

Dynamics of content propagation in BitTorrent-like P2P file exchange systems

Andrés Ferragut, Fabián Kozynski and Fernando Paganini
Universidad ORT Uruguay

Abstract—Peer to peer file exchange systems such as BitTorrent are changing the way in which content is distributed in the Internet. Service capacity for a certain content adjusts dynamically as a function of peer population, thus achieving scalability. This dynamic behavior has been the subject of recent analytical studies.

In this paper, we propose a partial differential equation model for BitTorrent-like systems, in which a fluid variable represents content and its distribution in the system is taken into account. This model allows for a variety of file sharing disciplines; we identify equilibrium properties that must hold regardless of this choice, and others that depend on a notion of efficiency. The equilibrium properties of many specific disciplines are described. Through a discretized ordinary differential equation model, we also present results on the stability of the equilibrium for a particular sharing policy that is suitable to model current BitTorrent systems.

I. INTRODUCTION

A large proportion of Internet traffic in recent years is attributed to *peer-to-peer* (P2P) systems. These applications move away from the traditional client-server model for data exchange, replacing it with a system in which every peer is both a client and a server, sharing its own upload capacity to enhance the collective download performance. The main advantage of this approach is scalability: as demand for a certain content becomes large, so does the available supply. One of the dominant P2P file sharing systems is BitTorrent [1], where a swarm of peers is dynamically formed in relation to a particular content of interest. Peers are classified in *seeders* who own the full content, and *leechers* who own part of it, and exchange rules are set up to elicit collaborative behavior among them. Some basic elements of BitTorrent are reviewed in Section II.

This new file-sharing paradigm has raised interesting modeling questions. In particular, the dynamics of the swarm population has non-trivial properties, given that service times depend on network capacity, which itself depends on the population size. A first model in this regard was proposed by [8]: here, populations of leechers and seeders are described by a continuous time Markov chain, and its equilibrium distribution is studied numerically. In [5], the authors proposed a differential equation model for this system, which corresponds to a fluid limit of the previous model, with some additional features. This allows for an analytical description of the resulting equilibrium and its stability, as we review

in Section II. The above models lump leechers into a single quantity, irrespective of how much content they possess. In [6], a Markov model that includes the degree of advance is considered, with a specific choice for the exchange rates among peers; again, numerical studies are performed.

In this paper we generalize the idea in [6] to provide a finer characterization of how content propagates in a P2P network. In particular, by taking a fluid approach in the content variable, we will express in Section III the population dynamics in terms of a partial differential equation, which covers BitTorrent-like systems under very general assumptions on how upload capacity is shared among peers. We identify certain equilibrium properties that are independent of this choice, as well as others that are contingent on a suitably defined notion of efficiency. In Section IV we analyze some specific sharing disciplines and show how they include and generalize the previously cited models.

In Section V we turn to dynamic studies, focusing on the sharing model that is most similar to current BitTorrent practice. We give results on local stability for its equilibrium, using an ordinary differential equation discretization of the previous dynamics. Conclusions are presented in Section VI.

II. BACKGROUND AND RELATED WORK

Consider a set of Internet users who have common interest in a certain *content* (i.e. a file or a set of files). In a BitTorrent-like system, such content is organized and subdivided in smaller size pieces named *chunks*, to be exchanged among peers. An incoming peer first obtains from a *tracker* information that identifies a set of peers that have the whole or part of the content; from here on, a series of control messages exchange a *bitfield* with detailed information on which chunks are in possession of each of these other peers. The new peer then becomes a *leecher*, requesting chunks from others; at the same time, as soon as the user has some pieces of the file, other peers can request them from it. At some point, the peer completes the download and becomes a *seeder*, who is no longer interested in downloading content, but contributes uploading to the remaining leeches. This goes on until the peer decides to leave the system.

A first dynamic model for such a P2P network was given in [8]. This Markov chain model has as state variables the populations of leechers and seeders in the system, behaving as follows: peers arrive as a Poisson process of rate λ , stay in a leecher queue until completing the download, and then in a seeder queue for an exponential time of parameter γ . The leecher queue is served in a processor sharing discipline

This work was supported by AFOSR-US under grant FA9550-09-1-0504, and ANII-Uruguay, grants PR_FCE_2009_1_2158 and BE_INI_2010_2106.
E-mail: ferragut@ort.edu.uy.

by those uploading content, the mean upload time for the file being $\frac{1}{\mu}$. Taking the latter time to be exponential leads to a continuous time Markov chain, which is then analyzed numerically.

As an initial observation regarding this model, we distinguish two cases that will have impact in what follows:

- (i) If $\gamma < \mu$, then the mean time $\frac{1}{\gamma}$ that a peer spends as a seeder suffices (with excess) to upload one copy of the file, thus generating a replacement seeder. In this situation, which we will call *seeder-sustained*, the P2P system could reproduce the content even without the upload contribution of leechers.
- (ii) If $\gamma > \mu$, then the opposite holds: seeders leave quickly, so the upload contribution of leechers is essential to reach an equilibrium. We call this regime *globally-sustained*; here is where the power of P2P sharing is most significant.

In [5], the authors proposed a differential equation model for this system, which corresponds to a fluid limit of the Markov model in [8], with some additional features. The main one is to make explicit the download capacity limit which is necessary to have a bounded equilibrium in the seeder-sustained case. We now recall this model, with some new notation of our own.

Let $x(t)$ denote the number of leechers at time t , and $y(t)$ the number of seeders. Therefore, the total upload capacity in the system in files per second becomes¹

$$\bar{\Omega}_{up} = \mu(y + x). \quad (1)$$

Let Ω_{down} denote the total download rate of the system, in files per second. This determines the rate at which leechers turn into seeders², so the overall dynamics, including leech arrivals and seed departures is:

$$\dot{x} = \lambda - \Omega_{down}(x, y), \quad (2a)$$

$$\dot{y} = \Omega_{down}(x, y) - \gamma y. \quad (2b)$$

To complete the model, the expression

$$\Omega_{down}(x, y) = \min\{\bar{\Omega}_{up}, cx\} \quad (3)$$

states that Ω_{down} is constrained (only) by either the available upload capacity or the maximum download rate per peer, defined as c . We assume that $c > \mu$ (download capacity exceeds upload capacity), common in practice.

It is easily checked that at an equilibrium of (2) the number of seeders must be $y^* = \frac{\lambda}{\gamma}$. As for the number of leechers, there are two cases distinguished by a critical value of γ , defined as $\gamma_{cr} := \mu \frac{c}{c-\mu}$.

- (i') If $\gamma < \gamma_{cr}$, the equilibrium point satisfies

$$x^* = \frac{\lambda}{c}, \quad \Omega_{down}(x^*, y^*) = cx^* < \bar{\Omega}_{up},$$

¹In [5], [8] the leecher contribution is weighed by an efficiency parameter $\eta < 1$; since subsequent analysis in [5] shows $\eta \approx 1$ in nearly all practical situations, we will not include it here.

²For simplicity, we assume here that leechers do not leave the system before completing their download.

so here the bottleneck of the system is the download capacity.

- (ii') If $\gamma > \gamma_{cr}$, the equilibrium point is

$$x^* = \lambda \left(\frac{1}{\mu} - \frac{1}{\gamma} \right), \quad \Omega_{down}(x^*, y^*) = \bar{\Omega}_{up} < cx^*,$$

so the bottleneck is the upload capacity.

With a slight modification, the above cases correspond respectively to the seeder-sustained (i) and globally-sustained (ii) cases mentioned before. Indeed, as $c \rightarrow \infty$, $\gamma_{cr} \rightarrow \mu$ and the identification becomes exact.

It is shown in [5] that equilibrium points in both cases are locally asymptotically stable; moreover, in [4] it is proved that the switched system defined by (1), (2), (3) is globally asymptotically stable for any choice of λ, μ, γ and $c > \mu$.

The main limitation of the above models is that the chosen state does not differentiate between peers who have different amounts of downloaded content; rather, the entire download rate is assigned indiscriminately to leecher termination. A more accurate model should describe the dynamics of a leecher gaining content, and how this degree of advance in turn influences the download rate, that need not be homogeneous among peers.

At the other extreme of detail would be a model that takes as state the complete bitfield of each peer, i.e. the list of chunks it possesses. This kind of model was analyzed in [2], [3], but becomes unmanageable as the number of chunks grows due to the state dimensionality.

An intermediate point is to consider the amount of content each peer possesses as an index for the state; a Markov model in this direction was studied by [6]. In what follows, we will propose and study a fluid model of this nature, that can be thought as a generalization of the models in [5], [6].

III. A FLUID MODEL FOR CONTENT DOWNLOAD IN BITTORRENT SYSTEMS

As mentioned before, the file of interest in a BitTorrent system is divided into small size chunks that are downloaded independently. We will propose now a fluid model in which the chunk size is considered infinitesimally small, and the download rate of a given peer may depend on the amount of file it has already obtained from the system.

To do so, let us denote by $F(t, \sigma)$ the number of leechers that have completed downloading less than σ units of the file. We let σ vary continuously between 0 and 1, with $\sigma = 1$ representing the entire file. $F(t, \sigma)$ acts therefore as a cumulative distribution of leechers in different stages of download, with $F(t, 0) = 0$ and $F(t, 1) = x(t)$, the total number of leechers.

Let $\omega(F, y, \sigma)$ denote the download speed of a single leecher, that depends on the current distribution F , the number of seeders y , and is also allowed to depend on σ , the amount of data already in the leecher's possession. The evolution of F after an infinitesimal amount of time dt verifies:

$$F(t + dt, \sigma) = F(t, \sigma) + \lambda dt - [F(t, \sigma) - F(t, \sigma - \omega dt)].$$

The preceding equation tells us that the number of peers with content less than σ after a time dt is incremented by the new arrivals λdt , and decremented by the number of peers that advance beyond σ in their download stage. The latter is given by those that at time t have content between $\sigma - \omega dt$ and σ , which explains the last term in the above equation.

By reordering terms, dividing by dt and letting $dt \rightarrow 0$ we have the following evolution equation for F , in the form of a transport partial differential equation:

$$\frac{\partial F}{\partial t} = \lambda - \omega(F, y, \sigma) \frac{\partial F}{\partial \sigma}. \quad (4)$$

As for the evolution of the seeders, the generalization of (2b) is:

$$\dot{y} = \omega(F, y, 1) \frac{\partial F}{\partial \sigma} \Big|_{\sigma=1} - \gamma y, \quad (5)$$

corresponding to the leechers finishing download minus the seeder departures.

To completely specify the model, we have to choose a particular $\omega(F, y, \sigma)$. This corresponds to choosing a way to distribute the available system capacity between leechers with different amounts of content. In particular, since the bandwidth is provided by the peers present in the system, a first constraint on ω is:

$$\Omega_{down}(t) := \int_0^1 \omega(F, y, \sigma) \frac{\partial F}{\partial \sigma} d\sigma \leq \bar{\Omega}_{up}(t), \quad (6)$$

where $\bar{\Omega}_{up}(t) = \mu(y(t) + F(t, 1))$ as before.

Secondly, we impose downlink capacity constraint for each user:

$$\omega(F, y, \sigma) \leq c \quad \forall \sigma. \quad (7)$$

We now analyze the model of (4), (5) with the above constraints (6), (7).

A. Evolution of the unfinished work

Let us define the following quantity, which represents the amount of data that must be provided to the leechers present for them to finish download.

Definition 1 (Unfinished work):

$$u(t) = \int_0^1 (1 - \sigma) \frac{\partial F}{\partial \sigma} d\sigma. \quad (8)$$

Integrating by parts in (8), we have an alternative way to express the amount of unfinished work as:

$$u(t) = \int_0^1 F(t, \sigma) d\sigma. \quad (9)$$

Assuming sufficient regularity in F and using (4) we can calculate the following evolution equation for u :

$$\begin{aligned} \dot{u} &= \frac{\partial}{\partial t} \int_0^1 F(t, \sigma) d\sigma = \int_0^1 \lambda - \omega(F, y, \sigma) \frac{\partial F}{\partial \sigma} d\sigma \\ &= \lambda - \Omega_{down}(t). \end{aligned} \quad (10)$$

The right hand side of (10) is similar to the one in equation (2a). In fact, we can interpret (2a) as an extreme case of our model where the torrent has only one chunk. In this case, the

number of leechers $x(t) = F(t, 1)$ of (2a) coincides with u , the amount of unfinished work. Similarly, (2b) can be seen as a specialization of (5) to the case of one chunk.

When the distribution of content is considered, however, the model becomes different; in particular, the seeder generation rate is determined only by the leechers who are about to complete their download.

B. Equilibrium properties

We now turn our attention to the equilibrium properties of our model. Let us denote by $F^*(\sigma)$ and y^* the values of the state at equilibrium. We have the following proposition, whose proof is direct from the dynamics:

Proposition 1: The number of seeders in equilibrium is

$$y^* = \frac{\lambda}{\gamma}.$$

The following result concerns the download time:

Proposition 2: Assume that in equilibrium each peer is assured a minimal download rate, i.e. $\omega(F^*, y^*, \sigma) \geq \varepsilon > 0 \quad \forall \sigma$. Then the download time satisfies:

$$\lambda \bar{T}_{leecher} = F^*(1). \quad (11)$$

Proof: Notice that in equilibrium a given leecher downloads an amount of content $d\sigma$ in a time $1/\omega(F^*, y^*, \sigma)$ and therefore we can calculate the download time as:

$$\bar{T}_{leecher} = \int_0^1 \frac{1}{\omega(F^*, y^*, \sigma)} d\sigma = \int_0^1 \frac{1}{\lambda} \frac{\partial F^*}{\partial \sigma} d\sigma = \frac{F^*(1)}{\lambda}. \quad \blacksquare$$

Equation (11) is a fluid version of Little's law of queueing theory.

Let us now focus in the efficiency of the file sharing and its consequences on the equilibrium. We have the following definition:

Definition 2 (Efficient sharing): A bandwidth sharing model $\omega(F, y, \sigma)$ is called *efficient* if $\Omega_{down}(t)$ satisfies:

$$\Omega_{down}(t) = \min\{\bar{\Omega}_{up}(t), cF(t, 1)\}$$

at all times. If the above is satisfied in the equilibrium we call the equilibrium *efficient*.

Note that the constraints (6), (7) in ω imply that Ω_{down} is always less or equal that the right hand side. The rationale behind Definition 2 is that if Ω_{down} is strictly less that both $\bar{\Omega}_{up}$ and $cF(t, 1)$, then the system has spare upload capacity, and has some leechers below their download capacity, thus a potential exchange opportunity is not being used.

We have the following bound for the number of leechers in equilibrium:

Lemma 1: The number of leechers in equilibrium verifies:

$$F^*(1) \geq \max\left\{\frac{\lambda}{c}, \lambda\left(\frac{1}{\mu} - \frac{1}{\gamma}\right)\right\} \quad (12)$$

independently of the shape of ω . The equality holds if and only if the equilibrium is efficient.

Proof: By equation (10) in equilibrium we must have:

$$\lambda = \Omega_{down}^* \leq \min\{\bar{\Omega}_{up}^*, cF^*(1)\}. \quad (13)$$

In particular we have $F^*(1) \geq \lambda/c$. Recalling the definition of $\bar{\Omega}_{up}$ we have

$$\lambda \leq \mu(y^* + F^*(1)).$$

Using that $y^* = \lambda/\gamma$ by Proposition 1 we get:

$$F^*(1) \geq \lambda \left(\frac{1}{\mu} - \frac{1}{\gamma} \right).$$

Combining the above equations we have the desired inequality. Note that equality for $F^*(1)$ holds if and only if equality holds in (13), which by definition is only achieved when the equilibrium is efficient. ■

We now analyze efficient bandwidth sharing policies at equilibrium. Recall that for the model (2) we defined $\gamma_{cr} := \mu \frac{c}{c-\mu}$. For our more general model we have the following result, analogous to the case (i') discussed earlier:

Proposition 3: If the bandwidth sharing is efficient and $\gamma < \gamma_{cr}$, then in equilibrium $\omega(F^*, y^*, \sigma) = c \forall \sigma$.

Proof: Suppose that $\omega(F^*, y^*, \sigma) < c$ for some value of σ . Then, $\Omega_{down}^* < cF^*(1)$. In equilibrium, equation (10) guarantees also that $\Omega_{down}^* = \lambda$, and we conclude that $F^*(1) > \frac{\lambda}{c}$. Since we assumed that the sharing is efficient, and $\Omega_{down}^* < cF^*(1)$ it must be by Definition 2:

$$\lambda = \Omega_{down}^* = \bar{\Omega}_{up}^* = \mu(y^* + F^*(1)) > \lambda\mu \left(\frac{1}{\gamma} + \frac{1}{c} \right).$$

And thus we require:

$$\mu \left(\frac{1}{\gamma} + \frac{1}{c} \right) < 1 \Leftrightarrow \frac{1}{\gamma} < \frac{1}{\mu} - \frac{1}{c} = \frac{1}{\gamma_{cr}}$$

which contradicts our hypothesis $\gamma < \gamma_{cr}$. ■

For the case $\gamma > \gamma_{cr}$, a bandwidth sharing cannot operate in equilibrium fully saturated by download, as the following Proposition shows:

Proposition 4: If $\gamma > \gamma_{cr}$, then in equilibrium $\omega(F^*, y^*, \sigma) < c$ for some σ .

Proof: Recall that from Lemma 1:

$$F^*(1) \geq \lambda \left(\frac{1}{\mu} - \frac{1}{\gamma} \right),$$

and from $\gamma > \gamma_{cr}$ the right hand side is greater than λ/c . We conclude that $cF^*(1) > \lambda$ and recalling that from (10) $\lambda = \Omega_{down}^*$ we have:

$$cF^*(1) > \Omega_{down}^* = \int_0^1 \omega(F^*, y^*, \sigma) \frac{\partial F^*}{\partial \sigma} d\sigma$$

We conclude that:

$$\int_0^1 (c - \omega(F^*, y^*, \sigma)) \frac{\partial F^*}{\partial \sigma} d\sigma > 0$$

and thus $\omega(F^*, y^*, \sigma) < c$ for some σ . ■

We end this Section with the following results on the download time. If we combine Proposition 2 and Lemma 1 we arrive at the following result:

Proposition 5: The download time in equilibrium satisfies the following inequality:

$$\bar{T}_{leecher} \geq \max \left\{ \frac{1}{c}, \frac{1}{\mu} - \frac{1}{\gamma} \right\},$$

with equality if and only if the equilibrium is efficient.

It is interesting to interpret the above proposition in the case where $c \rightarrow \infty$, i.e. when the downlink capacity is not the bottleneck, and the torrent is globally sustained, i.e. $\mu < \gamma$. Recalling that $1/\gamma$ represents the time spent in the system as a seeder, the bound of Proposition 5 becomes:

$$\bar{T} = \bar{T}_{leecher} + \bar{T}_{seed} \geq \frac{1}{\mu}.$$

In this case, Proposition 5 tells us that the minimal time spent in the system by a given peer is the time needed to upload a full copy of the file, that is, to return to the system the same amount of content it has downloaded. This minimum is achieved if and only if the bandwidth sharing is efficient.

IV. BANDWIDTH SHARING MODELS

We will now analyze some properties under specific choices for $\omega(F, y, \sigma)$. From now on we will focus on the case of globally sustained torrents, i.e. $\mu < \gamma$ and with no downlink capacity limit ($c \rightarrow \infty$). Note that in this case $\gamma_{cr} = \mu$.

A. Processor sharing discipline

The simplest model for bandwidth sharing is the *processor sharing* model in which each leecher gets the same amount of bandwidth, namely:

$$\omega(F, y, \sigma) = \frac{\mu(y(t) + F(t, 1))}{F(t, 1)} = \frac{\Omega_{up}(t)}{F(t, 1)}$$

The processor sharing model is efficient and its equilibrium is obtained by noting that $\omega(F, y, \sigma)$ is independent of σ and therefore:

$$\frac{\partial F^*}{\partial \sigma} = \frac{\lambda}{\omega(F^*, y^*, \sigma)} = K.$$

Thus using the result of Lemma 1:

$$F^*(\sigma) = F^*(1)\sigma = \lambda \left(\frac{1}{\mu} - \frac{1}{\gamma} \right) \sigma, \quad y^* = \frac{\lambda}{\gamma}.$$

Note that in equilibrium the distribution of leechers is uniform between different states of download. In [5] the authors assume this as an hypothesis for the system, here it is a consequence of the sharing mechanism chosen. In Section V we will analyze the stability properties of this equilibrium.

B. A generalization of the processor sharing discipline

We can generalize the processor sharing discipline by introducing a weight function $g(\sigma) > 0$ that determines the profile of bandwidth distribution among leechers. By analogy with queueing systems, we call this the *discriminatory processor sharing* model, and in this case:

$$\omega(F, y, \sigma) = \frac{\mu(y(t) + F(t, 1))g(\sigma)}{\int_0^1 g(\sigma) \frac{\partial F}{\partial \sigma} d\sigma}$$

Note that this sharing discipline is also efficient. In equilibrium we have:

$$\frac{\partial F^*}{\partial \sigma} = \frac{\lambda}{\omega^*} = \frac{\lambda \int_0^1 g(\sigma) \frac{\partial F^*}{\partial \sigma} d\sigma}{\mu(y^* + F^*(1))g(\sigma)} = \frac{K}{g(\sigma)}, \quad (14)$$

where K must verify $\int_0^1 \frac{K}{g(\sigma)} d\sigma = F^*(1) = \lambda \left(\frac{1}{\mu} - \frac{1}{\gamma} \right)$.

The introduction of $g(\sigma)$ allows us to model different sharing regimes, where download rates are dependent on the degree of advance. Note that due to efficiency, the total download time is *not* altered by this choice; what changes is the equilibrium distribution of leechers in different download stages.

C. Random selection model

We now turn our attention to a different sharing regime which takes into account a different aspect of the BitTorrent protocol, the tit-for-tat incentive mechanism. Under this exchange mechanism, a leecher exchanges data with another provided that each has something that the other one needs. To model this consider two leechers with content σ_1 and σ_2 . An approximation of the probability that two leechers at random can exchange data is $p(\sigma_1, \sigma_2) = \sigma_1(1 - \sigma_2)\sigma_2(1 - \sigma_1)$. Assume now that a leecher with content σ chooses another leecher at random. By appropriately modelling the matching probability, and assuming the uplink bandwidth of seeders is equally shared, the following model can be derived for $\omega(F, y, \sigma)$:

$$\omega(F, y, \sigma) = \frac{\mu}{F(t, 1)} \left(y(t) + \sigma(1 - \sigma) \int_0^1 s(1 - s) \frac{\partial F}{\partial \sigma} ds \right). \quad (15)$$

Equation (15) can be interpreted as a fluid version of the model proposed in [6]. We note now that this bandwidth sharing is not efficient in the sense of Definition 2. We can bound Ω_{down} for the ω given in (15) by:

$$\Omega_{down} \leq \mu \left(y(t) + \frac{1}{16} F(t, 1) \right) < \Omega_{up}.$$

The above model predicts that the system significantly under-utilizes the upload capacity of leechers, a fact not supported by real life and simulation evidence (see e.g. [5]): thus the model proposed in (15) is pessimistic. The main reason is that in real BitTorrent, peer exchange opportunities are not the result of random encounters: in contrast, each peer knows exactly the chunks available at other peers, and therefore tries to exchange data in a more directed fashion. There is a high chance that exchange opportunities will be found to exploit the available upload capacity.

V. STABILITY ANALYSIS

A. Stability of the processor sharing model

In this Section we will analyze the stability of the processor sharing model of Section IV-A. As in Section IV we assume that the torrent is globally sustained, i.e. $\mu < \gamma$ and for simplicity we assume $c \rightarrow \infty$, although the results remain valid for finite c whenever $\gamma > \gamma_{cr}$.

The model is given by equations (4) and (5) with:

$$\omega(F, y, \sigma) = \frac{\mu(y + F(t, 1))}{F(t, 1)}.$$

For tractability reasons, we will deal with a discretized version of the PDE (4). Let us define $n_i(t) = F(t, i/M) - F(t, (i-1)/M)$ for $i = 1, \dots, M$, i.e. the number of leechers

with content $\sigma \in [\frac{i-1}{M}, \frac{i}{M}]$. M is the number of intermediate states of download, and can be thought as the number of file chunks of the torrent.

Note that $\frac{\partial F}{\partial \sigma} \Big|_{\sigma=i/M} \approx Mn_i$ and that:

$$\omega(F, y, i/M) = \frac{\mu(y(t) + \sum_{i=1}^M n_i)}{\sum_{i=1}^M n_i}.$$

Defining $x = \sum_{i=1}^M n_i$, we arrive at the following discretized dynamics for the PDE:

$$\dot{n}_1 = \lambda - \frac{M\mu(y+x)}{x} n_1, \quad (16a)$$

$$\dot{n}_i = \frac{M\mu(y+x)}{x} (n_{i-1} - n_i), \quad i = 2, \dots, M, \quad (16b)$$

$$\dot{y} = \frac{M\mu(y+x)}{x} n_M - \gamma y. \quad (16c)$$

Note that the case $M = 1$ corresponds to the model described in (2). We have the following proposition, which is verified directly:

Proposition 6: The equilibrium of this system is:

$$n_i^* = \frac{\lambda}{M} \left(\frac{1}{\mu} - \frac{1}{\gamma} \right), \quad y^* = \frac{\lambda}{\gamma}.$$

We will analyze the local behavior of the dynamics around the equilibrium. In order to do so it is convenient to choose all units relative to μ , taking $\mu = 1$ and replacing λ, γ by $\lambda/\mu, \gamma/\mu$. With this choice, we are interested in the case $\gamma > 1 = \mu$ and we define $\alpha = \gamma - 1 > 0$. We need the following lemmas, whose proof is omitted for brevity:

Lemma 2: The Jacobian of the dynamics (16) around the equilibrium is the following $(M+1) \times (M+1)$ matrix:

$$J = \frac{1}{\alpha} \begin{bmatrix} 1 - M(1+\alpha) & 1 & \dots & 1 & -\alpha \\ M(1+\alpha) & -M(1+\alpha) & 0 & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & \dots & M(1+\alpha) & -M(1+\alpha) & 0 \\ -1 & \dots & -1 & (M(1+\alpha) - 1) & -\alpha^2 \end{bmatrix}$$

Lemma 3: The determinant of J is:

$$\det(J) = (-\alpha)^{1-M} (M(1+\alpha))^M,$$

and in particular $\det(J) \neq 0$ for any $\alpha > 0$ and $M \geq 1$, thus J does not have 0 as eigenvalue.

We are now ready to state the main result of this Section:

Theorem 1: The dynamics (16) are locally asymptotically stable for any $M \geq 1$ whenever $\alpha > 2$ (i.e. $\gamma > 3\mu$).

Proof: The proof relies on the Gerschgorin Circle Theorem [7] applied to the matrix J . Recall that the eigenvalues of a matrix $A = (a_{ij})$ are contained in the set:

$$G = \bigcup_i D(a_{ii}, R_i),$$

where $D(a_{ii}, R_i)$ is the disk in the complex plane centered in a_{ii} with radius $R_i = \sum_{j \neq i} |a_{ij}|$.

Direct application of the Gerschgorin Theorem to the matrix J does not give stability unless α grows linearly in M . To refine this result, consider the following transformation:

$$A = T^{-1}(\alpha J)T \quad \text{with} \quad T = \text{diag}(1, \dots, 1, \beta), \quad (17)$$

where $\beta > 0$ is a free parameter that we will choose later. Note that the eigenvalues of A are the same as those of J scaled by $\alpha > 0$, so if the eigenvalues of A lie in the region $\{Re(z) < 0\}$, the matrix J will be stable.

Applying the transformation of (17) only affects the first and last rows of αJ . We now apply the Gerschgorin result to the matrix A . The circles for $i = 2, \dots, M$ are given by:

$$a_{ii} = -M(1 + \alpha) < 0, \quad R_i = M(1 + \alpha),$$

and thus $D(a_{ii}, R_i)$ is entirely contained in the region $\{Re(z) < 0\} \cup \{0\}$ for any $i = 2, \dots, M$.

For the first row we have:

$$a_{11} = 1 - M(1 + \alpha) < 0, \quad R_1 = M - 1 + \alpha\beta,$$

and thus imposing $a_{11} + R_1 < 0$ we have that the circle lies in the region $\{Re(z) < 0\}$ whenever $\beta < M$. The last row gives:

$$a_{M+1, M+1} = -\alpha^2 < 0, \quad R_{M+1} = \frac{M-1}{\beta} + \frac{M(1+\alpha)-1}{\beta},$$

and again imposing $a_{M+1, M+1} + R_{M+1} < 0$ we have that the circle lies in the region $\{Re(z) < 0\}$ whenever

$$\beta > \frac{2(M-1)}{\alpha^2} + \frac{M}{\alpha}.$$

Note that in order to find a value of β that satisfies both conditions we need:

$$\alpha^2 - \alpha - \frac{2(M-1)}{M} > 0.$$

In particular, for any $\alpha \geq 2$ this is valid for any M .

Combining the above with Lemma 3, we conclude that the matrix J is Hurwitz and therefore the system is locally asymptotically stable. ■

Numerical analysis shows that the system is indeed stable if $\alpha < 2$, and even when $\alpha \rightarrow 0$ ($\gamma \downarrow \mu$). In particular, for $M = 2$ this can be verified via the Routh-Hurwitz criterium. The case $M > 2$ is open for future work. In Figure 1 we plot the evolution of the eigenvalues of J with $\alpha > 0$ for $M = 4$, which captures the general behavior of the dominant poles.

B. Local stability for $\gamma < \gamma_{cr}$

In Section III-B we showed that an efficient torrent, with for $\gamma < \gamma_{cr}$, must satisfy $\omega(F^*, y^*, \sigma) = c, \forall \sigma$. The equilibrium of (4), (5) is therefore:

$$F^*(\sigma) = \frac{\lambda}{c}\sigma, \quad y^* = \frac{\lambda}{\gamma},$$

and using the same approach that in V-A, the discretized dynamics have the following Jacobian around the equilibrium:

$$J = \begin{bmatrix} -Mc & 0 & 0 & \dots & 0 \\ Mc & -Mc & 0 & \dots & 0 \\ 0 & Mc & -Mc & \dots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & \dots & 0 & Mc & -\gamma \end{bmatrix},$$

which is clearly a Hurwitz matrix. Therefore, in the case where $\gamma < \gamma_{cr}$ and the torrent is efficient, the system is locally asymptotically stable.




Fig. 1. Complex plane plot of the eigenvalues of J with varying α for $M = 4$.

VI. CONCLUSIONS AND FUTURE WORK

In this paper we have analyzed the dynamics of a P2P file sharing protocol like BitTorrent. We derived a PDE that describes the evolution of the content distribution in the system. We described the general properties of the equilibrium of the model. In particular, we showed that the download time in equilibrium of a given torrent is independent of the sharing mechanism, provided it is efficient. We also analyzed how the model is suitable to describe different sharing situations, generalizing previous works on the subject. We also derived local stability results for the processor sharing discipline under very general conditions, generalizing the results of [5].

In future work, we plan to address the local stability of the processor sharing discipline in the case where $\gamma \downarrow \mu$, as well as the generalized sharing discipline proposed. Another interesting line of work is the fairness between users. As we mentioned, the download time in our fluid model is determined by external parameters, but we can expect the real system to operate randomly around this equilibrium. Determining the impact of these fluctuations in the fairness of the system will be analyzed in future work.

REFERENCES

- [1] B. Cohen, *Incentives build robustness in BitTorrent*, 1st Workshop on the Economics of Peer-2-Peer Systems, Berkeley, 2003.
- [2] G. Kesidis, T. Konstantopoulos, and P. Sousi, *Modeling file-sharing with bittorrent-like incentives*, ICASSP, IEEE International Conference on Acoustics, Speech and Signal Processing - Proceedings, vol. 4, 2007, pp. IV1333-IV1336.
- [3] L. Massoulié and M. Vojnović, *Coupon replication systems*, IEEE/ACM Transactions on Networking **16** (2008), no. 3, 603-616.
- [4] D. Qiu and W. Sang, *Global stability of peer-to-peer file sharing systems*, Computer Communications **31** (2008), no. 2, 212-219.
- [5] D. Qiu and R. Srikant, *Modeling and performance analysis of BitTorrent-like peer-to-peer networks*, ACM SIGCOMM Computer Communication Review **34** (2004), no. 4, 367-378.
- [6] Y. Tian, Di Wu, and KW Ng, *Performance analysis and improvement for BitTorrent-like file sharing systems*, Concurrency and Computation: Practice and Experience **19** (2007), no. 13, 1811-1835.
- [7] Richard S. Varga, *Matrix iterative analysis*, Springer-Verlag, NY, 2000.
- [8] X. Yang and G. De Veciana, *Service capacity of peer-to-peer networks*, Proceedings of IEEE/Infocom, 2004.