two classes of peers, with the classes denoted by $i = 1$ and $i = 2$. All the peers in both classes want to obtain the single file F. Without loss of generality, we take the file size to be equal to 1. Each class has seeds and downloaders (leechers). Seeds have all of the file, whereas downloaders have only portions of the file. When a downloader obtains the whole file, it immediately becomes a seed. Let $y_i(t)$ and $x_i(t)$ denote the number of seeds and downloaders, respectively, for class-i peers at time t. In this paper, we are particularly interested in the steady-state behavior of y_i and x_i, $i = 1, 2$. We need to also define the following:

- Let λ_i be the rate at which new class-i downloaders arrive. Whenever a new class-i downloader arrives, then x_i is incremented by 1.
- Let μ_i be the upload bandwidth of a peer from class i.
- Let c_i be the download bandwidth of a peer from class i. We make the realistic assumption that $c_i \geq \mu_i$, which is consistent with the contemporary access technologies. Whenever a class-i peer has fully downloaded the file, x_i is decremented by 1 and y_i is incremented by 1.
- As in [7], we allow downloaders to abort downloading before fully obtaining the file. Let θ_i be the rate at which class-i downloaders abort. Whenever a class-i downloader aborts, x_i is decremented by 1.
- Let γ_i denote the rate at which class-i seeds leave the system. Whenever a class-i seed leaves the system, y_i is decremented by 1.
- Let $\eta_i \in (0, 1)$ denote the efficiency of class-i downloaders, which is the average amount of a downloader's upload bandwidth that is being used for content distribution. This parameter was first introduced in [7] in the single-class case.

We now discuss the resource allocation policy. A peer (seed or downloader) will upload chunks to multiple peers simultaneously. The aggregate rate at which a class-i seed peer uploads is μ_i; the aggregate rate at which a class-i downloader peer uploads is $\eta_i \mu_i$. A peer will allocate its upload rate between the two classes of peers. For a class-i peer, let $\alpha_i(x_1, x_2)$

(resp. $1-\alpha_i(x_1,x_2)$) be the fraction of its upload rate that is allocated to class-i peers, that is, to peers in its own class (resp. to peers in the other class) when there are x_1 class-1 downloaders present and x_2 class-2 downloaders present. Thus, $\alpha_i(x_1,x_2)$ lies between 0 and 1. We refer to $(\alpha_1(x_1,x_2),\alpha_2(x_1,x_2))$ as a **dynamic allocation policy**.

In this paper we limit our attention to **static allocation policies**, namely, policies of the form $\alpha_i(x_1,x_2)=\alpha_i$ for all x_1 and x_2 for $i=1,2$. We will consider dynamic policies in a future work.

Our model of the two-class multiclass P2P network is now complete. Figure 1 summarizes the states and rates in the system.

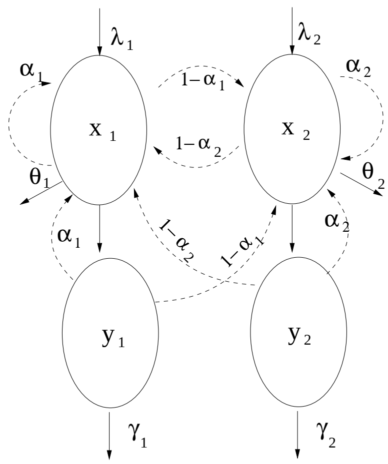

Figure 1: General model for a two-class P2P file dissemination system.

We now develop a system of differential equations for the fluid-version of the above multiclass model.

At time t, the total upload rate provided by class-i peers to peers of class i is $\alpha_i\mu_i(\eta_i x_i(t)+y_i(t))$ and to peers of the other class is $(1-\alpha_i)\mu_i(\eta_i x_i(t)+y_i(t))$. Therefore, the total upload rate provided by class-i peers is $\mu_i(\eta_i x_i(t)+y_i(t))$. Let $k=3-j$ designate the other class. The total download rate provided to peers of class i cannot exceed $c_i x_i(t)$ so that the total flow rate out of state $x_i(t)$ is $\min(c_i x_i(t), \alpha_i\mu_i(\eta_i x_i(t)+y_i(t))+(1-\alpha_k)\mu_k(\eta_k x_k(t)+y_k(t)))$, to which we must add $\theta_i x_i(t)$, the total flow rate at which downloaders leave the system without having downloaded the entire file. By definition, the flow rate into state $x_i(t)$ is λ_i.

Hence, the time-evolution of $(x_1(t), x_2(t))$ is governed by the following differential equations

$$\frac{dx_i(t)}{dt} = \lambda_i - \theta_i x_i(t) - \min\left(c_i x_i(t),\ \alpha_i \mu_i(\eta_i x_i(t) + y_i(t)) + (1-\alpha_k)\mu_k(\eta_k x_k(t) + y_k(t))\right) \quad (1)$$

for $i = 1, 2$ and $k = 3 - i$.

Similarly, we find that the total flow rate into state $y_i(t)$ is given by the total rate at which downloaders become seeds, namely $\mu_i(\eta_i x_i(t) + y_i(t)) + (1-\alpha_k)\mu_k(\eta_k x_k(t) + y_k(t))$ as explained above, while the total flow rate out of state $y_i(t)$ is simply $\gamma_i y_i(t)$. This gives the following equations for the time-evolution of $(y_1(t), y_2(t))$

$$\frac{dy_i(t)}{dt} = \min\left(c_i x_i(t),\ \alpha_i \mu_i(\eta_i x_i(t) + y_i(t)) + (1-\alpha_k)\mu_k(\eta_k x_k(t) + y_k(t))\right) - \gamma_i y_i(t) \quad (2)$$

for $i = 1, 2$ and $k = 3 - i$.

Equations (1)-(2) fully define the system dynamics.

We will be particularly interested in the case where downloaders leave the system at once when they have completed their download, namely $1/\gamma_1 = 1/\gamma_2 = 0$. There are two reasons why we will be considering this situation. First, it will yield much more tractable equations, as shown next. Second, this case represents a worst-case situation, where peers are not willing to cooperate, and leave the system as soon as they have downloaded the file. In this case, they never become seeds, which implies $y_i(t) = 0$ for all $t > 0$. As a result, system (1) reduces to

$$\frac{dx_i(t)}{dt} = \lambda_i - \theta_i x_i(t) - \min\left(c_i x_i(t),\ \alpha_i \beta_i x_i(t) + (1-\alpha_k)\beta_k x_k(t)\right) \quad (3)$$

for $i = 1, 2$ and $k = 3 - i$, where

$$\beta_i := \mu_i \eta_i. \quad (4)$$

Note that

$$c_i > \beta_i, \quad i = 1, 2 \quad (5)$$

since we have assumed that $c \geq \mu_i$ and $0 < \eta_i < 1$. In matrix form (3) writes

$$\dot{\mathbf{x}}(t) = A_{\sigma(\mathbf{x}(t))}\, \mathbf{x}(t) + \mathbf{b} \quad (6)$$

with $\mathbf{x}(t) := (x_1(t), x_2(t))^T$ and $\mathbf{b} = (-\lambda_1, -\lambda_2)^T$ (as usual $\mathbf{v}^T$ denotes the transpose vector of the vector $\mathbf{v}$). In (6) σ is an integer-value mapping, taking values in $\sigma \in \{1, 2, 3, 4\}$, given by

$$\sigma(\mathbf{x}) = 1+2\times\mathbf{1}(c_1x_1 \geq \alpha_1\beta_1x_1+(1-\alpha_2)\beta_2\eta_2x_2)+\mathbf{1}(c_2x_2 \geq \alpha_2\beta_2\eta_2x_2+(1-\alpha_1)\beta_1\eta_1x_1) \quad (7)$$

for $\mathbf{x} = (x_1, x_2)$, where $\mathbf{1}(A)$ denotes the indicator function of the event A (i.e. $\mathbf{1}(A) = 1$ if A holds and zero otherwise). The mapping σ is called a *switching condition* and a system like (6) is called a *switched system* [6, 5]. The 2-by-2 matrices A_i, $i = 1, \ldots, 4$, can easily be identified from (3).

The model where $1/\gamma_1 = 1/\gamma_2 = 0$ will be referred to as the *no-seed* model. A natural question is: how do downloaders ever get any chunk if there are no seeds? Here, we make a distinction between two notions of seeds. A BitTorrent-like system needs, at startup time, at least one seed, for as long as it needs to upload (at least) a complete copy of the file. Though this bootstrap seed is mandatory to make the file available, it may leave long before the system reaches a steady-state. Therefore, its role is limited to starting the torrent, and is negligible on the long-term. Note that the general system (1)-(2), as well as the single-class model in [7], also neglect this bootstrap seed, since the system may have a nonzero solution even if $y_i(0) = 0$ for $i = 1, 2$. Downloaders which have a complete copy of the file, on the other hand, will have an impact on the steady-state since they belong to the long-term dynamics of the system. These regular seeds are considered in (1)-(2), whereas the no-seed model assumes they leave the system immediately.

We conclude this section by introducing the cost functions that we will consider throughout the paper. Let ϕ_i be the *download cost* of peers of class i, which is defined as the expected download time given that the peer completes the download. An analytic expression for ϕ_i can easily be derived as follows. Assume that $x_i(t)$ has a stationary regime, denoted by $\bar{x}_i$. By Little's formula, the expected download time T_i for peers of class i is given by $T_i = \bar{x}_i/\lambda_i$. On the other hand, the stationary probability p_i that a class-i peer completes its download is $p_i = (\lambda_i - \theta_i\bar{x}_i)/\lambda_i$. Therefore, the download cost for peers of class i takes the form

$$\phi_i = \frac{\bar{x}_i}{\lambda_i - \theta_i\bar{x}_i}, \quad i = 1, 2. \quad (8)$$

In the next two sections we shall address two different problems corresponding to different subsets of (static) allocation policies: $(\alpha_1, \alpha_2) = (\alpha, 1 - \alpha)$, referred to as the service differentiation problem (Section 3), and $(\alpha_1, \alpha_2) = (\alpha, \alpha)$, referred to as the bandwidth diversity problem (Section 4). Both problems will be considered for no-seed models.

3 Resource Allocation Policy for Service Differentiation

In this section we address the service differentiation problem for the no-seed model (unless otherwise mentioned). For the sake of simplicity we further restrict the analysis to the case where all peers have the same download/upload bandwidths and the same efficiency parameters. In other words, we assume that $1/\gamma_i = 0$, $c_i = c$, $\mu_i = \mu$ and $\eta_i = \eta$ for $i = 1, 2$. We define $\beta := \mu\eta$.

We recall that the service differentiation problem corresponds to the situation where $\alpha_1 = 1 - \alpha_2 = \alpha$ (see end of Section 2).

Our goal is to resolve the resulting system of differential equations (see below) and determine the download cost (defined in (8)) of the two classes of peers. In particular, we shall show that differential service can indeed be provided to the two classes of peers via the allocation parameter α.

Under the above assumptions the system of differential equations (3) governing the dynamics of $(x_1(t), x_2(t))$ simplifies to

$$\begin{aligned} \tfrac{dx_1(t)}{dt} &= \lambda_1 - \theta_1 x_1(t) - \min\left(cx_1(t),\, \alpha\beta(x_1(t) + x_2(t))\right) \\ \tfrac{dx_2(t)}{dt} &= \lambda_2 - \theta_2 x_2(t) - \min\left(cx_2(t),\, (1-\alpha)\beta(x_1(t) + x_2(t))\right). \end{aligned} \tag{9}$$

In matrix notation this system is given by (6) with the switching condition ($\mathbf{x} = (x_1, x_2)$)

$$\sigma(\mathbf{x}) = 1 + 2 \times \mathbf{1}(cx_1 \geq \alpha\beta(x_1 + x_2)) + \mathbf{1}(cx_2 \geq (1-\alpha)\beta(x_1 + x_2)). \tag{10}$$

We introduce the new parameters

$$\begin{aligned} a_1 &:= \max\left(0, 1 - \frac{c\lambda_2(\theta_1 + \beta)}{D}\right) \\ a_2 &:= \min\left(1, \frac{c\lambda_1(\theta_2 + \beta)}{D}\right) \end{aligned}$$

with $D := \beta(\lambda_1(\theta_2 + c) + \lambda_2(\theta_1 + c))$.

Proposition 3.1 below computes the equilibrium point of the switched system (9).

Proposition 3.1 (Equilibrium point for service differentiation) *Regardless of the initial condition* $\mathbf{x}(0)$*, the system of equations (9) has a unique equilibrium point* $\bar{\mathbf{x}}$ *given by*

$$\bar{\mathbf{x}}^T = \begin{cases} \left(\dfrac{\lambda_1 - \alpha\frac{\lambda_2\beta}{\theta_2+c}}{\theta_1+\alpha\beta}, \dfrac{\lambda_2}{\theta_2+c}\right) & \text{if } 0 \le \alpha < a_1 \\[2ex] \left(\dfrac{\lambda_1(\theta_2+(1-\alpha)\beta) - \lambda_2\alpha\beta}{\theta_2(\theta_1+\alpha\beta)+\theta_1(1-\alpha)\beta}, \dfrac{\lambda_2(\theta_1+\alpha\beta) - \lambda_1(1-\alpha)\beta}{\theta_2(\theta_1+\alpha\beta)+\theta_1(1-\alpha)\beta}\right) & \text{if } a_1 \le \alpha \le a_2 \\[2ex] \left(\dfrac{\lambda_1}{\theta_1+c}, \dfrac{\lambda_2 - (1-\alpha)\frac{\lambda_1\beta}{\theta_1+c}}{\theta_2+(1-\alpha)\beta}\right) & \text{if } a_2 < \alpha \le 1. \end{cases} \tag{11}$$

$\diamond$

Proof. We first check that if $\lim_{t\to\infty} \mathbf{x}(t)$ exists, then it is given by (11).

Assume that $\lim_{t\to\infty} \mathbf{x}(t) = \bar{\mathbf{x}}$. Letting $t \to \infty$ in (6) yields

$$A_{\sigma(\bar{\mathbf{x}})}\, \bar{\mathbf{x}} + \mathbf{b} = 0 \tag{12}$$

where σ is given in (10). We consider separately the four systems of linear equations obtained from (12) when (a) $\sigma(\bar{\mathbf{x}}) = 1$, (b) $\sigma(\bar{\mathbf{x}}) = 2$, (c) $\sigma(\bar{\mathbf{x}}) = 3$ and (d) $\sigma(\bar{\mathbf{x}}) = 4$.

(a) $\sigma(\bar{\mathbf{x}}) = 1$ or equivalently $c\bar{x}_1 < \alpha\beta(\bar{x}_1 + \bar{x}_2)$ and $c\bar{x}_2 < (1-\alpha)\beta(\bar{x}_1 + \bar{x}_2)$.

The download rate is the bottleneck for both classes of peers. We find

$$\bar{\mathbf{x}}^T = \left(\frac{\lambda_1}{\theta_1+c}, \frac{\lambda_2}{\theta_2+c}\right). \tag{13}$$

(b) $\sigma(\bar{\mathbf{x}}) = 2$ or equivalently $c\bar{x}_1 < \alpha\beta(\bar{x}_1 + \bar{x}_2)$ and $c\bar{x}_2 \ge (1-\alpha)\beta(\bar{x}_1 + \bar{x}_2)$.

The bottleneck is the download rate for class-1 peers and the upload rate for class-2 peers. We find

$$\bar{\mathbf{x}}^T = \left(\frac{\lambda_1}{\theta_1+c}, \frac{\lambda_2 - (1-\alpha)\frac{\lambda_1\beta}{\theta_1+c}}{\theta_2+(1-\alpha)\beta}\right). \tag{14}$$

(c) $\sigma(\bar{\mathbf{x}}) = 3$ or equivalently $c\bar{x}_1 \ge \alpha\beta(\bar{x}_1 + \bar{x}_2)$ and $c\bar{x}_2 < (1-\alpha)\beta(\bar{x}_1 + \bar{x}_2)$.

The bottleneck is the download rate for peers of class 2 and the upload rate for peers of class 1. In this case

$$\bar{\mathbf{x}}^T = \left(\frac{\lambda_1 - \alpha\frac{\lambda_2\beta}{\theta_2+c}}{\theta_1+\alpha\beta}, \frac{\lambda_2}{\theta_2+c}\right). \tag{15}$$

(d) $\sigma(\bar{\mathbf{x}}) = 4$ or equivalently $c\bar{x}_1 \ge \alpha\beta(\bar{x}_1 + \bar{x}_2)$ and $c\bar{x}_2 \ge (1-\alpha)\beta(\bar{x}_1 + \bar{x}_2)$.

The bottleneck is the download rate for both classes of peers. The equilibrium point is

$$\bar{\mathbf{x}}^T = \left(\frac{\lambda_1(\theta_2 + (1-\alpha)\beta) - \lambda_2 \alpha\beta}{\theta_2(\theta_1 + \alpha\beta) + \theta_1(1-\alpha)\beta}, \frac{\lambda_2(\theta_1 + \alpha\beta) - \lambda_1(1-\alpha)\beta}{\theta_2(\theta_1 + \alpha\beta) + \theta_1(1-\alpha)\beta} \right). \tag{16}$$

In the following, we call "type-i equilibrium" the equilibrium found when $\sigma(\bar{\mathbf{x}}) = i$.

The next step is to check if a type-i equilibrium may exist, namely, if $\sigma(\bar{\mathbf{x}}) = 1$ (resp. $\sigma(\bar{\mathbf{x}}) = 2$, $\sigma(\bar{\mathbf{x}}) = 3$, $\sigma(\bar{\mathbf{x}}) = 4$) when $\bar{\mathbf{x}}$ is given by (13) (resp. (14), (15), (16)).

It is easily seen that a type-1 equilibrium may only exist if $c \le \beta$. Since this condition is never met (use (5) with $c_i = c$ and $\beta_i = \beta$) we conclude that there is no type-1 equilibrium.

Recall that $0 \le \alpha \le 1$. We prove in Appendix A that a type-2 equilibrium may only exist if $a_2 < \alpha \le 1$. The same type of analysis shows that a type-3 equilibrium may only exist if $0 \le \alpha < a_1$, and that a type-4 equilibrium may only exist if $a_1 \le \alpha \le a_2$.

This concludes the proof that, if $\lim_{t\to\infty} \bar{\mathbf{x}}(t) = \bar{\mathbf{x}}$ exists, then $\bar{\mathbf{x}}$ is given by (11) (regardless of the initial condition).

We now turn to the proof that $\lim_{t\to\infty} \bar{\mathbf{x}}(t)$ exists. To this end, we investigate the nature of the equilibrium of each of the linear systems $\dot{\mathbf{x}}(t) = A_i\, \mathbf{x}(t) + \mathbf{b}$, for $i = 2, 3, 4$, with

$$A_2 = \begin{pmatrix} -(\theta_1 + c) & 0 \\ -(1-\alpha)\beta & -(\theta_2 + (1-\alpha)\beta) \end{pmatrix} \quad A_3 = \begin{pmatrix} -(\theta_1 + \alpha\beta) & -\alpha\beta \\ 0 & -(\theta_2 + c) \end{pmatrix}$$

and

$$A_4 = \begin{pmatrix} -(\theta_1 + \alpha\beta) & -\alpha\beta \\ -(1-\alpha)\beta & -(\theta_2 + (1-\alpha)\beta) \end{pmatrix}.$$

Recall that the equilibrium of the system $\dot{\mathbf{x}}(t) = A_i\, \mathbf{x}(t) + \mathbf{b}$ is stable if and only if all eigenvalues of the matrix A_i have strictly negative real parts [4]. It is easily seen that A_2 and A_3 have two strictly negative eigenvalues, given by $(-(\theta_1 + c), -(\theta_2 + (1-\alpha)\beta))$ and $(-(\theta_1 + \alpha\beta), -(\theta_2 + c))$, respectively. The same property holds for A_4. To see this, denote by $D(c, r)$ the closed disc of center c and radius r in the complex plane. From Geršgorin circle theorem [2, p. 344] we know that both eigenvalues of A_4 lie in the region $D(-\theta_1 - \alpha\beta, \alpha\beta) \cup D(-\theta_2 - (1-\alpha)\beta, (1-\alpha)\beta)$, from which the result follows.

We have now proved the local stability of the equilibrium of each linear subsystem of (9). However, a rigorous proof of the global stability of (9) would require more attention. For brevity, we do not address this problem here. The interested reader can refer to [5] for the stability of linear switched systems.

In summary, we have shown that for a given value of α, a unique equilibrium exists, is given in (11), and is stable. This completes the proof. ∎

How can we achieve a target QoS ratio k?

It is now possible to achieve service differentiation using parameter α as follows.

The goal is to differentiate the download costs ϕ_1 and ϕ_2 of class-1 and class-2 peers, respectively. These costs are given in the next proposition.

Proposition 3.2 (Download costs for service differentiation) *In a no-seed model, the download cost ϕ_i of class-i peers in the service differentiation problem is given by:*

$$\begin{array}{lll}
\phi_1 = \dfrac{\lambda_1(\theta_2+c)-\alpha\lambda_2\beta}{\alpha\beta(\lambda_2\theta_1+\lambda_1(\theta_2+c))}, & \phi_2 = \dfrac{1}{c} & \text{if } 0 \le \alpha < a_1 \\
\phi_1 = \dfrac{\lambda_1(\theta_2+\beta)-\alpha\beta(\lambda_1+\lambda_2)}{\alpha\beta(\lambda_2\theta_1+\lambda_1\theta_2)}, & \phi_2 = \dfrac{\lambda_2\theta_1-\lambda_1\beta+\alpha\beta(\lambda_1+\lambda_2)}{(1-\alpha)\beta(\lambda_2\theta_1+\lambda_1\theta_2)} & \text{if } a_1 \le \alpha \le a_2 \\
\phi_1 = \dfrac{1}{c}, & \phi_2 = \dfrac{\lambda_2(\theta_1+c)-\lambda_1\beta+\alpha\lambda_1\beta}{(1-\alpha)\beta(\theta_2\lambda_1+\lambda_2(\theta_1+c))} & \text{if } a_2 < \alpha \le 1.
\end{array}$$

◇

First, note that in the service differentiation problem, we considered the static allocation policy $(\alpha, 1-\alpha)$. Since the two classes have the same bandwidth characteristics (i.e. $c_1 = c_2$, $\mu_1 = \mu_2$) and the same efficiency parameters ($\eta_1 = \eta_2$), this policy results in a download cost tradeoff governed by α. This tradeoff is represented in Figure 2.

There are at least two ways to achieve service differentiation. The first one is to guarantee a subscribed download cost for one class (e.g. $\phi_1 = \Phi$ for peers of class 1) with no constraint on the download cost of the other class. This can be done by assigning to the parameter α the (unique) root in $[0, 1]$ of the linear mapping $\alpha \to \phi_1 - \Phi$, where ϕ_1 is given in Proposition 3.2.

The second one is to achieve a target download cost ratio k between first- and second-class peers, namely

$$\frac{\phi_2}{\phi_1} = k. \tag{17}$$

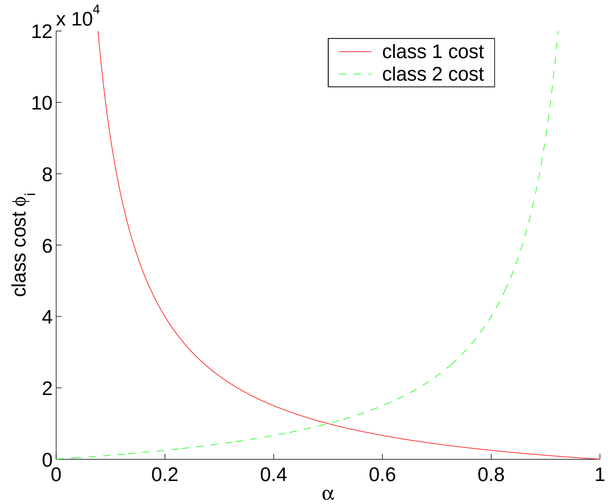

Figure 2: Download cost tradeoff ($\lambda_1 = \lambda_2 = 10^{-1}$, $\theta_1 = \theta_2 = \beta = 10^{-4}$, $c = 10^{-3}$)

The parameter α is then obtained as the (unique) root in $[0, 1]$ of the (either linear or quadratic) mapping $\alpha \to \phi_2/\phi_1 - k$. For a given set of parameters (see caption), Figure 3 reports the value of α that satisfies (17) as a function of k, for $k \in [1, 300]$.

We conclude that service differentiation in BitTorrent-like networks can easily be achieved through the single parameter α.

What if users stay connected after the download?

All the results obtained so far in this section have been derived under the assumption that there are no seeds in the system. As already observed this case can be seen as a worst-case scenario, where peers are selfish and leave the system as soon as they have downloaded the file.

In this section, we relax the no-seed assumption. In other words, we assume that downloaders do not leave the system immediately after they have downloaded the file, but continue to upload chunks to the other peers for some time of average duration $1/\gamma_i > 0$ for class-i peers.

In this more general setting the time-evolution of the system is given by the system of

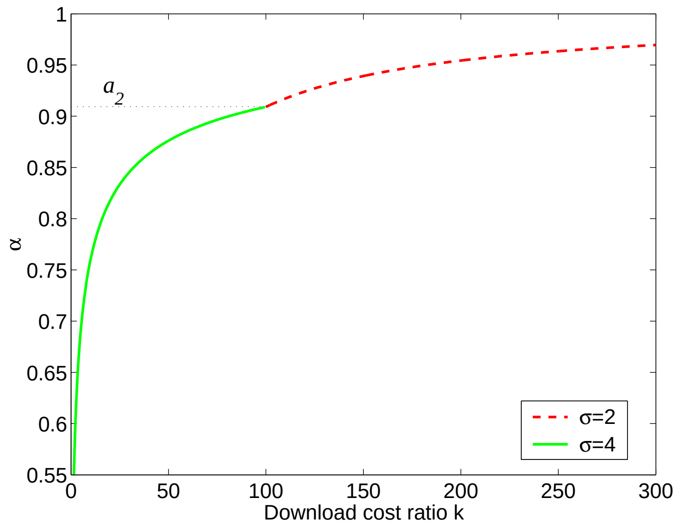

Figure 3: Selection of α for a target cost ratio k ($\lambda_1 = \lambda_2 = 10^{-1}_{,}$ $\theta_1 = \theta_2 = \beta = 10^{-4}_{,}$ $c = 10^{-3}$)

differential equations (1)-(2), with $(\alpha_1, \alpha_2) = (\alpha, 1-\alpha)$. We still assume that $\mu_1 = \mu_2$, $c_1 = c_2$ and $\eta_1 = \eta_2$ (these assumptions could be relaxed). The analysis of this system is much more complex than that of the no-seed model. While it is still easy to compute the stationary solutions of (1)-(2) in explicit form, it is much more complex to study the existence and stability of these solutions. However, there is no difficulty to numerically compute the steady-state of these equations once numerical values have been assigned to the system parameters.

This has been done for the following set of parameters: $\lambda_1 = \lambda_2 = 10^{-1}$ peers/s, $\theta_1 = \theta_2 = \mu = 10^{-4}s^{-1}, c = 10^{-3}s^{-1}, \eta_1 = \eta_2 = 0.9$. These parameters are rounded values of typical values estimated using the statistics in [3] in particular. We also assumed $\gamma_1 = \gamma_2 = \gamma$.

For given values of γ and $\alpha \in (0,1)$ we have computed the ratio of download costs $R = \phi_2/\phi_1$ for the seed model and the ratio of download costs $r = \phi_2/\phi_1$ for the no-seed model.

We have found that for $\gamma = c$, the relative error $|R - r|/R$ averages 1%. For $\gamma \geq c$, this relative error rapidly decreases, making the no-seed model very-well suited for the service differentiation problem. For $\gamma < c$, the relative error rapidly increases, making a numerical estimation of α necessary, using(1)-(2).

4 Bandwidth Diversity

We now address the bandwidth diversity problem for the no-seed model ($1/\gamma_i = 0$ for $i = 1, 2$). We consider two classes of peers with different bandwidths (e.g., ADSL users and corporate users). Recall that the bandwidth diversity problem we consider is characterized by $\alpha_1 = \alpha_2 = \alpha$ (see Section 2).

Our first objective is to determine the download cost for each class of peers. Then, we will find a static allocation policy (α, α) that minimizes the maximum download cost of both classes. With a slight abuse of notation a static allocation policy (α, α) will simply be referred to as an allocation α from now on.

Under the aforementioned assumptions the system of differential equations (3) becomes

$$\begin{aligned} \tfrac{dx_1}{dt} &= \lambda_1 - \theta_1 x_1 - \min\left(c_1 x_1,\ \alpha\beta_1 x_1 + (1-\alpha)\beta_2 x_2\right) \\ \tfrac{dx_2}{dt} &= \lambda_2 - \theta_2 x_2 - \min\left(c_2 x_2,\ (1-\alpha)\beta_1 x_1 + \alpha\beta_2 x_2\right). \end{aligned} \tag{18}$$

In matrix notation the system (18) is given by (6), with the switching condition

$$\sigma(\mathbf{x}) = 1 + 2 \times \mathbf{1}(cx_1 \geq \alpha\beta_1 x_1 + (1-\alpha)\beta_2 x_2) + \mathbf{1}(cx_2 \geq (1-\alpha)\beta_1 x_1 + \alpha\beta_2 x_2). \tag{19}$$

For the sake of compactness we introduce the new parameters

$$a_3 := \frac{\lambda_2\beta_2(\theta_1 + c_1) - \lambda_1(c_1\theta_2 + \beta_1\beta_2)}{\lambda_2\beta_2(\theta_1 + c_1) - \lambda_1(\beta_1\theta_2 + 2\beta_1\beta_2 - c_1\beta_2)} \quad a_4 := \frac{\lambda_1\beta_1(\theta_2 + c_2) - c_2\lambda_2(\theta_1 + c_1)}{\lambda_1\beta_1(\theta_2 + c_2) - \beta_2\lambda_2(\theta_1 + c_1)} \tag{20a}$$

$$a_5 := \frac{\lambda_1\beta_1(\theta_2 + c_2) - \lambda_2(c_2\theta_1 + \beta_2\beta_1)}{\lambda_1\beta_1(\theta_2 + c_2) - \lambda_2(\beta_2\theta_1 + 1\beta_2\beta_1 - c_2\beta_1)} \quad a_6 := \frac{\lambda_2\beta_2(\theta_1 + c_1) - c_1\lambda_1(\theta_2 + c_2)}{\lambda_2\beta_2(\theta_1 + c_1) - \beta_1\lambda_1(\theta_2 + c_2)} \tag{20b}$$

$$d := (\theta_2 + \alpha\beta_2)(\theta_1 + \alpha\beta_1) - (1-\alpha)^2\beta_1\beta_2. \tag{20c}$$

We also define the elementary conditions

$$(\mathcal{C}1): \quad \lambda_1(c_1\theta_2 + \beta_1\beta_2) \leq \lambda_2\beta_2(\theta_1 + c_1) \quad \text{and} \quad 0 \leq \alpha < a_3 \tag{21}$$

$$(\mathcal{C}2): \quad c_2\lambda_2(\theta_1 + c_1) \geq \lambda_1\beta_1(\theta_2 + c_2) \quad \text{or} \quad a_4 \leq \alpha \leq 1 \tag{22}$$

$$(\mathcal{C}3): \quad \lambda_2(c_2\theta_1 + \beta_2\beta_1) \leq \lambda_1\beta_1(\theta_2 + c_2) \quad \text{and} \quad 0 \leq \alpha < a_5 \tag{23}$$

$$(\mathcal{C}4): \quad c_1\lambda_1(\theta_2 + c_2) \geq \lambda_2\beta_2(\theta_1 + c_1) \quad \text{or} \quad a_6 \leq \alpha \leq 1. \tag{24}$$

Furthermore, let us define the following set of conditions

$$\begin{aligned}(\mathcal{D}2) &= (\mathcal{C}1) \cup (\mathcal{C}2)\\(\mathcal{D}3) &= (\mathcal{C}3) \cup (\mathcal{C}4)\\(\mathcal{D}4) &= (\text{not } (\mathcal{C}1)) \cup (\text{ not } (\mathcal{C}3)).\end{aligned}$$

The above definitions imply that $(\mathcal{D}4) \cap (\mathcal{D}2) = (\mathcal{D}4) \cap (\mathcal{D}3) = \emptyset$, where $\emptyset$ denotes the empty set. However, $(\mathcal{D}2) \cap (\mathcal{D}3)$ is not necessarily empty, so that $(\mathcal{D}2)$ and $(\mathcal{D}3)$ may hold simultaneously for some sets of parameters. Finally, we define the two-dimensional vectors $\mathbf{x}_i$, $i = 2, 3, 4$, by

$$\begin{aligned}\mathbf{x}_2 &= \left(\frac{\lambda_1}{c_1+\theta_1}, \frac{\lambda_2 - (1-\alpha)\beta_1 \frac{\lambda_1}{c_1+\theta_1}}{\theta_2 + \alpha\beta_2}\right)\\ \mathbf{x}_3 &= \left(\frac{\lambda_1 - (1-\alpha)\beta_2 \frac{\lambda_2}{c_2+\theta_2}}{\theta_1 + \alpha\beta_1}, \frac{\lambda_2}{\theta_2 + c_2}\right)\\ \mathbf{x}_4 &= \left(\frac{\lambda_1(\theta_2+\alpha\beta_2) - (1-\alpha)\lambda_2\beta_2}{d}, \frac{\lambda_2(\theta_1+\alpha\beta_1) - (1-\alpha)\lambda_1\beta_1}{d}\right)\end{aligned}$$

where d is defined in (20c). Proposition 4.1 below investigates the steady-state behavior of the switched system (18).

Proposition 4.1 (Equilibrium for bandwidth diversity) *The system of differential equations (18) has a unique equilibrium point $\bar{\mathbf{x}}$ given by*

$$\bar{\mathbf{x}} = \begin{cases} \mathbf{x}_2^T & \textit{regardless of } \mathbf{x}(0), \ \textit{if } (\mathcal{D}2) \textit{ holds and } (\mathcal{D}3) \textit{ does not hold}\\ \mathbf{x}_3^T & \textit{regardless of } \mathbf{x}(0), \ \textit{if } (\mathcal{D}3) \textit{ holds and } (\mathcal{D}2) \textit{ does not hold}\\ \mathbf{x}_4^T & \textit{regardless of } \mathbf{x}(0), \ \textit{if } (\mathcal{D}4) \textit{ holds}\\ \mathbf{x}_2^T \textit{ or } \mathbf{x}_3^T & \textit{depending on } \mathbf{x}(0), \ \textit{if } (\mathcal{D}2) \textit{ and } (\mathcal{D}3) \textit{ hold simultaneously.}\end{cases} \tag{25}$$

$\diamond$

Proof. As in Section 3, we first assume that $\lim_{t\to\infty} \mathbf{x}(t)$ exists and check that it is given by (25). Letting $t \to \infty$ in (6) yields (12), where σ is now given by (19).

We consider separately the four systems of linear equations obtained from (12) when (a) $\sigma(\bar{\mathbf{x}}) = 1$, (b) $\sigma(\bar{\mathbf{x}}) = 2$, (c) $\sigma(\bar{\mathbf{x}}) = 3$ and (d) $\sigma(\bar{\mathbf{x}}) = 4$.

(a) $\sigma(\bar{\mathbf{x}}) = 1$ or equivalently $c_1\bar{x}_1 < \alpha\beta_1\bar{x}_1 + (1-\alpha)\beta_2\bar{x}_2$ and $c_2\bar{x}_2 < (1-\alpha)\beta_1\bar{x}_1 + \alpha\beta_2\bar{x}_2$.

The download rate is the bottleneck for both classes of peers. We find

$$\bar{\mathbf{x}}^T = \left(\frac{\lambda_1}{\theta_1 + c_1}, \frac{\lambda_2}{\theta_2 + c_2}\right). \tag{26}$$

(b) $\sigma(\bar{\mathbf{x}}) = 2$ or equivalently $c_1\bar{x}_1 < \alpha\beta_1\bar{x}_1 + (1-\alpha)\beta_2\bar{x}_2$ and $c_2\bar{x}_2 \geq (1-\alpha)\beta_1\bar{x}_1 + \alpha\beta_2\bar{x}_2$.

The bottleneck is the download rate for class-1 peers and the upload rate for class-2 peers. We find

$$\bar{\mathbf{x}}^T = \left(\frac{\lambda_1}{\theta_1 + c_1}, \frac{\lambda_2 - (1-\alpha)\beta_1 \frac{\lambda_1}{c_1+\theta_1}}{\theta_2 + \alpha\beta_2}\right). \tag{27}$$

(c) $\sigma(\bar{\mathbf{x}}) = 3$ or equivalently $c_1\bar{x}_1 \geq \alpha\beta_1\bar{x}_1 + (1-\alpha)\beta_2\bar{x}_2$ and $c_2\bar{x}_2 < (1-\alpha)\beta_1\bar{x}_1 + \alpha\beta_2\bar{x}_2$.

The bottleneck is the download rate for peers of class 2 and the upload rate for peers of class 1. In this case

$$\bar{\mathbf{x}}^T = \left(\frac{\lambda_1 - (1-\alpha)\beta_2 \frac{\lambda_2}{c_2+\theta_2}}{\theta_1 + \alpha\beta_1}, \frac{\lambda_2}{\theta_2 + c_2}\right). \tag{28}$$

(d) $\sigma(\bar{\mathbf{x}}) = 4$ or equivalently $c_1\bar{x}_1 \geq \alpha\beta_1\bar{x}_1 + (1-\alpha)\beta_2\bar{x}_2$ and $c_2\bar{x}_2 \geq (1-\alpha)\beta_1\bar{x}_1 + \alpha\beta_2\bar{x}_2$.

The bottleneck is the download rate for both classes of peers. The stationary solution is

$$\bar{\mathbf{x}}^T = \left(\frac{\lambda_1(\theta_2 + \alpha\beta_2) - (1-\alpha)\lambda_2\beta_2}{d}, \frac{\lambda_2(\theta_1 + \alpha\beta_1) - (1-\alpha)\lambda_1\beta_1}{d}\right) \tag{29}$$

where d is defined in (20c).

The next step is to check if a type-i equilibrium may exist, namely, if $\sigma(\bar{\mathbf{x}}) = 1$ (resp. $\sigma(\bar{\mathbf{x}}) = 2$, $\sigma(\bar{\mathbf{x}}) = 3$, $\sigma(\bar{\mathbf{x}}) = 4$) when $\bar{\mathbf{x}}$ is given by (26) (resp. (27), (28), (29)).

It is easily seen that a type-1 equilibrium may only exist if $c_1\lambda_1 + c_2\lambda_2 \leq \beta_1\lambda_1 + \beta_2\lambda_2$. Since $c_i > \beta_i$ for $i = 1, 2$ (see (5)) we conclude that there is no type-1 equilibrium, where both classes would saturate their download capacity.

A simple analysis, similar to that in Appendix A, shows that a type-2 equilibrium only exists if (21) and (22) are true, and that a type-3 equilibrium only exists if (23) and (24) are true. For the existence conditions of a type-4 equilibrium, we also use the stability condition (30) below, in addition to $\sigma(\bar{\mathbf{x}}) = 4$, to get the following condition: (not (21)) and (not (23)). It happens that conditions for $\sigma = 2$ and $\sigma = 3$ are not mutually exclusive. When they are simultaneously satisfied (i.e., $(\mathcal{D}2) \cap (\mathcal{D}3)$ holds) then the steady-state is given either by (27) or (28) depending on the initial conditions.

We now turn to the proof that $\lim_{t\to\infty} \bar{\mathbf{x}}(t)$ exists. Namely, we investigate the nature of the equilibrium of each of the linear systems $\dot{\mathbf{x}}(t) = A_i\, \mathbf{x}(t) + \mathbf{b}$, for $i = 2, 3, 4$, with

$$A_2 = \begin{pmatrix} -\theta_1 - c_1 & 0 \\ -(1-\alpha)\beta_1 & -\theta_2 - \alpha\beta_2 \end{pmatrix} \quad A_3 = \begin{pmatrix} -\theta_1 - \alpha\beta_1 & -(1-\alpha)\beta_2 \\ 0 & -\theta_2 - c_2 \end{pmatrix}$$

and

$$A_4 = \begin{pmatrix} -\theta_1 - \alpha\beta_1 & -(1-\alpha)\beta_2 \\ -(1-\alpha)\beta_1 & -\theta_2 - \alpha\beta_2 \end{pmatrix}.$$

It is easily seen that the matrices A_2 and A_3 have two strictly negative eigenvalues. The eigenvalues of the matrix A_4 are the roots in λ of the polynomial

$$\begin{aligned} det(A_4 - \lambda I) &= (\theta_1 + \alpha\beta_1 + \lambda)(\theta_2 + \alpha\beta_2 + \lambda) - (1-\alpha)^2\beta_1\beta_2 \\ &= \lambda^2 + \lambda(\theta_1 + \alpha\beta_1 + \theta_2 + \alpha\beta_2) + d \end{aligned}$$

where I denotes the 2×2 identity matrix. The roots of this polynomial have strictly negative real parts if and only if their product is strictly positive and their sum is strictly negative, which is equivalent to

$$d > 0. \tag{30}$$

This shows that all equilibria are stable, which concludes the proof. ∎

We now compute the download costs ϕ_1 and ϕ_2 associated with each equilibrium point found in Proposition 4.1.

In order to simplify the notation, we introduce the following two-dimensional vectors

$$\begin{aligned} \varphi_2 &= \left(\frac{1}{c_1}, \quad \frac{\lambda_2(\theta_1 + c_1) - (1-\alpha)\lambda_1\beta_1}{\theta_2\lambda_1\beta_1 + \alpha(\lambda_2\beta_2(\theta_1 + c_1) - \lambda_1\beta_1\theta_2)}\right) \\ \varphi_3 &= \left(\frac{\lambda_1(\theta_2 + c_2) - (1-\alpha)\lambda_2\beta_2}{\theta_1\lambda_2\beta_2 + \alpha(\lambda_1\beta_1(\theta_2 + c_2) - \theta_1\lambda_2\beta_2)}, \quad \frac{1}{c_2}\right) \\ \varphi_4 &= \left(\frac{(\lambda_1\theta_2 - \lambda_2\beta_2 + \alpha\beta_2(\lambda_1 + \lambda_2))}{\beta_2(\lambda_2\theta_1 - \lambda_1\beta_1) + \alpha(\lambda_1\beta_1(\theta_2 + 2\beta_2) - \theta_1\lambda_2\beta_2)}, \right. \\ &\qquad \left. \frac{\lambda_2\theta_1 - \lambda_1\beta_1 + \alpha\beta_1(\lambda_1 + \lambda_2)}{\beta_1(\lambda_1\theta_2 - \lambda_2\beta_2) + \alpha(\lambda_2\beta_2(\theta_1 + 2\beta_1) - \theta_2\lambda_1\beta_1)}\right). \end{aligned}$$

The next proposition partially characterizes the download costs ϕ_1 and ϕ_2.

Proposition 4.2 (Download costs for bandwidth diversity) *In a no-seed model, the vector of download costs* (ϕ_1, ϕ_2) *in the bandwidth diversity problem is given by*

$$(\phi_1, \phi_2) = \begin{cases} \varphi_2 & \textit{regardless of } \mathbf{x}(0), \textit{ if } (\mathcal{D}2) \textit{ holds and } (\mathcal{D}3) \textit{ does not hold} \\ \varphi_3 & \textit{regardless of } \mathbf{x}(0), \textit{ if } (\mathcal{D}3) \textit{ holds and } (\mathcal{D}2) \textit{ does not hold} \\ \varphi_4 & \textit{regardless of } \mathbf{x}(0), \textit{ if } (\mathcal{D}4) \textit{ holds} \\ \varphi_2 \textit{ or } \varphi_3 & \textit{depending on } \mathbf{x}(0), \textit{ if } (\mathcal{D}2) \textit{ and } (\mathcal{D}3) \textit{ hold simultaneously.} \end{cases}$$

◇

Proof. The proof directly follows from Proposition 4.1 and (8). ∎

How can we minimize the highest download cost?

In the bandwidth diversity problem, several optimization problems could be considered. For instance, one may wish to find an allocation α that yields the same download costs. Another objective could be to minimize a linear combination of the download costs. However, as shown in Proposition 4.2, it is difficult to analytically determine ϕ_1 and ϕ_2 whenever $(\mathcal{D}2) \cap (\mathcal{D}3) \neq \emptyset$, and thereby to solve the above optimization problems.

Instead, we will seek to minimize the maximum download cost over all initial states and over all classes. To this end, we introduce the mapping $\alpha \to E(\alpha)$, called the *envelope function*, defined by

$$E(\alpha) = \max_{\sigma \in \{2,3,4\}} \max_{i \in \{1,2\}} \phi_i.$$

Our objective is to minimize the envelope function as a function of α.

We now use Proposition 4.2 to calculate the value of α that minimizes $E(\alpha)$. In Figures 4 and 5, the envelope function is represented along with the possible values of (ϕ_1, ϕ_2) for α in $(0, 1)$, for two different set of physical parameters.

In Figure 4, we observe that $E(\alpha)$ is minimal for a single value of α, when $\sigma = 4$ and $\phi_1 = \phi_2$. In this case, the exact value of α that minimizes the maximum download cost can be found by solving $\phi_1 = \phi_2$ using Proposition 4.2. Note that in Figure 4, we have both type-2 and 3 equilibria for $\alpha \leq 0.32$. The steady-state is then determined by initial conditions.

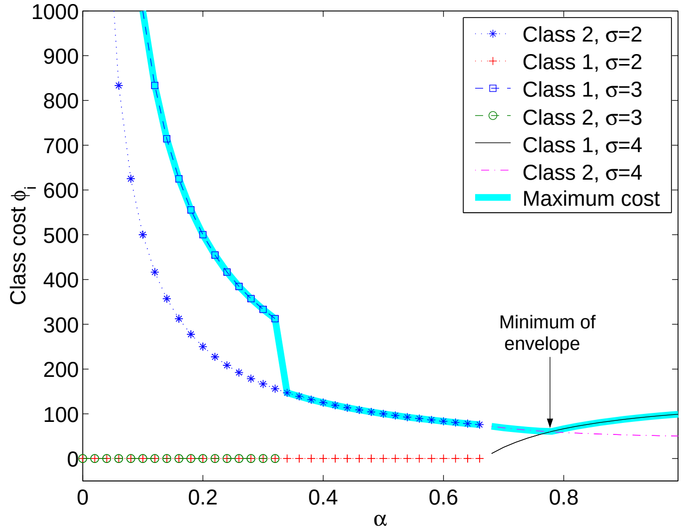

Figure 4: Minimum of maximum download cost achieved for $\alpha \approx 0.78$. $(\lambda_1 = \frac{\lambda_2}{2} = 10, \beta_1 = \frac{\beta_2}{2} = 10^{-2}, c_1 = \frac{c_2}{2} = 400, \theta_1 = \theta_2 = 10^{-5})$

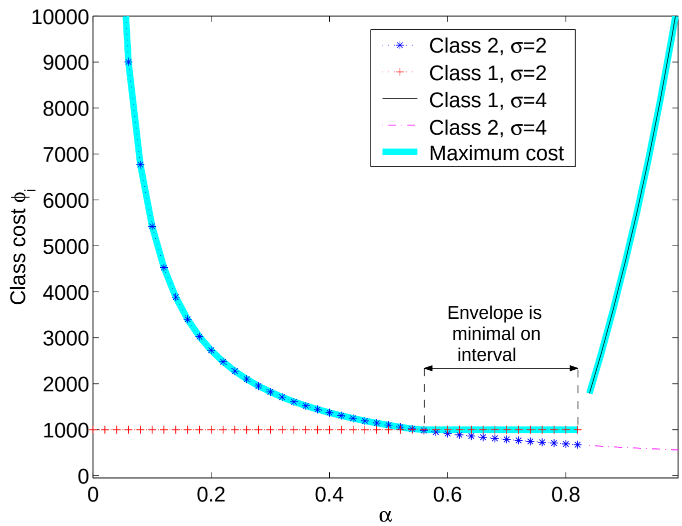

Figure 5: Minimum of maximum download cost achieved for a whole interval $[0.5502, 0.8207]$. $(\lambda_1 = \frac{\lambda_2}{4} = 10^{-1}, \theta_1 = 2\theta_2 = \beta_1 = \frac{\beta_2}{20} = 10^{-4}, c_1 = \frac{c_2}{2} = 10^{-3})$

In Figure 5, $E(\alpha)$ is minimal on a whole interval on which it is equal to the constant download cost ϕ_1 when $\sigma = 2$. In this case, the interval can also be determined using Proposition 4.2, by solving $\phi_1 = \phi_2$ for $\sigma = 2$ for the lower bound, and by determining the maximum value of α that satisfies (21) and (22) for the upper bound. Note that in Figure 5, condition (23) and (24) are never satisfied simultaneously with this set of physical parameters, since we do not have a type-3 equilibrium.

In any case, finding the value of α that minimizes the worst download cost, amounts to solve a linear or quadratic equation $\phi_1 = \phi_2$ using the appropriate expression in Proposition 4.2.

We conclude that for a given physical set of parameters, it is possible to account for bandwidth diversity in BitTorrent-like networks through parameter α.

5 Conclusions and Perspectives

In this paper we presented a simple multiclass fluid model for BitTorrent-like distribution systems. We successfully applied this model to account for two specific problems: service differentiation and bandwidth diversity. We mainly focused our attention to the special case where peers selfishly leave the system immediately after their download ("no-seed case"). For both the service differentiation and bandwidth diversity problems, we have defined a single parameter α that defines a resource allocation strategy. We showed how this parameter affects the steady-state of the system and provided closed-form expressions for the successful download time in each case. In addition, we showed how this parameter α can be chosen so as to achieve a target quality of service ratio (download time ratio) for the service differentiation problem. We also quantified the impact of the no-seed assumption on this result through a numerical resolution of the general problem. For the bandwidth diversity problem, we also showed how it is possible to choose parameter α so as to minimize the highest download time between two classes of peers.

Many open problems remain. In particular, we intend to compare the results of our model to a simulation of a real P2P file dissemination system. Another problem for further research is the study of dynamic resource allocation, where α would depend on the class population.

A Service Differentiation: Type-2 Equilibrium

In this appendix we show that a type-2 equilibrium exists for $\alpha \in [0,1]$ if and only if $a_2 < \alpha \leq 1$, where a_2 is defined in Section 3.

By definition, a type-2 equilibrium exists if $\bar{\mathbf{x}} = (\bar{x}_1, \bar{x}_2)$ given in (14) is such that $\sigma(\bar{\mathbf{x}}) = 2$, to which we need to add the condition that $\bar{x}_2 \geq 0$ (note that $\bar{x}_1$ is always nonnegative). Equivalently, we need to find the values of α in $[0,1]$ such that

$$\begin{aligned} c\xi &< \alpha\beta\left(\xi + \frac{\lambda_2 - (1-\alpha)\beta\xi}{\theta_2 + (1-\alpha)\beta}\right) \\ c\frac{\lambda_2 - (1-\alpha)\beta\xi}{\theta_2 + (1-\alpha)\beta} &\geq (1-\alpha)\beta\left(\xi + \frac{\lambda_2 - (1-\alpha)\beta\xi}{\theta_2 + (1-\alpha)\beta}\right) \\ \lambda_2 - (1-\alpha)\beta\xi &\geq 0 \end{aligned}$$

where we have set $\xi := \lambda_1/(\theta_1 + c)$. The first two conditions express the identity $\sigma(\bar{\mathbf{x}}) = 2$ and the third condition expresses the constraint $\bar{x}_2 \geq 0$.

Straightforward algebra shows that these conditions are simultaneously met for $\alpha \in [0,1]$ if and only if

$$\alpha > \frac{c\lambda_1(\theta_2 + \beta)}{D} \quad \text{and} \quad \alpha \geq \max\left(1 - \frac{c\lambda_2(\theta_1 + c)}{D}, 1 - \frac{\lambda_2(\theta_1 + c)}{\lambda_1\beta}\right) \tag{31}$$

where we recall that $D = \beta(\lambda_1(\theta_2 + c) + \lambda_2(\theta_1 + c))$.

Let us first compare $c\lambda_1(\theta_2 + \beta))/D$ to $1 - c\lambda_2(\theta_1 + c))/D$. We have

$$\begin{aligned} \frac{c\lambda_1(\theta_2 + \beta)}{D} - \left(1 - \frac{c\lambda_2(\theta_1 + c)}{D}\right) &= \frac{1}{D}\left(c\lambda_1(\theta_2 + \beta) + c\lambda_2(\theta_1 + c) - D\right) \\ &= \frac{1}{D}\left(c\left(\lambda_1(\theta_2 + \beta) + \lambda_2(\theta_1 + c)\right) - \beta(\lambda_1(\theta_2 + c) + \lambda_2(\theta_1 + c))\right) \\ &= \frac{1}{D}(c - \beta)(\lambda_1\theta_2 + \lambda_2(\theta_1 + c)). \end{aligned}$$

We have observed earlier in the proof of Proposition 3.1 that $c > \beta$, which shows that $c\lambda_1(\theta_2 + \beta))/D > 1 - c\lambda_2(\theta_1 + c))/D$.

We now compare $c\lambda_1(\theta_2+\beta))/D$ to $1-\lambda_2(\theta_1+c)/(\lambda_1\beta)$. We have

$$\begin{aligned}
&\frac{c\lambda_1(\theta_2+\beta)}{D} - \left(1-\frac{\lambda_2(\theta_1+c)}{\lambda_1\beta}\right) \\
&= \frac{1}{D\lambda_1\beta}\left(c\lambda_1^2\beta(\theta_2+\beta) - D\lambda_1\beta + \lambda_2(\theta_1+c)D\right) \\
&= \frac{1}{D\lambda_1\beta}\big(\lambda_1\beta(c\lambda_1(\theta_2+\beta) - \beta\lambda_1(\theta_2+c) - \beta\lambda_2(\theta_1+c) + \lambda_2(\theta_1+c)(\theta_2+c)) \\
&\quad +\beta\lambda_2^2(\theta_1+c)^2\big) \\
&= \frac{1}{D\lambda_1}\left(\lambda_1(\lambda_1\theta_2(c-\beta) + \lambda_2(\theta_1+c)(\theta_2+c-\beta)) + \lambda_2^2(\theta_1+c)^2\right) > 0
\end{aligned}$$

since $c > \beta$.

In summary we have shown that the conditions $\sigma(\bar{\mathbf{x}}) = 2$ and $\bar{x}_2 \geq 0$ will simultaneously hold for $\alpha \in [0,1]$ if and only if $\alpha > \min(1, c\lambda_1(\theta_2+\beta))/D) = a_2$, which is the announced result. ∎

References

[1] Cohen, B. BitTorrent. http://www.bitconjurer.org/BitTorrent/.

[2] Horn, R. A., and Johnson, C. R. *Matrix Analysis.* Cambridge University Press, 1985.

[3] Izal, M., Urvoy-Keller, G., Biersack, E., Felber, P., Hamra, A. A., and Garcés-Erice, L. Dissecting bittorrent: Five months in a torrent's lifetime. In *Proceedings of Passive and Active Measurement Workshop (PAM2004)* (Antibes, France, April 2004).

[4] Khalil, H. K. *Nonlinear systems.* MacMillan, 1992.

[5] Liberzon, D. *Switching in systems and control.* Birkhäuser, 2003.

[6] Morse, A. S., Ed. *Control Using Logic-Based Switching.* London: Springer-Verlag, 1997.

[7] Qiu, D., and Srikant, R. Modeling and performance analysis of bittorrent-like peer-to-peer networks. In *Proceedings of ACM Sigcomm* (Portland, OR, Aug 2004).