Monotonicity and global stability in download dynamics of content-sharing networks

Fernando Paganini and Andres Ferragut
Universidad ORT Uruguay

***Abstract*— This paper analyzes previously-proposed dynamic models of content-sharing networks from the point of view of global stability. Our focus is on models that track populations of participating peers as a function of the download progress achieved, described in previous research by a Partial Differential Equation where this progress is a fluid variable. We use such a model to identify conditions on the rate allocation function that make the dynamics preserve a suitable ordering of the state. This enables the application of tools of monotone dynamical systems. This is formally done with finite-dimensional, ordinary differential equation model in which the content fraction index is discrete. Our results apply both to the case of homogeneous upload bandwidths in the participating population, as well as the heterogeneous case with multiple populations of each bandwidth class.**

I. Introduction

In recent years, over-the-net content sharing has become one of the main sources of traffic on the global Internet. Different mechanisms have appeared to empower users to share content (files) mostly through peer-to-peer (P2P) connections that mutually exchange file pieces. Among them, the BitTorrent protocol [2] provides an important example, and we adopt its terminology: a subset of peers called *seeders* already possess the file and act as the servers to others, while *leechers* are the downloading clients. However the latter become, after downloading their first pieces, themselves servers to other peers, thus contributing their upload bandwidth during content download.

Such a *swarm* of peers is itself *dynamic*, as new peers arrive and others depart after completing the download. Since departures depend on the service capacity, which itself depends on the number of uploading peers, a *feedback* structure appears in the population dynamics. An early model that identifies this feedback is [14], using an ordinary differential equation (ODE); the model has roots in stochastic models of a similar nature [16]. Local [14] and global [13] stability results are derived for these ODE dynamics.

These initial models are coarse in the sense that the state does not track the peers' download progress. To obtain a finer description, in [4], [11] we proposed using a Partial Differential Equation (PDE) which discriminates the population of leechers as a function of the fraction of content file downloaded, expressed as a fluid variable. This intermediate resolution avoids the exponential complexity of tracking individual pieces (an alternative pursued in [7], [10], [18]), yet provides tighter predictions of the dynamics as compared to models with a single variable for the overall population. In particular we analyzed in [11] the resulting equilibrium, established its local stability through linearization, and showed through comparisons with BitTorrent simulations the increased predictive power of the model when analyzing variability and transient response.

All the above references work with a homogeneous population, in the sense that the upload bandwidth provided by each peer is the same. If there are multiple classes of peers with different parameters, an analysis of the resource allocation is required [3], [8]. It is found that while a set of homogeneous peers can be modeled through a *processor-sharing* allocation (where all leechers receive equal rate), service differentiation can and arguably must occur in the heterogeneous case. [9] suggests a *proportional* allocation, where every leecher receives from fellow leechers as much as it contributes to the network. In terms of the dynamics of such heterogeneous networks, a first ODE model for the two-class case was given in [1]. Recently in [12] we considered multi-class versions of both ODE and PDE models under proportional reciprocity; for the latter we generalized the equilibrium and local stability analysis.

What has been lacking for our PDE models with download progress is proof of *global* stability of the equilibrium, which is in general challenging since these are nonlinear dynamics with a distributed state. In this paper we recognize a key tool that enables such global results to be obtained, namely the *monotonicity* of trajectories with respect to a natural ordering. In particular we identify general conditions on the rate-allocation functions that give rise to such order preservation, as well boundedness of trajectories, from which stability theorems can be pursued using the theory of monotone dynamical systems [6]. Our conditions cover both the homogeneous and heterogeneous scenarios, and apply in particular to the processor sharing and proportional reciprocity models. While they are motivated very transparently with the PDE model, for our theory we will avoid the complexity of an infinite dimensional state and work with an ODE version of the dynamics where the content fraction variable is discrete.

The rest of the paper is organized as follows. In Section II we review some results on monotone dynamical systems, and our PDE models. In Section III we identify conditions on the rate allocation that yield a monotone, bounded dynamics for the PDE. We then turn to a finite-dimensional version in Section IV to give precise results on monotonicity and global stability. In Section V we extend these results to the multi-class setting. Conclusions are given in Section VI.

Research supported by AFOSR-US under Grant FA9550-12-1-0398.
E-mail:paganini@ort.edu.uy.

II. Background

A. Monotone systems

The theory of monotone dynamical systems refers to dynamics on a set X in which a partial ordering relationship has been established. An extensive recent monograph on this topic is [6], we extract from it some elements to be applied in the present paper.

Given a Banach space Y and a closed, convex cone $Y_+ \subset Y$, an ordering can be defined by $x \le x' \Longleftrightarrow x' - x \in Y_+$; Strict ordering $x \ll x'$ means $x' - x \in \text{Int}(Y_+)$, the interior of the cone, assumed non-empty. A *semiflow* $\Phi_t : X \to X$ for $t \ge 0$ on a set $X \subset Y$, is monotone or order preserving if $x(0) \le x'(0)$ implies $x(t) = \Phi_t(x(0)) \le x'(t) = \Phi_t(x'(0))$ for all t. Strong monotonicity means the ordering becomes strict for $t > 0$.

Strong monotonicity greatly narrows down the possibilities for such dynamics; for instance, there can be no stable periodic orbits. And, for semiflows with orbits of compact closure, most initial conditions will lead to convergent trajectories. If, further, for an invariant region one can establish there is a unique equilibrium, it must attract the entire region. We state the following result from [6][1].

Proposition 1 ([6], Corollary 1.20): Let Y be an ordered Banach space, and Φ a strongly monotone semiflow on $X \subset Y$ with orbits of compact closure. If X is open and contains a single equilibrium point p, all trajectories in X converge to p.

For the case $X \subset \mathbb{R}^n$, a natural ordering to be used in this paper is the one defined by $Y_+ = \mathbb{R}^n_+$; namely, $x \le x'$ iff $x_i \le x'_i$ for all i, and $x \ll x'$ when all inequalities are strict. Suppose the semiflow Φ is given by solutions to a smooth differential equation $\dot{x} = f(x)$ defined on X. Chapter 3 of [6] analyzes monotonicity in this situation, in terms of the system Jacobian. In particular, if the matrix $A(x) := \frac{\partial f}{\partial x}$ is *Metzler* (this means $a_{ij} \ge 0$ for $i \ne j$) at any x, then the dynamics are called *cooperative*, and the corresponding flow is monotone. Strong monotonicity requires in addition that the matrix $A(x)$ be *irreducible*: this means there is no permutation of variables where the matrix becomes block triangular. Equivalently, if one considers a graph with vertices in $\{1, \dots, n\}$ and a directed edge from i to j whenever $a_{ij} > 0$, the resulting graph must be strongly connected (there is a directed path between any two nodes).

B. Population and download dynamics in p2p file-sharing

As discussed before, content-sharing is implemented by the bidirectional exchange of small pieces of a certain file between peers. In this paper we focus on the scenario where the population of seeders (who own the entire content) is fixed at y_0, whereas leechers interested in the content form a variable population $x(t)$: they arrive at the swarm at a certain rate, stay as long as necessary to complete the download, and then immediately leave the system. This simplified scenario is nevertheless common in practice.

[1] In fact Corollary 1.20 from [6] is a stronger version of this statement.

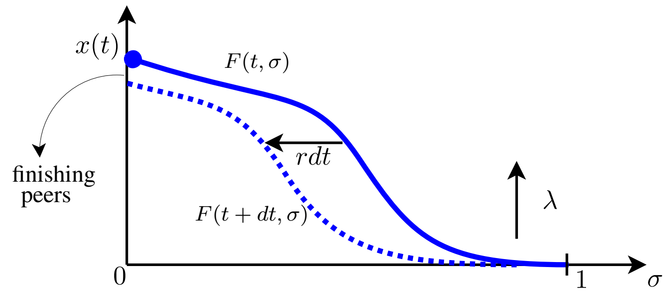

Fig. 1. Fluid state evolution.

To develop a clean and intuitive picture of the dynamics, in [4], [11] it was found convenient to use a fluid variable $\sigma \in [0, 1]$ for the download fraction, and express the resulting dynamics by means of a partial differential equation (PDE), as is now reviewed. We consider first the single class case, where all leechers are of the same kind, see Section V for the multi-class version.

Define the real-valued variable $F(t, \sigma)$ that represents (in fluid terms) the population of leechers that, at time t, have pending download of at least σ. Thus $F(t, \sigma)$, non-increasing in σ, acts as the complementary cumulative distribution of the leecher population, with $F(t, 0) = x(t)$, the total leecher count and $F(t, 1) \equiv 0$. Figure 1 shows an instance of such fluid download profiles. The dynamic model from [11] is:

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

Here λ is the arrival rate of new leechers, which we assume arrive with no prior content[2]; this provides an upward drift to the entire population distribution. $r(F, \sigma) \ge 0$ represents the download rate per peer, in units of files per second, which regulates the speed at which the function $F(t, \sigma)$ is transported in the direction of $\sigma = 0$ as download progresses. The notation expresses the dependence on the function state F, and a possible differentiation across σ; state-dependence arises because download speeds are determined by upload speeds of others; this is the cause of nonlinear feedback present in these dynamics. Specifically, the overall upload capacity of the set of peers is

$$R_{up} = \mu(x + y_0), \tag{2}$$

where the parameter μ represents the individual upload rate of each peer (assuming for now homogeneity across peers). This upload capacity must necessarily be no smaller than the aggregate download rate, integrated over the population:

$$R_{down} := \int_0^1 r(F, \sigma) \left[-\frac{\partial F}{\partial \sigma} \right] d\sigma; \tag{3}$$

(the minus sign arises due to F being a complementary CDF). The functional $r(F, \sigma)$ that maps population profiles F in $[0, 1]$ to rate profiles specifies how the download

[2] This is done for simplicity, our results of this paper generalize to the case considered in [11] where leechers have a distribution $H(\sigma)$ of content to download upon arrival.

bandwidth is distributed among the population, allowing for differentiation (at most) across σ. A simple yet important case for the rate functional is

$$r(F,\sigma) = \frac{R_{up}}{x} = \frac{\mu(x+y_0)}{x} \quad \forall\sigma. \tag{4}$$

This is a *processor sharing* allocation, i.e. the upload bandwidth is efficiently and uniformly distributed among leechers, independently of their download stage. In particular, $R_{down} = R_{up}$ and this implies there are no other bottlenecks. Empirical evidence indicates that this is quite an accurate model for upload-constrained BitTorrent systems under the homogeneity assumption (cf. the discussion in [5], [11]). Other allocations have been analyzed in [4].

III. Monotonicity and stability in the PDE model

In this section we will identify conditions for which the file-sharing dynamics can exhibit monotonicity, bounded trajectories, and a unique equilibrium. For this task the PDE model provides a clear intuition, which we present here, somewhat informally; in the next section we prove rigorous results with a discretized ODE version.

A. Monotonicity

As a natural order in the space of trajectories we consider pointwise inequality[3].

$$F(t,\cdot) \geq \tilde{F}(t,\cdot) \iff F(t,\sigma) \geq \tilde{F}(t,\sigma) \ \forall\sigma \in [0,1]. \tag{5}$$

One such pair of profiles is depicted in Figure 2. We now identify a condition on the rate functional to guarantee this order is preserved.

Assumption 1: r is *decreasing* with the respect to the pointwise order:

$$F \geq \tilde{F} \Longrightarrow r(F,\sigma) \leq r(\tilde{F},\sigma) \ \forall\sigma.$$

We argue hat this assumption ensures the monotonicity of the flow with respect to the pointwise order. In reference to Fig. 2: Note that both curves $F \geq \tilde{F}$ are subject to the same upward drift due to the arrivals term in (1), but the transport term is in general different, because rate depends on population. The condition $r \leq \tilde{r}$ implies the resulting drift vectors for each of the two states must be as depicted in the figure, with the upper curve having a more upward direction, thus locally preserving the inequality $F \geq \tilde{F}$.

B. Boundedness

The second issue for a global stability result is to ensure boundedness of trajectories; this turns out to hold provided the file sharing achieves at least a minimal level of efficiency, as is now specified.

Assumption 2: The rate functional $r(F,\sigma)$ is κ-efficient for some $0 < \kappa \leq 1$. We define this to mean that $R_{down} > \kappa\mu x$, i.e. the total download rate exceeds a positive fraction κ of the leecher upload bandwidth.

[3] This corresponds to the cone of positive functions; we will not be precise about the Banach space in question, since our theory will be developed in finite dimensions.

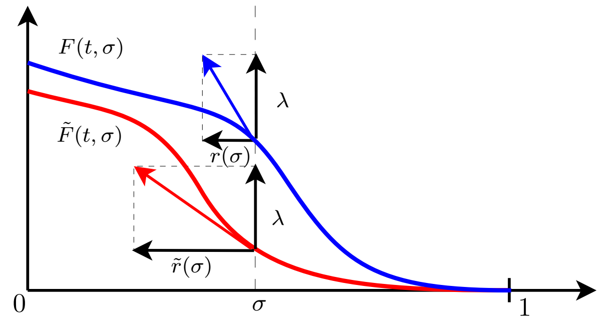

Fig. 2. Monotonicity property of the PDE dynamics

To show this assumption imposes a bound on trajectories, define the *unfinished work* of a profile $F(\sigma)$ by

$$u(F) = \int_0^1 \sigma \left[-\frac{\partial F}{\partial \sigma}\right] d\sigma = \int_0^1 F(\sigma) d\sigma, \tag{6}$$

the last identity following from integration by parts.

We claim that the set $X_K = \{F : u(F) < K\}$ is invariant under the dynamics for $K \geq \frac{\lambda}{\kappa\mu}$. To see this integrate (1) over $\sigma \in [0,1]$ to obtain

$$\dot{u} = \lambda - R_{down} < \lambda - \kappa\mu x, \tag{7}$$

where we have used κ-efficiency. If, starting in X_K, a trajectory were to reach the boundary $u(t) = K$, this implies $x(t) \geq \frac{\lambda}{\kappa\mu}$ (note $x(t) \geq u(t)$ since $F(t,\cdot)$ is decreasing), and thus from (7) that $\dot{u} < 0$, in contradiction with u reaching the boundary from below. This establishes the claim.

C. Equilibria and stability

The remaining question is to characterize the set of equilibria. To make the problem interesting we will assume henceforth that $\lambda > \mu y_0$; this means that the seeder upload rate alone is insufficient to sustain the load, so any equilibrium will necessarily have positive population of leechers. In the perfectly efficient case with $R_{down} = \mu(x+y_0)$ the population would be $x^* = F^*(0) = \frac{\lambda}{\mu} - y_0$, as we can see setting $\dot{u} = 0$ in (7). In a less efficient situation the population would be larger.

Even if x^* is determined, there could in general be multiple profiles $F^*(\sigma)$ that set to zero the right-hand side of (1). For certain rate functionals we can further show that this equilibrium profile is unique. In particular, the processor sharing rate functional (4) is decreasing and perfectly efficient, and has a unique equilibrium. This follows from the fact that $r(F,\sigma)$ in (4) does not depend on σ, so the equilibrium condition is

$$\frac{\partial F^*}{\partial \sigma} = -\frac{\lambda}{r^*},$$

constant in σ. Using boundary conditions the unique equilibrium is the linear function $F^*(\sigma) = x^*(1-\sigma)$ (uniform distribution) with $x^* = \lambda/\mu - y_0$.

Other rate functionals also meet the uniqueness test. Assuming this happens, we have the ingredients for a global stability result using monotone systems. Pursuing this result precisely in the realm of PDEs requires defining an appropriate Banach space, with compactness of orbit closures, and refining inequalities to be strict in the appropriate sense, for Proposition 1 to be applicable. This is presently out of our scope, we will instead turn to a finite-dimensional setting.

IV. Monotonicity and global stability for a discretized ODE model

To avoid technicalities associated with PDEs we now formulate an ordinary differential equation model, discretizing the σ variable in M points, defining the states

$$z_j(t) = F\left(t, j/M\right), \quad j = 0, \ldots, M-1.$$

Through the approximation

$$\left.\frac{\partial F}{\partial \sigma}\right|_{\sigma=\frac{j}{M}} \approx M(z_{j+1} - z_j),$$

the earlier model (1) indicates the dynamics

$$\dot{z}_j = \lambda + M r_j(z)(z_{j+1} - z_j), \quad j = 0, \ldots, M-1, \tag{8}$$

where by convention $z_M = 0$. The state is defined in the cone of vectors with decreasing components

$$\mathcal{Z} := \left\{ z \in \mathbb{R}^M : z_0 \geq z_1 \geq \cdots \geq z_{M-1} \geq 0 \right\}. \tag{9}$$

Here the function $r_j(z) \geq 0$ provides the download rate for peers at the j-th download stage; it could vary with j, as well as depend on the entire profile z. The download constraint is now written as

$$\sum_{j=1}^{M} r_j(z)(z_j - z_{j+1}) \leq \mu(z_0 + y_0).$$

Remark 1: Of course, file content could be quantified discretely from the start, writing directly (8) without reference to any PDE. We note however that:

(i) The fluid PDE model provides a very intuitive description of the dynamics, with pictures as in Figure 2 which helped us reason about monotonicity.

(ii) The PDE suggests the natural state variables to use. Namely z_j is the accumulated population of leechers with pending download larger than j/M. This will turn out to be a simpler coordinate system for monotonicity studies, as compared to the non-accumulated population quantities considered in [4].

As an additional comment, we note the distinction between the above ODE and population models in [14], which do not discriminate by download stage and would thus correspond to the case $M = 1$, a significantly coarser model.

Introduce now the interior of the cone $\mathcal{Z}$,

$$\text{Int } \mathcal{Z} := \left\{ z \in \mathbb{R}^M : z_0 > z_1 > \cdots > z_{M-1} > 0 \right\}. \tag{10}$$

Int $\mathcal{Z}$ is invariant under (8) provided that $r_j(z) > 0$ (strictly) on Int $\mathcal{Z}$. To see this start from an initial condition in Int $\mathcal{Z}$, assume that a certain z_j later reaches zero; then so must z_l, $l \geq j$ and thus (8) gives $\dot{z}_j = \lambda > 0$, a contradiction with z_j reaching zero. Assume instead that $z_j - z_{j+1}$ reaches zero for $j \leq M-2$ at some point in time, with positive z_j; by redefining j if necessary we can suppose $z_j = z_{j+1} > z_{j+2}$, which implies

$$\dot{z}_j - \dot{z}_{j+1} = -M r_{j+1}(z)(z_{j+2} - z_{j+1}) > 0,$$

a contradiction with $z_j - z_{j+1}$ reaching zero from above.

The assumptions on the rate functions $r_j(z)$ are now stated, using as guideline the PDE counterparts:

1) $r_j(z)$ is decreasing with respect to the componentwise ordering in z. Assuming it is smooth, this translates to the condition

$$\frac{\partial r_j}{\partial z_k} \leq 0 \quad \forall j, k. \tag{11}$$

1') We will also impose a stronger requirement on derivatives with respect to the total population $z_0 = x$:

$$\frac{\partial r_j}{\partial z_0} < 0 \quad \forall j, z \in \text{Int } \mathcal{Z}. \tag{12}$$

2) $r_j(z)$ is κ-efficient:

$$\sum_j r_j(z)(z_{j+1} - z_j) > \kappa \mu z_0. \tag{13}$$

A. Monotonicity

We first prove that assumption (11) implies the flow is monotone, through the system Jacobian. We denote by $g(z) = (g_j(z))_{j=0}^{M-1}$ the vector field in (8).

Proposition 2: Assume (11) holds. For $z \in$ Int $\mathcal{Z}$ the Jacobian matrix $A(z) = \frac{\partial g}{\partial z}$ is Metzler, i.e.

$$\frac{\partial g_j}{\partial z_k} \geq 0 \text{ for } k \neq j. \tag{14}$$

If in addition (12) holds, then $A(z)$ is irreducible in Int $\mathcal{Z}$.

Proof: For $k \neq j$, $k \neq j+1$ we have

$$\frac{\partial g_j}{\partial z_k} = M \frac{\partial r_j}{\partial z_k}(z) \cdot (z_{j+1} - z_j) \geq 0.$$

Since both factors are non-positive in $\mathcal{Z}$ due to (11), then (14) follows. For $k = j+1 < M$, we have

$$\frac{\partial g_j}{\partial z_k} = M r_j(z) + M \frac{\partial r_j}{\partial z_k}(z) \cdot (z_{j+1} - z_j) \geq 0,$$

the additional term being non-negative as well. We conclude that the Jacobian is Metzler as claimed.

Assuming now an interior point where $\frac{\partial r_j}{\partial z_0} < 0$, note that this implies $r_j(z) > 0$ (otherwise r_j would become negative in a neighborhood), thus $\frac{\partial g_j}{\partial z_{j+1}} > 0$. Also, $z \in$ Int $\mathcal{Z}$ implies

$$\frac{\partial g_j}{\partial z_0} = M \frac{\partial r_j}{\partial z_0}(z) \cdot (z_{j+1} - z_j) > 0 \quad \forall j > 0.$$

We see then that the elements a_{j0} for $j > 0$, and $a_{j,j+1}$ for $j < M-1$ are strictly positive. This implies that the directed graph associated with $A(z)$ is strongly connected: one can move up sequentially in j and back to 0 with positive transitions. Hence $A(z)$ is irreducible. ■

Invoking the material on Section II-A, the dynamics produce a strongly monotone flow.

B. Boundedness

Let us now show that the second assumption on efficiency ensures boundedness of trajectories.

Proposition 3: Consider the dynamics (8). Under assumption (13), the set $X_K = \left\{ z \in \text{Int } \mathcal{Z} : \sum_j z_j < K \right\}$ for $K \geq M \frac{\lambda}{\kappa\mu}$ is positively invariant under the flow.

Proof: From the hypothesis we obtain

$$\sum_j \dot{z}_j = M\lambda - M \sum_j r_j(z)(z_{j+1} - z_j) < M(\lambda - \kappa\mu z_0). \quad (15)$$

If, starting from X_K, a trajectory reaches the boundary, it must be that $\sum_j z_j = K$, since we already established the invariance of Int $\mathcal{Z}$. Now since z is a decreasing vector we must have $z_0 \geq \frac{K}{M} \geq \frac{\lambda}{\kappa\mu}$. But this makes the left-hand side of (15) negative, a contradiction since $\sum_j z_j$ approached the threshold from below. This establishes the invariance. ∎

Note also the above argument implies there can be no equilibrium points outside such X_K.

C. Global stability

We complete the analysis by stating a global stability result for the case where the equilibrium is unique.

Theorem 4: Suppose the rate functions $r_j(z)$ satisfy (11), (12), and (13). If in addition there is a unique $z^* \in$ Int $\mathcal{Z}$ such that

$$\lambda = M r_j(z^*)(z_j^* - z_{j+1}^*), \quad j = 0, \dots, M-1,$$

then this point is a global attractor of the dynamics (8).

Proof: For any initial condition we can choose an invariant set X_K that contains it, and restrict our attention to the dynamics on this bounded open set, which must contain the equilibrium. We have a strictly monotone flow with orbits of compact closure, with a single equilibrium point in X_K. Proposition 1 implies there is global convergence to equilibrium. ∎

As a particular case, we obtain global stability for the processor sharing rate function introduced in (4).

Corollary 5: The system (8) with the rate functions

$$r_j(z) = r(z_0) = \mu \frac{z_0 + y_0}{z_0}, \quad j = 0, \dots, M-1,$$

has the globally attracting equilibrium

$$z_j^* = \left(\frac{\lambda}{\mu} - y_0 \right) \left(1 - \frac{j}{M} \right). \quad (16)$$

Proof: Clearly $r_j(z)$ is decreasing and in particular

$$\frac{\partial r_j}{\partial z_0} = -\frac{\mu y_0}{(z_0)^2} < 0;$$

it is also perfectly efficient ($R_{down} = R_{up} = \mu(z_0 + y_0)$); this implies as in (15) that $\sum_j \dot{z}_j = \lambda - \mu y_0 - \mu z_0$, from which there is a unique possible z_0^* in equilibrium. Looking now at the individual derivatives in (8) we find at equilibrium

$$\lambda + M r(z_0^*)(z_{j+1}^* - z_j^*) = 0$$

from where all partial increments are $z_{j+1}^* - z_j^*$ must be equal, yielding the formula in (16). In particular the equilibrium is unique, so Theorem 4 can be invoked. ∎

V. Extension to heterogeneous networks

The models discussed so far apply to the situation where all peers have a common bandwidth access parameter μ. When this parameter is allowed to vary across peers, resource allocation does not correspond to processor sharing; instead, service differentiation can arise due to the reciprocity mechanisms of p2p. In this section we study this more general case with a multi-class model, where peers in each of n classes have upload bandwidths $\{\mu^i\}_{i=1}^n$, and arrival rate λ^i. Seeders are fixed and have an upload rate μ^0.

We will assume each class has a population profile as a function of downloaded content. Adopting initially a continuous variable σ to represent file fraction, let $F^i(t, \sigma)$ represent the population of leechers of class i with pending download larger than σ. Its dynamics has the form

$$\frac{\partial F^i}{\partial t} = \lambda^i + r^i(F, \sigma) \frac{\partial F^i}{\partial \sigma}, \quad \sigma \in [0, 1]. \quad (17)$$

Here the overall state $F(t, \sigma)$ is the vector of profile functions $(F^i(t, \sigma))_{i=1}^n$, and the download rate functional r^i depends in general on the entire state across all classes.

Our state ordering (5) extends to this case as follows:

$$F \geq \tilde{F} \Leftrightarrow F^i(\sigma) \geq \tilde{F}^i(\sigma) \; \forall \sigma \in [0, 1], \; i = 1, \dots, n. \quad (18)$$

The following are the appropriate generalizations of the assumptions considered in Section III for the rate functional $r(F, \sigma) = (r^i(F, \sigma))_{i=1}^n$:

1) r is *decreasing* with the respect to the pointwise order:

$$F \geq \tilde{F} \Longrightarrow r(F, \sigma) \leq r(\tilde{F}, \sigma) \; \forall \sigma.$$

2) r is *κ-efficient* in each component, for $0 < \kappa \leq 1$:

$$R_{down}^i = \int_0^1 r^i(F, \sigma) \left[-\frac{\partial F^i}{\partial \sigma} \right] d\sigma > \kappa \mu^i x^i.$$

This means that class i receives at least a fraction κ of what it contributes to the network, which in fact embeds some level of fairness in the allocation, as opposed to just efficiency of total bandwidth use.

It can be argued, analogously to Section III, that these conditions ensure respectively, monotonicity and boundedness for the flow in (17). Once again, under uniqueness of the equilibrium we should have global stability. Below we state precise results with a finite dimensional model.

A. Global stability for a multi-class ODE

For each class i define M state variables $z_j^i, j = 0, \dots M-1$, corresponding to the accumulated populations $F^i(t, \frac{j}{M})$. The overall state z belongs to

$$\mathcal{Z}^n := \left\{ z = (z_j^i) \in \mathbb{R}^{nM} : z_0^i \geq \dots \geq z_{M-1}^i \geq 0 \; \forall i \right\}; \quad (19)$$

when thought as a column vector, the variables of z are ordered lexicographically in (i, j). Int $\mathcal{Z}^n$ is analogous, with strict inequalities. The dynamics is

$$\dot{z}_j^i = \lambda^i + M r_j^i(z)(z_{j+1}^i - z_j^i), \quad \begin{array}{l} i = 1, \dots, n; \\ j = 0, \dots, M-1. \end{array} \quad (20)$$

We lay out the assumptions on the rate functions $r_j^i(z)$.

1) $r_j^i(z)$ is decreasing for componentwise ordering:

$$\frac{\partial r_j^i}{\partial z_k^l} \leq 0 \quad \forall (i,j),(l,k). \tag{21}$$

1') Rates are strictly decreasing with respect to total class populations:

$$\frac{\partial r_j^i}{\partial z_0^l} < 0 \quad \forall (i,j), l, \quad \forall z \in \text{Int } \mathcal{Z}^n. \tag{22}$$

2) $r_j^i(z)$ is κ-efficient in each class:

$$R_{down}^i = \sum_j r_j^i(z)(z_{j+1}^i - z_j^i) > \kappa \mu^i z_0^i. \tag{23}$$

We denote by $g(z) = (g_j^i(z))$ the overall vector field, again using lexicographical order in i,j. We now state the corresponding results extending the work in Section IV. Proofs are analogous and omitted due

Proposition 6 (Montonicity): Assume (21) holds. For $z \in \mathcal{Z}^n$ the Jacobian matrix $A(z) = \frac{\partial g}{\partial z}$ is Metzler, i.e.

$$\frac{\partial g_j^i}{\partial z_k^l} \geq 0 \text{ for } (l,k) \neq (i,j). \tag{24}$$

If in addition (22) holds, then $A(z)$ is irreducible in Int $\mathcal{Z}^n$.

Proposition 7 (Boundedness): Consider the dynamics (20). Let $K = (K^i)_{i=1}^n, \ K^i \geq \frac{\lambda^i}{\kappa \mu^i}$. Under assumption (23), the set $X_K^n = \left\{ z \in \text{Int } \mathcal{Z} : \sum_j z_j^i < K^i, \ i = 1, \ldots, n \right\}$ is positively invariant under the flow.

Theorem 8 (Global stability): Suppose the rate functions $r_j^i(z)$ satisfy (21), (22), and (23). If in addition there is a unique $z^* \in \text{Int } \mathcal{Z}^n$ such that

$$\lambda^i = M r_j^i(z^*)(z_j^{*i} - z_{j+1}^{*i}), \qquad \begin{array}{l} i = 1, \ldots, n; \\ j = 0, \ldots, M-1. \end{array} \tag{25}$$

then this point is a global attractor of the dynamics (20).

B. Application to the proportional reciprocity rate function

We apply the theory to an important example of multi-class rate allocation, the *proportional reciprocity* model:

$$r_j^i(z) = \frac{\mu^0 y_0}{\sum_{l=1}^n z_0^l} + \mu^i \tag{26}$$

Here the seeder capacity is equally distributed among all downloading peers, but leechers receive from other leechers exactly as much as they give. This reciprocity proposal has been discussed in many references [3], [9], [12], [15], [17], and shown to approximately represent the tit-for-tat component of BitTorrent's file-sharing. (26) is monotonically decreasing and satisfies κ-efficiency with $\kappa = 1$.

Proposition 9: Consider the system (20) with the rate function (26) for each i, with $j = 0, \ldots, M-1$.

Suppose that $\sum_i \lambda^i > \mu^0 y_0$ (meaning that seeders alone cannot cope with the service demands). Then there exists a unique, globally attracting equilibrium of the form

$$z_j^{*i} = \left(\frac{\lambda^i}{\mu^i + \alpha} \right) \left(1 - \frac{j}{M} \right), \qquad \begin{array}{l} i = 1, \ldots, n; \\ j = 0, \ldots, M-1, \end{array}$$

where $\alpha > 0$ is the unique solution to $\sum_i \lambda_i \frac{\alpha}{\mu_i + \alpha} = \mu^0 y_0$.

The main step of the proof is establishing existence and uniqueness of equilibrium; this was done in [12] for the case of a *single* population state per class ($M = 1$), and can easily be extended to the present case.

VI. Conclusions

In this paper we obtained global stability results for the dynamics of content-sharing networks. These apply to resource allocation policies which assign decreasing resources per peer at any download stage or class when populations increase, with a minimum amount of efficiency in the allocation for each class. Under these conditions we can invoke the powerful theory of monotone dynamical systems to prove global stability. In particular the question is settled for models (processor sharing, proportional allocation) commonly applied to respectively homogeneous and heterogeneous p2p networks; in future work we will study other file-sharing policies from this perspective.

References

[1] F. Clévenot and P. Nain, "A simple model for the analysis of SQUIRREL," in *Proc. of IEEE Infocom, Hong Kong*, 2004.

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

[11] 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.

[12] F. Paganini, A. Ferragut, and M. Zubeldía, "Dynamics of heterogeneous peer-to-peer networks," in *Proc. of 52th IEEE Conf. on Decision and Control, Firenze, Italy*, 2013.

[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] 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.

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

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

[18] 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.