Dynamics of heterogeneous peer-to-peer networks

Fernando Paganini, Andrés Ferragut and Martín Zubeldía
Universidad ORT Uruguay

***Abstract*— The most tractable models of population dynamics in peer-to-peer file sharing systems apply to a single class of peers with homogeneous network access parameters. When upload bandwidths are heterogeneous, reciprocity mechanisms lead to non-uniform download rates and a more complex multi-class dynamics. We consider first a model where mutual download bandwidths are allocated in proportion to the upload speed, plus a uniformly distributed server component. For an ordinary differential equation model of the multi-class peer populations, we characterize the equilibrium and establish its global stability, invoking results from monotone systems. We also analyze a partial differential equation model that tracks download progress of the populations; we establish the local asymptotic stability of the equilibrium. Finally, we extend the ODE model to include a mix of proportional and uniform bandwidth allocation, which better describes the mechanisms of BitTorrent systems; again we characterize equilibrium configurations and give a partial result on local stability.**

I. Introduction

Peer-to-peer (P2P) file-sharing networks have become a popular means to distribute content over the Internet. They are based on the principle that clients downloading a file can themselves contribute their upload bandwidth to serve others, thus achieving a valuable self-scaling property as supply of server bandwidth increases with demand. One of the prevailing P2P systems is BitTorrent [2], which incorporates a reciprocity incentive: peers will orient their upload towards others from whom they have downloaded the most. This is designed to avoid free-riding; for game theoretic studies of the inherent incentives we refer to [9].

P2P networks are inherently dynamic: the population of peers participating in the sharing of a certain file will vary over time, as peers arrive and leave the swarm. The speed of departures is determined by download rates, which themselves depend on populations; understanding this feedback dynamics has been a topic of active research. A Markov queueing model was proposed in [18], and led subsequently to an ordinary differential equation model [14], which has been successful in estimating equilibrium populations, and establishing stability [13]. These models are coarse in the sense that the state does not discriminate download progress of the swarm; in this regard, in [4], [12] it is shown how download quantities can be tracked by a partial differential equation model that leads to tighter dynamic predictions than the models of [14]. The above models are reviewed in Section II. Other references that attempt a finer tracking of file pieces are [7], [11], [20].

This work was supported in part by ANII-Uruguay under Grant FCE_2_2011_1_7052, and by AFOSR-US under Grant FA9550-12-1-0398.
E-mail:paganini@ort.edu.uy.

The above analytical work on P2P dynamics focuses on the case of *homogeneous* peers, with common access bandwidth parameters. Here it is natural to assume that the real-time download bandwidth is evenly distributed, a "processor-sharing" assumption that simplifies mathematics and shows good agreement with simulations of homogeneous BitTorrent systems. When access parameters are *heterogeneous*, however, bandwidth allocation can, and arguably should, be uneven: a recent reference for this design space is [3]. In Section III we review an idealized proportional reciprocity scheme where each peer receives as much as it gives. The practical tit-for-tat mechanisms of BitTorrent do induce a certain service differentiation of this kind, whereas other aspects of the protocol incorporate a processor-sharing component, as studied empirically in [8]. When it comes to analytical work on population dynamics, far less is known; some early work that extends the fluid model of [14] with a two-class dynamics is [1], [10], however bandwidth allocation is not related to reciprocity schemes in these references.

The object of this paper is to analyze the population dynamics of heterogeneous P2P swarms where resource allocation reflects the reciprocity incentives. In Section IV we analyze the case of proportional reciprocity with an ordinary differential equation model; we characterize the unique equilibrium point of the system and prove it is globally stable by means of results in monotone systems [6]. In Section V we extend the model to a multi-class version of the PDE in [12]; again under proportional reciprocity, we describe the equilibrium and in this case analyze the local aymptotic stability through small-gain and Lyapunov arguments. In Section VI we consider the situation where bandwidth is distributed according to a mixed policy: a fraction of proportional allocation combined with a fraction of processor sharing. Using an ODE model, again we characterize the equilibrium and provide a partial result on its local stability. Conclusions are given in Section VII.

II. Background on P2P dynamics

In a P2P system, *content* is disseminated by subdividing it into small *chunks*, and enabling peers to exchange such units bidirectionally. Thus every peer present is a server; those who are also clients are referred to as *leechers*, whereas *seeders* are those peers present in the system only to altruistically distribute content. Populations of both types of peers may vary dynamically as arrivals and departures from the swarm take place, and also leechers may turn into seeders upon termination of their download, as has been considered in many models [14], [18]. However a very common scenario in practice is to have a few persistent seeders who act as overall servers for the content, and selfish leechers who just abandon the swarm upon termination. In this paper we will restrict our attention to this simplified scenario, that nevertheless captures the essence of p2p file-sharing.

In this section we review models for peer population in the *homogeneous* case, where all peers have upload bandwidth μ (in files/second, i.e. content is normalized to unit size), and a much larger download capacity that is never a bottleneck.

A. Ordinary differential equation model

Following [14], let $x(t) \geq 0$ denote the population of leechers in the system, taken to be a real variable, with arrival rate λ peers/second. Let y_0 be the fixed population of seeders. The total upload rate $\mu(y_0 + x)$ in files/second determines the departure rate of leechers, so the population dynamics is

$$\dot{x} = \lambda - \mu(y_0 + x). \tag{1}$$

Clearly, if $\lambda < \mu y_0$ the trajectories of (1) converge to zero in finite time, at which point the equation must be projected for x to stay at the zero boundary. This, however, is not an interesting case, since servers y_0 on their own could satisfy the download demand without any p2p contribution.

The more important case is $\lambda > \mu y_0$. Here trajectories will remain positive and converge to the equilibrium point

$$x^* = \frac{\lambda}{\mu} - y_0. \tag{2}$$

B. Partial differential equation model

The dynamic model (1) contains limited information on the system state: the leecher population is counted by a single variable, without any detail on the progress in their download. To characterize the latter, without the complex detail of keeping track of specific chunks, in [4], [5], [12] it is proposed to track populations as a function of the *fraction* of downloaded content, treated as a continuous variable.

Let $\sigma \in [0, 1]$ represent fraction of the content file, assumed of unit size. Define the real-valued variable $F(t, \sigma)$ that represents the population of leechers that at time t, have pending download of at least σ. Thus $F(t, \sigma)$, nonincreasing in σ acts as the complementary cumulative distribution of the leecher population, with $F(t, 0) = x(t)$, the total leecher count and $F(t, 1) = 0$. The dynamic model from [12] is:

$$\frac{\partial F}{\partial t} = \lambda + r\frac{\partial F}{\partial \sigma}, \quad \sigma \in [0, 1]. \tag{3}$$

Here the arrival rate λ increases the entire distribution, assuming leechers arrive with no prior content. r represents the download rate per peer: it regulates the speed at which the function $F(t, \sigma)$ is transported in the direction of $\sigma = 0$. In our scenario of fixed seeders, the simplest expression for the rate is

$$r = \frac{\mu(x + y_0)}{x}. \tag{4}$$

This assumes: (i) *Efficiency*, i.e. the entire upload bandwidth $\mu(x+y_0)$ is available for download; (ii) a *processor sharing* discipline, i.e. the upload bandwidth is uniformly distributed among leechers. Empirical evidence indicates that this is quite an accurate model for BitTorrent systems under the homogeneity assumption (see [5], [12]).

The equilibrium of (3) under (4) is a uniform distribution of download, $F^*(\sigma) = x^*(1 - \sigma)$, where x^* is given by (2). So ODE and PDE models coincide in their prediction of the equilibrium population; when it comes to dynamics, however, it is shown in [5], [12] that the finer state description of the PDE model provides more accurate predictions.

III. Resource allocation in heterogeneous p2p

Consider now the situation where peers participating in the swarm have a *heterogeneous* access to the network, specified by a set of possible upload rates $\{\mu_i\}_{i=1}^n$. Assume there are x_i leechers at each of the classes, how is their total upload capacity $\sum_i \mu_i x_i$ distributed among the leecher population? This can be taken a design question (which allocation *should* there be?) or as a descriptive question about current P2P systems (e.g. BitTorrent). A recent reference that explores the design space is [3], exhibiting a tradeoff between *performance* (minimizing mean download time) and *fairness*, understanding the latter to mean parity between how much peers give to, and receive from, the network.

An important requirement for the allocation is a decentralized implementation, i.e. a set of mutual exchange rules peers can follow to achieve it, without the intervention of a central authority. We now review a proportional reciprocity scheme that can achieve the aforementioned fairness, assuming a fine control of mutual rates. We then compare it with the result of the more practical mechanisms of BitTorrent.

A. Proportional reciprocity

Consider in this section a fixed, finite set of n leechers, with upload bandwidths $\{\mu_i\}_{i=1}^n$, possibly with repetitions. Let k be a discrete-time index that represents an exchange slot, and let $z_{ij}^{(k)}$ denote the bandwidth devoted by peer i to peer j in the k-th slot. Assuming efficiency we have

$$z_{ii}^{(k)} = 0, \qquad \sum_j z_{ij}^{(k)} = \mu_i.$$

The total download rate of peer j from all leechers is

$$r_j^{(k)} = \sum_i z_{ij}^{(k)}.$$

In matrix terms: if $Z^{(k)}$ is the $n \times n$ matrix with components $z_{ij}^{(k)}$, it has zero diagonal and satisfies the relationships

$$Z^{(k)}\mathbf{1} = \mu, \qquad [Z^{(k)}]^T\mathbf{1} = r^{(k)};$$

here $\mathbf{1}$, μ, and $r^{(k)}$ are column vectors with components 1, μ_i, and $r_i^{(k)}$ respectively, and ^{T} denotes transpose.

Based on received rates, peers must select their allocation for the following slot; a natural rule considered in [9], [17], [19] is *proportional reciprocity*: give to others in proportion to what is received from them. Mathematically

$$z_{ij}^{(k+1)} = \mu_i.\frac{z_{ji}^{(k)}}{r_i^{(k)}},$$

or $Z^{(k)} = \mathrm{diag}(\mu_i / r_i^{(k)})[Z^{(k+1)}]^T$. This means to transpose the matrix and renormalize rows to have sum μ. Modulo the transpose operation we have then an iterative row and column renormalization of a non-negative matrix, a topic studied classically by Sinkhorn [15], who established conditions for convergence. This connection was made in [17]. In particular, provided the equations

$$Z\mathbf{1} = \mu, \qquad Z^T\mathbf{1} = \mu$$

are jointly feasible with a matrix Z of the prescribed structure, row and column renormalization will converge. If one only imposes a zero diagonal structure, feasibility will hold provided no single μ_i is greater than the sum of the rest. This is a mild restriction if the peer population is large. Under these circumstances, the odd and even subsequences of $Z^{(k)}$ converge, and the rates $r^{(k)} \stackrel{k\to\infty}{\longrightarrow} \mu$.

The conclusion is that proportional reciprocity allocates each peer, asymptotically, a download equal to its upload. In the homogeneous case this implies a processor-sharing model, but in the heterogeneous case service becomes differentiated in proportion to contributions to the swarm.

B. BitTorrent's tit-for-tat and optimistic unchoke

The above rule, while elegant, is not easily implemented in practice: it requires maintaining open connections with all other peers, and regulating their rate in a differentiated fashion, something not simple to implement with Internet TCP connections. The alternative implemented by BitTorrent is a ranking scheme: classify peers according to bandwidth received from them in a recent period, and then "unchoke" (allow a transfer to) those peers which occupy the highest places. It turns out that such scheme also closely approximates a proportional allocation, since peers with this rule tend to form *cliques* with others of similar upload bandwidth (for a theoretical justification of this fact see [3], Thm 2.).

In addition to the above tit-for-tat rule, BitTorrent introduces an *optimistic unchoke* to another peer at random: this is done to explore the set of peers. This portion of the upload bandwidth is then distributed in an egalitarian fashion in the swarm. The result is therefore a combination of proportional and processor sharing allocations. Empirical studies that approximately validate these claims are presented in [8].

IV. ODE MODEL UNDER PROPORTIONAL RECIPROCITY

In this section we begin our theoretical analysis of a heterogeneous p2p network. We consider a swarm of p2p leechers downloading a common content file, classified in n groups according to their upload bandwidths $\{\mu_i\}$. Proportional reciprocity is assumed, so each peer receives from other leechers a service rate equal to its own upload rate. As just explained, this model is consistent with the tit-for-tat portion of the BitTorrent exchange mechanism. A generalization that covers the egalitarian portion inherent in the "optimistic unchoke" will be discussed in Section VI.

We will also assume that there is a fixed set of seeders $y = y_0$, each of which has upload bandwidth μ_0 [1]. Let x_i denote the fluid population of leechers in class i, of upload bandwidth μ_i. An ordinary differential equation model for the population dynamics is:

$$\dot{x}_i = \lambda_i - \left[\frac{\mu_0 y_0}{\sum_{j=1}^n x_j} + \mu_i \right] x_i. \tag{5}$$

Here $\lambda_i > 0$ is the arrival rate of leechers of class i. We are assuming that the leecher sharing portion of the download satisfies proportional reciprocity, and thus the download rate per peer is equal to the class upload bandwidth; however the seeder portion of the download is equally distributed among all the peer population. This model parallels (1) for the single-class case; note however that we now have a more complex, nonlinear ODE.

Remark 1: Since the state represents populations, the above equation applies only to the positive orthant $\mathbb{R}^n_+$. In regard to the boundary: if a nonzero x has one coordinate $x_i = 0$, the corresponding right-hand side of (5) is $\lambda_i > 0$, hence the flow moves back to $x_i > 0$. Equivalently, an initial condition $x_i(0) > 0 \ \forall i$ will never reach one of these boundary faces of the orthant. The only degenerate point (where our equation is not well defined) is $x = 0$: the next assumption implies that this point is never reached from positive initial conditions.

Assumption 1: $\sum_i \lambda_i > \mu_0 y_0$. This means that the seeders alone cannot cope with the service demands.

We now establish the existence and uniqueness of an equilibrium point of the dynamics under this condition.

Proposition 1: Under Assumption 1, the dynamics (5) has a unique equilibrium $x^* = (x_i^*)$ with $x_i^* > 0$ for each i.

Proof: First note that adding (5) over i gives

$$\sum_i \dot{x}_i = \sum_i \lambda_i - \mu_0 y_0 - \sum_i \mu_i x_i;$$

this rules out the possibility that x_i tends to zero for every i, since before this happens the sum $\sum_i x_i$ would begin to increase through Assumption 1. So trajectories will remain in the interior of the orthant. We look for equilibrium points x^* in this region. A first necessary condition follows from the preceding equation:

$$\sum_i \lambda_i - \mu_0 y_0 = \sum_i \mu_i x_i^*. \tag{6}$$

To obtain further necessary equilibrium conditions impose

$$\begin{aligned} 0 &= \dot{x}_i x_j^* - \dot{x}_j x_i^* \\ &= \lambda_i x_j^* - \lambda_j x_i^* - (\mu_i - \mu_j) x_i^* x_j^*, \end{aligned}$$

which implies that

$$\frac{\lambda_i}{x_i^*} - \mu_i = \frac{\lambda_j}{x_j^*} - \mu_j \quad \forall\ i, j; \tag{7}$$

[1] Heterogeneous seeders could be included with essentially no change, this is avoided for simplicity.

we denote the above quantity by α. Then (6) implies that

$$\alpha = \frac{\mu_0 y_0}{\sum_i x_i^*} > 0. \tag{8}$$

Therefore α represents the equilibrium fraction of seeder bandwidth available per leecher, common to all classes.

Solving now (7) gives

$$x_i^* = \frac{\lambda_i}{\mu_i + \alpha}, \tag{9}$$

which can be substituted in (6) to express the necessary condition for equilibrium as a single equation in α:

$$\sum_i \lambda_i \frac{\alpha}{\mu_i + \alpha} = \mu_0 y_0. \tag{10}$$

Denote the left-hand side of (10) as $g(\alpha)$. It is strictly increasing in $\alpha \geq 0$, $g(0) = 0$, $g(+\infty) = \sum_i \lambda_i > \mu_0 y_0$. Therefore there exists a unique root $\alpha > 0$ to (10). Substituting in (9), there is a unique point $x^* > 0$ that satisfies the conditions for equilibrium. It is straightforward to show conversely that satisfying (9-10) is sufficient to have an equilibrium of (5). ■

We will now establish the global stability of this equilibrium point. The key observation is that we are in the realm of *monotone dynamical systems* [6], because the vector field in (5) (let us denote it by $h(x)$) satisfies the *cooperative* condition

$$\frac{\partial h_i}{\partial x_j} \geq 0 \quad \text{for } i \neq j. \tag{11}$$

Flows satisfying this property[2] are monotonic, in the sense of preserving vector inequalities, which greatly narrows down the possibilities for the dynamics. We now state the main result of this section.

Theorem 2: Under the conditions of Proposition 1, the equilibrium is globally asymptotically stable.

Proof: We begin by verifying condition (11). Indeed we have for $i \neq j$:

$$\frac{\partial}{\partial x_j}\left(\lambda_i - \left[\frac{\mu_0 y_0}{\sum_l x_l} + \mu_i\right] x_i\right) = \frac{\mu_0 y_0 x_i}{(\sum_l x_l)^2} \geq 0.$$

It follows from [6] (Theorems 3.5 and 3.2) that the corresponding flow is monotone: for two initial conditions satisfying the vector inequality $x(0) \leq \hat{x}(0)$, the corresponding solutions satisfy $x(t) \leq \hat{x}(t) \quad \forall t$. Furthermore, a strong monotonicity holds in the interior of the orthant.

We now establish that the n-dimensional open interval $X = \{x \in \mathbb{R}_+^n : 0 < x_i < \lambda_i/\mu_i \text{ for each } i\}$ is positively invariant under the flow. Indeed, assume $x(0) \in X$; we already know trajectories remain strictly positive. For it to reach the boundary of X at some time t would require some component to satisfy $x_i = \lambda_i/\mu_i$ and $\dot{x}_i \geq 0$. But at such x_i the dynamics (5) gives

$$\dot{x}_i = -\frac{\mu_0 y_0}{\sum_l x_l} x_i < 0,$$

a contradiction. On the other hand, assume the initial condition is outside X. First we claim the state remains bounded. In fact: coordinates (if any) satisfying $x_i(0) < \lambda_i/\mu_i$ will remain within this bound as before, while those $i : x_i \geq \lambda_i/\mu_i$ will have $\dot{x}_i \leq 0$. We can thus write the bound

$$\frac{\mu_0 y_0}{\sum_i x_i} \geq \epsilon \quad \text{for some } \epsilon > 0.$$

But then any coordinate i with $x_i \geq \lambda_i/\mu_i$ satisfies $\dot{x}_i \leq -\epsilon x_i$, and therefore the set X is reached in finite time.

We can thus restrict our attention to the dynamics on the bounded open interval X. We have a strictly monotone flow with orbits of compact closure, with a single equilibrium point in X. Corollary 1.20 in [6] implies there is global convergence to equilibrium. ■

[2]Equivalently, the Jacobian matrix $\frac{\partial h}{\partial x}$ is *Metzler*, with non-negative off-diagonal elements.

V. Multi-class PDE model

In this section we will work with the partial differential equation model reviewed in Section II-B that keeps track of download progress in addition to population, and extend it to cover the multi-class situation. Specifically, we consider a content file of unit size, and the continuous variable σ representing file fraction: let $F_i(t, \sigma)$ be the fluid population of leechers of class i that have at time t a pending download of at least σ. The total class population is $F_i(t, 0) = x_i$. The corresponding PDE model takes the form

$$\frac{\partial F_i}{\partial t} = \lambda_i + \underbrace{\left(\frac{\mu_0 y_0}{\sum_{j=1}^n x_j} + \mu_i\right)}_{r_i(F, y_0)} \frac{\partial F_i}{\partial \sigma}, \quad \sigma \in [0, 1]. \tag{12}$$

The equilibrium analysis from the previous section extends readily to this case. Note from (12) that at equilibrium, $\frac{\partial F_i}{\partial \sigma}$ must be constant in σ, so we have the uniform distribution

$$F_i^*(\sigma) = x_i^*(1 - \sigma),$$

where x_i^* satisfies the same equilibrium conditions of the ODE case. Under Assumption 1 we have a unique equilibrium point, characterized by (9-10). We will show *local* stability of this equilibrium through linearization.

We use $\tilde{x}_i$, $\tilde{r}_i$ to denote incremental scalar variables, and lowercase notation for the function-valued incremental variable $f_i(t, \sigma) = F_i(t, \sigma) - F_i^*(\sigma)$; we have $f_i(t, 1) \equiv 0$ and $f_i(t, 0) = \tilde{x}_i$. The linearization is

$$\frac{\partial f_i}{\partial t} = r_i^* \frac{\partial f_i}{\partial \sigma} - x_i^* \tilde{r}_i \tag{13}$$

where the incremental rate is

$$\tilde{r}_i = -\frac{\mu_0 y_0}{(\sum_j x_j^*)^2} \sum_j \tilde{x}_j = -\frac{\alpha}{\sum_j x_j^*} \sum_j \tilde{x}_j,$$

with α from (8). The second term in (13) is now expressed as $-x_i^* \tilde{r}_i = r_i^* \kappa_i \sum_j \tilde{x}_j$, where

$$\kappa_i := \frac{\alpha}{r_i^*} \frac{x_i^*}{\sum_j x_j^*}.$$

Note $r_i^* = \alpha + \mu_i > \alpha$, so $\sum_i \kappa_i < 1$. Let us further denote

$$u_i := \kappa_i \sum_j \tilde{x}_j. \tag{14}$$

It will also be convenient to define $\tau_i = (r_i^*)^{-1} = \frac{x_i^*}{\lambda_i}$ (equilibrium download time per peer). It follows that the linearized dynamics is the feedback interconnection of:

- A set of parallel blocks G_i, with input u_i and output $\tilde{x}_i$, characterized by the infinite-dimensional dynamics

$$\frac{\partial f_i}{\partial t}(t,\sigma) = \frac{1}{\tau_i}\frac{\partial f_i}{\partial \sigma}(t,\sigma) + \frac{1}{\tau_i}u_i(t), \tag{15a}$$
$$\tilde{x_i}(t) = f_i(t,0), \tag{15b}$$
$$0 \equiv f_i(t,1). \tag{15c}$$

- The static mapping (14), represented in matrix form by

$$u = K\mathbf{1}\mathbf{1}^T\tilde{x}, \tag{16}$$

where K is the diagonal matrix $\mathrm{diag}(\kappa_i)$.

In [5] the single-class counterpart of the above dynamics was analyzed. It was shown that the block (15) has transfer function

$$\hat{G}_i(s) = \frac{1-e^{-\tau_i s}}{\tau_i s},$$

and in particular satisfies $\|\hat{G}_i(s)\|_\infty = \sup_{\omega\in\mathbb{R}} |\hat{G}_i(j\omega)| = 1$. In the present case, the loop transfer function $K\mathbf{1}\mathbf{1}^T\mathrm{diag}(G_i(s))$ is of rank one, so its input-output stability is equivalent to that of the scalar loop gain

$$L(s) = \mathbf{1}^T\mathrm{diag}(G_i(s))K\mathbf{1} = \sum_i \kappa_i \hat{G}_i(s).$$

It follows that

$$|L(j\omega)| \le \sum_i \kappa_i |\hat{G}_i(j\omega)| \le \sum_i \kappa_i < 1,$$

hence input-output stability holds by a small-gain argument.

The *internal* stability of the linearized dynamics can also be established through a Lyapunov argument, which is now skteched without proof due to space limitations.

For each of the infinite dimensional systems in (15), with local state f_i, introduce the storage functional

$$V_i(f_i) = \tau_i \int_0^1 \sigma \left[\frac{\partial f_i}{\partial \sigma}\right]^2 d\sigma. \tag{17}$$

The following dissipation inequality can be established under the dynamics (15):

$$\dot{V}_i = u_i^2 - \int_0^1 \left[\frac{\partial f_i}{\partial \sigma}\right]^2 d\sigma \le u_i^2 - \tilde{x}_i^2. \tag{18}$$

Closing now the feedback loop (16) and imposing $\sum_i \kappa_i < 1$ implies that the global Lyapunov functional

$$V(f) = \sum_i \kappa_i^{-1} V_i(f_i)$$

is decreasing along trajectories. A further refinement of this argument leads to the following asymptotic stability result:

Theorem 3: The dynamics (15-16) is asymptotically stable in the sense of $L_2[0,1]$; in particular

$$\|f_i(t,\sigma)\|_2 \overset{t\to\infty}{\longrightarrow} 0 \quad \text{for each } i.$$

VI. Generalized ODE model: including a processor sharing component

As discussed in Section III-B, BitTorrent systems implement a mixture of tit-for-tat reciprocity, which approximates a proportional allocation due to the formation of cliques, and optimistic unchoke, which provides an egalitarian file sharing. In this section we use ODE models to analyze this mixed behavior. The modified dynamics is now presented:

$$\dot{x}_i = \lambda_i - \underbrace{\left[\frac{\mu_0 y_0 + \theta\sum_{j=1}^n \mu_j x_j}{\sum_{j=1}^n x_j} + (1-\theta)\mu_i\right]}_{r_i} x_i \tag{19}$$

Here the parameter $\theta \in (0,1)$ controls the fraction of upload bandwidth that all peers devote to the egalitarian file-sharing. e.g. $\theta = 1/4$ in a typical BitTorrent implementation. Thus peers of class j will contribute a bandwidth $\theta\mu_j x_j$ to the common upload pool, which together with the seeder bandwidth $\mu_0 y_0$ will be uniformly distributed in the leecher swarm, hence the first term of the rate per peer r_i above. The second term above results from a reciprocity scheme, involving the fraction $(1-\theta)$ of the bandwidth. Assuming proportionality as before each leecher will receive from here a bandwidth equal to its contribution. A model for rates with this structure was considered in [3]. We now study the corresponding dynamics.

Proposition 4: Under Assumption 1, the dynamics (19) has a unique equilibrium $x^* = (x_i^*)$ with $x_i^* > 0$ for each i.

Proof: We generalize the proof of Proposition 1. Aggregating (19) over i gives again

$$\sum_i \dot{x}_i = \sum_i \lambda_i - \mu_0 y_0 - \sum_i \mu_i x_i,$$

therefore we can once again rule out the possibility of trajectories going to zero, and any equilibrium must satisfy (6). We further impose

$$\begin{aligned} 0 &= \dot{x}_i x_j^* - \dot{x}_j x_i^* \\ &= \lambda_i x_j^* - \lambda_j x_i^* - (1-\theta)(\mu_i - \mu_j)x_i^* x_j^*, \end{aligned}$$

which implies now that

$$\alpha := \frac{\lambda_i}{x_i^*} - (1-\theta)\mu_i = \frac{\lambda_j}{x_j^*} - (1-\theta)\mu_j \quad \forall\ i,j; \tag{20}$$

An analogous calculation leads to

$$\alpha = \frac{\mu_0 y_0 + \theta\sum_i \mu_i x_i^*}{\sum_i x_i^*} > 0. \tag{21}$$

α now represents the equilibrium download bandwidth each leecher receives from the "egalitarian" portion of the upload. Now (9) generalizes to

$$x_i^* = \frac{\lambda_i}{(1-\theta)\mu_i + \alpha}, \tag{22}$$

and further operations lead to the necessary condition

$$\sum_i \lambda_i - \sum_i \mu_i x_i^* = \underbrace{\sum_i \frac{\lambda_i(\alpha - \theta\mu_i)}{\alpha + (1-\theta)\mu_i}}_{g(\alpha)} = \mu_0 y_0. \tag{23}$$

The modified function $g(\alpha)$ is still strictly increasing, $g(+\infty) > \mu_0 y_0$, and now with $g(0) < 0$. So there is still a single solution to (23), which results through (22) in a single equilibrium point $x^* > 0$. ∎

We now tackle the question of stability of the equilibrium. Unfortunately, the additional term in the dynamics does not preserve the monotone property which was crucial for our earlier global stability argument. What follows is a partial result on *local* stability.

Theorem 5: Consider the equilibrium x^* of the system (19) under the conditions of Proposition 4. Suppose that α in (21) satisfies

$$\theta \mu_i \leq 2\alpha \text{ for each } i. \tag{24}$$

Then the equilibrium is locally asymptotically stable.

Before proceeding we provide an interpretation for condition (24). The left-hand side is the upload bandwidth that a leecher of class i contributes to the swarm in an egalitarian (non-reciprocal) way; α represents the download portion every peer receives from the egalitarian upload. (24) means that no peer is contributing in equilibrium, more than twice of what it receives; so there is a bound on the level of imbalance in this non-reciprocal component of the file-sharing.

Proof: We linearize the dynamics around equilibrium. The incremental rate $\tilde{r}_i$ now becomes

$$\begin{aligned}\tilde{r}_i &= \frac{\theta \sum_j \mu_j \tilde{x}_j}{\sum_j x_j^*} - \frac{(\mu_0 y_0 + \theta \sum_j \mu_j x_j^*) \sum_j \tilde{x}_j}{(\sum_j x_j^*)^2} \\ &= \frac{\sum_j (\theta \mu_j - \alpha) \tilde{x}_j}{\sum_j x_j^*},\end{aligned}$$

where we have invoked (21). The linearized dynamics is

$$\begin{aligned}\dot{\tilde{x}}_i &= -r_i^* \tilde{x}_i - x_i^* \tilde{r}_i \\ &= -r_i^* \tilde{x}_i + \frac{x_i^* \alpha}{\sum_j x_j^*} \sum_j (1 - \frac{\theta \mu_j}{\alpha}) \tilde{x}_j. \end{aligned} \tag{25}$$

Setting $D = \text{diag}(r_i^*)$, v the vector of components $\frac{x_i^* \alpha}{\sum_j x_j^*}$ and w the vector of components $1 - \frac{\theta \mu_j}{\alpha}$, we can write the linearized dynamics as $\dot{\tilde{x}} = A\tilde{x}$ with

$$A = -D + vw^T.$$

Note that A would be a Metzler matrix if the components of w were non-negative, which happens when $\theta \mu_j \leq \alpha$. Our hypothesis (24) is however weaker than this, all we can claim is that $|w_j| \leq 1$ for all j. Still, this enables a diagonal dominance argument. For fixed j write

$$\begin{aligned} a_{jj} + \sum_{i \neq j} |a_{ij}| &\leq -r_j^* + \sum_{i=1}^n |v_i w_j| \\ &\leq -r_j^* + \sum_{i=1}^n \frac{x_i^* \alpha}{\sum_j x_j^*} \quad = -r_j^* + \alpha < 0, \end{aligned}$$

(in fact $r_j^* = \alpha + \mu_j(1-\theta)$). Invoking for instance the Gershgorin Circle Theorem [16] applied to the columns of A, we find that its eigenvalues must lie in circles of center a_{jj}, radius $\sum_{i \neq j} |a_{ij}|$, contained in the open left half-plane. Therefore A is a Hurwitz matrix. ∎

VII. Conclusions

We have analyzed the dynamics of P2P networks under heterogeneity in access bandwidth. For multi-class models that discriminate peer populations according to this parameter, we studied the effect of reciprocity schemes which provide asymmetric download speeds to these classes. For proportional schemes where download speeds from other leechers equal the upload rate, we provide a complete analysis of the equilibrium and its global stability in the case of ODE population models, and a local stability analysis for the finer PDE models that track download advance. For ODE models with a mixture of proportional and egalitarian file sharing, partial stability results were given. In future work we will seek to complete the analysis of this mixed case, and pursue global questions for the case of PDE models.

References

[1] F. Clévenot-Perronnin, "Fluid Models for Content Distribution Systems," Ph.D. dissertation, University of Nice-Sophia Antipolis, 2005.

[2] B. Cohen, "Incentives build robustness in BitTorrent," in *Proc. 1st Workshop on the Econ. of Peer-2-Peer Systems, Berkeley, CA*, 2003.

[3] B. Fan, J. C. Lui, and D.-M. Chiu, "The Design Trade-Offs of BitTorrent-Like File Sharing Protocols," *IEEE/ACM Trans. on Networking*, vol. 17, no. 2, pp. 365–376, 2009.

[4] A. Ferragut, F. Kozynski, and F. Paganini, "Dynamics of content propagation in BitTorrent-like P2P file exchange systems," in *Proc. of 50th IEEE Conf. on Decision and Control, Orlando, FL*, 2011.

[5] A. Ferragut and F. Paganini, "Content dynamics in P2P networks from queueing and fluid perspectives," in *Proc. of 24th International Teletraffic Congress, Krakow, Poland*, 2012.

[6] M. Hirsch and H.L.Smith, "Monotone dynamical systems," in *Handbook of Differential Equations: Ordinary Differential Equations*, P. D. A. Cañada and A. Fonda, Eds. Elsevier, 2006, vol. 2, pp. 239–357.

[7] G. Kesidis, T. Konstantopoulos, and P. Sousi, "Modeling file-sharing with BitTorrent-like incentives," in *Proc. of IEEE Int. Conf. on Acoustics, Speech and Signal Processing, Honolulu, HI*, 2007.

[8] A. Legout, N. Liogkas, E. Kohler, and L. Zhang, "Clustering and sharing incentives in BitTorrent systems," in *ACM SIGMETRICS Performance Evaluation Review*, vol. 35, no. 1, 2007, pp. 301–312.

[9] D. Levin, K. LaCurts, N. Spring, and B. Bhattacharjee, "BitTorrent is an Auction: Analyzing and Improving BitTorrent's Incentives," *ACM SIGCOMM Computer Comm. Review*, vol. 38, pp. 243–254, 2008.

[10] F. Lo Piccolo, G. Neglia, and G. Bianchi, "The effect of heterogeneous link capacities in BitTorrent-like file sharing systems," in *Intl. Workshop on Hot Topics in Peer-to-Peer Systems*, 2004, pp. 40–47.

[11] L. Massoulié and M. Vojnović, "Coupon replication systems," *IEEE/ACM Transactions on Networking*, vol. 16, no. 3, pp. 603–616, 2008.

[12] F. Paganini and A. Ferragut, "PDE models for population and residual work applied to peer-to-peer networks," in *Proc. of 46th Annual Conference on Information Sciences and Systems, Princeton, NJ*, 2012.

[13] D. Qiu and W. Sang, "Global stability of Peer-to-Peer file sharing systems," *Computer Comm.*, vol. 31, no. 2, pp. 212–219, 2008.

[14] D. Qiu and R. Srikant, "Modeling and performance analysis of BitTorrent-like peer-to-peer networks," *ACM SIGCOMM Computer Communication Review*, vol. 34, no. 4, pp. 367–378, 2004.

[15] R. Sinkhorn, "Diagonal Equivalence to Matrices with Prescribed Row and Column Sums. II," *Proc. of American Mathematical Society*, vol. 45, no. 2, pp. 195–198, 1974.

[16] R. S. Varga, *Matrix iterative analysis*. NY: Springer-Verlag, 2000.

[17] F. Wu and L. Zhang, "Proportional response dynamics leads to market equilibrium," in *Proc. of the 39th Annual ACM Symposium on Theory of Computing, San Diego, CA*, 2007.

[18] X. Yang and G. De Veciana, "Service capacity of peer-to-peer networks," in *Proc. of IEEE Infocom, Hong Kong*, 2004.

[19] ——, "Fairness, incentives and performance in peer-to-peer networks," *Performance Evaluation*, vol. 63, no. 3, pp. 175–194, 2006.

[20] J. Zhu and B. Hajek, "Stability of a peer-to-peer communication system," *IEEE Trans. on Information Theory*, vol. 58, no. 7, pp. 4693–4713, 2012.