Fluid Limits Applied to Peer to Peer Network Analysis

Laura Aspirot
Universidad de la República
Montevideo, Uruguay
aspirot@fing.edu.uy

Ernesto Mordecki
Universidad de la República
Montevideo, Uruguay
mordecki@cmat.edu.uy

Gerardo Rubino
INRIA Rennes - Bretagne Atlantique
Rennes, France
gerardo.rubino@inria.fr

Abstract—The objective of several techniques including fluid limits and mean field approximations is to analyze a stochastic complex system (e.g. Markovian) studying a simplified model (deterministic, represented by ordinary differential equations (ODEs)). In this paper, we explore models proposed for the analysis of BitTorrent P2P systems and we provide the arguments to justify the passage from the stochastic process, under adequate scaling, to a fluid approximation driven by an ODE. We also make the link between the stationary regime of the stochastic models and the fixed points of the associated ODEs. Finally, we analyze the asymptotic distribution of the scaled process.

Index Terms—fluid limits, mean fields, BitTorrent

I. INTRODUCTION

There are several examples of complex stochastic systems for which analytical expressions cannot be derived, or that are even difficult to simulate. However, in many cases they can be studied much more easily by analyzing deterministic systems, obtained as asymptotic approximations of the original ones. The complexity of the system may be due to its size, its dependence structure, etc. There are many such systems, as for instance TCP connections, wireless systems, or peer to peer networks. In the case of TCP connections sharing a bottleneck, wireless users sharing a channel, or in the peer to peer case, there is a common resource and many individuals. The behavior of each one depends on the state of the whole system, introducing dependence between all the individuals.

Starting from a stochastic model the objective is to find a deterministic approximation for the original process. This introduces the problem of finding the suitable scale for this approximation. For example, classical results in queuing theory consider a sequence of stochastic processes indexed by an integer N, where some key state variable appears divided by N, and the time variable is multiplied ("accelerated") by the same factor, obtaining a deterministic limit when N goes to infinity. In addition, in other areas such as in biology, or in the analysis of epidemic phenomena, a typical scaling consists in dividing by N, and in considering transition rates increasing with N (jumps are of order $1/N$ and transition rates of order N, that means that the product remains "constant" as N increases). For a survey about this topic see [1], and for a more general reference about limits of stochastic processes we suggest [2]. In this paper we follow mostly the approach of [2] and [3].

Let us describe another way to approximate complex stochastic systems by deterministic ones. This approach is called mean field approximations. This technique comes from physics, where it is used to study systems with a large number of interacting particles. When the number of particles increases each particle behaves as if it were under the action of a global force (the mean field). Applications to telecommunications appeared in the literature and were widely developed in the last decade. There are many works, considering different types of phenomena and different types of models (discrete or continuous time, discrete or continuous state space, etc.). For instance [4], [5], [6], [7] and the references therein cover a wide range of techniques and applications. More recently mean field methods have been applied to game theory and optimal control (see for example [8], [9] and references therein).

In mean field approximations we can distinguish two steps: the first one is focused on the occupation measure limit (i.e. the asymptotic proportion of individuals in each state) and the second one is focused on the decoupling assumption (asymptotically the state of each individual is independent from the others). A very frequent approach to the first step consists in proving the limit with the same techniques as in the fluid limits case. The proof of the asymptotic independence relies in different tools.

One of the main results, both in the case of fluid limits or mean fields, when a stochastic system is approximated by one modeled by an ODE, is that in some cases the stationary regime of the former can be analyzed by studying the ODE's fixed points. There are many issues on this topic discussed in [4], [5], [10].

Our object of study is the use of fluid limits for modeling peer to peer systems. In the literature we can find works on peer to peer systems using stochastic models [11], fluid models [12], [13] and fluid limits or mean field approximations [14], [15], [16]).

In this work we consider a fluid limit model for a BitTorrent network based on [12], [13]. In both papers the deterministic model is the starting point of the analysis. We, on the other hand, start from a stochastic one and justify the passage from one to the other.

The contributions of this work consist first in the mathematical justification that the deterministic fluid models in [12] and [13] are fluid limits of stochastic models, that we present.

En each case we define a stochastic model and construct a sequence of stochastic processes such that, under adequate scaling, converges to a deterministic model driven by an ODE. The second contribution is that we prove the existence of a stationary regime for each process in the sequence and then we prove that the sequence of processes in stationary regime converges to the ODE's fixed point. We finally describe the asymptotic distribution of the stochastic process. We prove that the difference between the scaled process and the deterministic one can be approximated by a gaussian process.

The remainder of the paper is structured as follows. In Section II first we present some well known models and then we describe our model. In Section III we provide our results about the approximation by a deterministic process, the existence and convergence of stationary regime and the asymptotic distribution. In Section IV we conclude this work.

II. Model

In this section we give a brief description of BitTorrent. Then we consider three BitTorrent models from the literature: a stochastic model in [11], and two fluid models in [12] and [13], in subsection II-A. At the end we state our stochastic models in subsection II-B.

BitTorrent is a peer to peer protocol, for file sharing over a network. BitTorrent divides the target file into small files (*chunks*). Each peer connects to others and downloads simultaneously different chunks. There are two types of peers: *leechers* and *seeds*. Leechers download parts of the file from other peers and upload parts of the file for other leechers. Seeds have all the file and only remain in the system to help leechers to get missing file parts (they are *altruist* nodes). We do not detail here the peer selection policy (see for example [12]) and other features that help in understanding the behavior, for example based on traffic measures (see [17], [18], [19]).

A. Stochastic and fluid models for BitTorrent in the literature

We first describe the stochastic model proposed by Yang and de Veciana in [11], that is the motivation for the fluid model in [12]. In [11] the BitTorrent network is described using a branching process for the transient regime and a Markov model for the stationary regime. For the Markov model, the following parameters are considered:

- $X(t)$: number of leechers at time t,
- $Y(t)$: number of seeds at time t,
- λ: arrival rate (Poisson) of peers,
- μ: uploading rate for each peer,
- γ: leaving rate for seeds;

with the following transition rates:

- $q((x,y),(x+1,y)) = \lambda$ (arrival of a new peer),
- $q((x,y),(x-1,y+1)) = \mu(x+y)$ (a leecher successfully finishes downloading the file),
- $q((x,y),(x,y-1)) = \gamma y$ (a seed leaves the network).

For $(0,y)$ there is no possible transition in the direction $q((x,y),(x-1,y+1))$ and the remaining rates are the same

as before. In [11] the stationary distribution is computed numerically.

Now we describe the fluid model proposed by Qiu and Srikant in [12]. A BitTorrent system is analyzed, using different tools. One of the approaches is the fluid description based on the stochastic model of [11]. The fluid model also considers two aspects that are not discussed in [11]: the first one is that leechers may leave the system before finishing their download and the second one is that capacity restriction, related to the time needed to finish a download, may be in the uploading capacity of peers (as in [11]) but also in the downloading capacity. The fluid model is stated as follows:

- $x(t)$: number of leechers at time t,
- $y(t)$: number of seeds at time t,
- λ: arrival rate (Poisson) of peers,
- μ: uploading rate for each peer,
- c: downloading rate for each peer,
- θ: leaving rate for leechers,
- γ: leaving rate for seeds,
- $\eta \in [0,1]$: efficiency factor, that takes into account the efficiency of the file sharing mechanism ([12] provides a detailed analysis of η),

The maximal total uploading rate is $\mu(\eta x + y)$, the maximal total downloading rate is cx, and the restriction may be in the upload or in the download. The effective downloading rate is thus $\min(cx, \mu(\eta x + y))$. The evolution of the number of leechers and seeds is described by the following ODE:

$$\begin{cases} x' = \lambda - \min(cx, \mu(\eta x + y)) - \theta x, \\ y' = \min(cx, \mu(\eta x + y)) - \gamma y. \end{cases} \quad (1)$$

There is a line $y = (c/\mu - \eta)\, x$ where the behavior of the system changes because of the term $\min(cx, \mu(\eta x + y))$, dividing the state space in two zones. The authors state that the average number of leechers and seeds in stationary regime are the values of the ODE's fixed point (x^*, y^*) and derive the average downloading time from an approximation of Little's law. They also show a good fitting with simulations of the BitTorrent protocol and with real traces, specially when the arrival rate λ is high.

Based on [12], Rivero and Rubino in [13] consider a fluid model for a BitTorrent network with different classes of peers. There are two classes of leechers: high tolerance leechers and low tolerance ones. The parameters are the following:

- $x_a(t)$: number of high tolerance leechers at time t,
- $x_b(t)$: number of low tolerance leechers at time t,
- $y(t)$: number of seeds at time t,
- λ_a: arrival rate of high tolerance leechers,
- λ_b: arrival rate of low tolerance leechers,
- μ: uploading rate for each peer,
- c: downloading rate for each peer,
- θ_a: leaving rate for high tolerance leechers,
- θ_b: leaving rate for low tolerance leechers, with $\theta_b > \theta_a$,
- γ: leaving rate for seeds,
- the efficiency factor is $\eta = 1$,

The fluid model for this system is:

$$\begin{cases} x_a' & = \lambda_a - \theta_a x_a - u_a, \\ x_b' & = \lambda_b - \theta_b x_b - u_b, \\ y' & = u_a + u_b - \gamma y, \end{cases} \qquad (2)$$

where

$$\begin{aligned} u_a & = \min\left(c x_a, \mu(x_a + x_b + y)\frac{x_a}{x}\right), \\ u_b & = \min\left(c x_b, \mu(x_a + x_b + y)\frac{x_b}{x}\right). \end{aligned}$$

From this equation there are also two different zones, divided by a plane, again due to the restriction in uploading and downloading capacity. A strategy to improve performance by giving priority to peers that will probably stay more time in the system as seeds, specially in bad resource conditions, is studied. In order to define the policy, the space is divided by planes in three zones, according to the capacity. For each zone a server policy is defined (giving priority for high tolerance leechers when capacity is not enough). A fluid model considering the different policies in each zone is stated. The study of fixed points allows to analyze the priority policy, compared with the non priority one.

In this paper we consider stochastic and fluid models that allow to analyze the number of peers in the system. Another approach to study BitTorrent systems is to consider the number and type of chunks that each peer possesses [14], [15], [16]. In particular they study the asymptotic behavior of a Markov process that converges to the solution of an ODE. In the three works the authors consider different asymptotic regimes, both in the number of peers and in the number of chunks. For instance, the authors of [16] study what they call *coupon replication* system, that models a file sharing BitTorrent-like mechanism. Their model consider many users, each one aiming to complete a collection of coupons. At each time two users meet and obtain one missing coupon from the other by replication (if they do not have the same coupons). The model is motivated by the BitTorrent mechanism, where each chunk is a coupon. Results from [2] are used to prove the approximation by an asymptotic deterministic model for the number of coupons hold by each user, when the number of coupons goes to infinity. However, a closed form formula is obtained only for some particular cases.

B. Stochastic models

In this subsection we introduce our stochastic models that describe the number of leechers and seeds for the systems studied in [12], [13]. From these microscopic descriptions of the models in [12], [13] we construct a sequence of processes that converges to deterministic limits. We consider a two-dimensional continuous time Markov chain for the number of leechers and seeds, so we describe the whole system. However, our model is motivated by a detailed description for the behavior of each peer. In Section III we prove that these limits verify equations (1) and (2) respectively, which are the starting points in [12], [13].
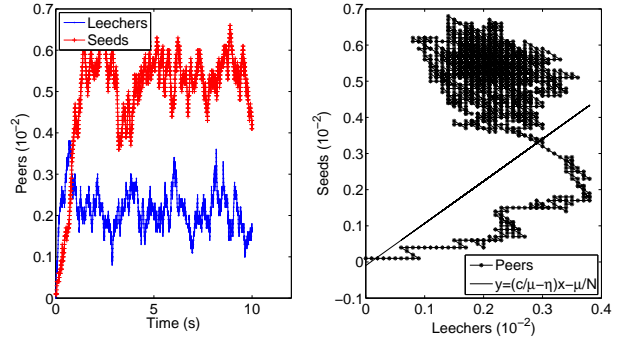

Figure 1. Evolution with time of the scaled number of leechers and seeds. (Model description in II-B and parameters in Table I). The line divides the plane in two zones with different behavior, considering capacity constraint in uploading or in downloading .

N	λ	μ	c	γ	θ	η
100	1	3	5	2	0.01	0.5

Table I
PARAMETER SET.

Let $\widetilde{X}^N(t)$ be the number of leechers and $\widetilde{Y}^N(t)$ the number of seeds at time t. We also assume that there is an additional fixed seed (the *total* number of seeds is thus $\widetilde{Y}^N(t) + 1$), so that the system never dies (results are the same for a finite fixed number of seeds). We specify the transitions at time t as follows:

- $\widetilde{X}^N(t)$: number of leechers at time t,
- $\widetilde{Y}^N(t) + 1$: number of seeds at time t,
- λN: arrival rate (Poisson) for peers (leechers),
- μ: uploading rate for each peer,
- c: downloading rate for each peer,
- $\eta \in [0, 1]$: efficiency factor,
- in the whole system a leecher becomes a seed with rate

$$\min\left(c\widetilde{X}^N(t),\ \eta\mu\widetilde{X}^N(t) + \mu\left(\widetilde{Y}^N(t) + 1\right)\right),$$

- the time in the system for a leecher before aborting is exponentially distributed with parameter θ,
- the time in the system for a seed before leaving is exponentially distributed with parameter γ.

For each N there is a line $y = (c/\mu - \eta)x - \mu$, where the behavior of the system changes because of the term $\min\left(c\widetilde{X}^N(t),\ \eta\mu\widetilde{X}^N(t) + \mu\left(\widetilde{Y}^N(t) + 1\right)\right)$, dividing the state space in two zones. In Figure 1 we show the evolution of the scaled number of leechers and seeds $(X^N(t), Y^N(t)) = \frac{1}{N}(\widetilde{X}^N(t), \widetilde{Y}^N(t))$, for the parameter set in Table I, with $(\widetilde{X}^N(0), \widetilde{Y}^N(0)) = (0, 1)$.

Let us compare our model with the BitTorrent Markov model previously proposed [11]. Differently to ours, they do not take into account the restriction in upload or download and the fact that peers may abandon the system before finishing their download. From a mathematical point of view, transition rates are continuous in our model, whereas in [11] there are discontinuities when $x = 0$. Due to this difference the same
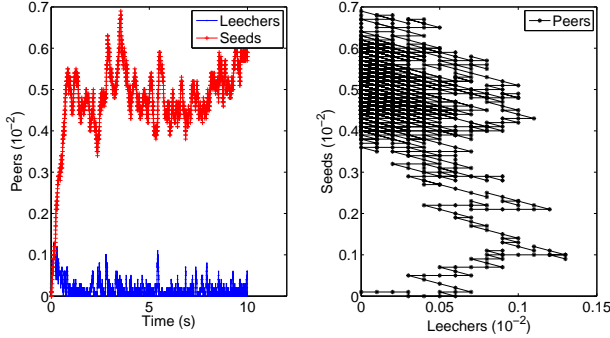
Figure 2. Evolution with time of the scaled number of leechers and seeds for the model in [11] (parameters in Table I).

techniques cannot be used to analyze them. To get some intuition on both models we compare them in Figure 1 and 2, where we show the evolution for the model in [11], when the arrival rate is λN, the scaled number of leechers and seeds. Note that there is a refracting barrier in $x = 0$ in Figure 2.

Now describe a stochastic microscopic model associated to the fluid model in [13]. Consider two classes of leechers. Let $\widetilde{X}_a^N(t)$ be the number of leechers of type a, $\widetilde{X}_b^N(t)$ the number of leechers of type b, and $\widetilde{Y}^N(t)$ the number of seeds at time t. The total number of seeds is $\widetilde{Y}^N(t) + 1$, so that the system never dies. We specify the transitions at time t as follows:

- $\lambda_a N$, $\lambda_b N$: arrival rate (Poisson) for peers (leechers) of type a and b respectively,
- a leecher of type a becomes a seed with rate

$$\min\left(c\widetilde{X}_a^N(t), \mu\widetilde{X}_a^N(t) + \mu(\widetilde{Y}^N(t) + 1)\frac{\widetilde{X}_a^N(t)}{\widetilde{X}^N(t)}\right),$$

and the same holds for a leecher of type b with rate

$$\min\left(c\widetilde{X}_b^N(t), \mu\widetilde{X}_b^N(t) + \mu(\widetilde{Y}^N(t) + 1)\frac{\widetilde{X}_b^N(t)}{\widetilde{X}^N(t)}\right),$$

- the time in the system for a leecher before aborting its download is exponentially distributed with parameter θ_a for leechers of type a and with parameter θ_b for leechers of type b, with $\theta_b > \theta_a$,
- the time in the system for a seed before leaving is exponentially distributed with parameter γ.

Regarding our model we have considered a Poisson arrival process and exponentially distributed times. There are references in the literature where this assumption is discussed and contrasted with real measurements [17], [19]. This point will be addressed in future work.

III. RESULTS

In this section we present results about deterministic approximations for the model in II-B. Convergence to an ODE, existence of stationary regime and convergence of the stationary regime are studied in III-A. A Central Limit Theorem is discussed in III-B.

A. Deterministic approximation

We justify the fluid approximation of our model stated in II-B, obtaining equation (1) as the limit when the arrival rate for peers goes to infinity. We also prove the existence of a stationary regime and the convergence in this regime to the ODE's fixed point. These issues are discussed in [12], [13], [16] without proofs and the whole system is directly analyzed from the study of ODE's fixed points.

Proposition 1. *Consider*

$$\left(X^N(t), Y^N(t)\right) = \frac{1}{N}\left(\widetilde{X}^N(t), \widetilde{Y}^N(t)\right)$$

and (x, y) *the solution to equation* (1) *with initial condition* $(x(0), y(0))$. *If*

$$\lim_{N\to\infty}\left(X^N(0), Y^N(0)\right) = (x(0), y(0))$$

then, for all $T > 0$,

$$\lim_{N\to\infty}\sup_{t\in[0,T]}\left\|\left(X^N(t), Y^N(t)\right) - (x(t), y(t))\right\| = 0 \ a.s.,$$

where a.s. means almost sure convergence.

Proof: The possible transitions in the N-th model, from state $(\widetilde{X}^N(t), \widetilde{Y}^N(t))$ are the following:

- a leecher arrives with rate $N\lambda$,
- a leecher becomes seed with rate

$$N\min\left(cX^N(t), \mu\left(\eta X^N(t) + Y^N(t) + 1\right)\right),$$

- a leecher aborts before downloading with rate $N\theta X^N(t)$,
- a seed leaves the system with rate $N\gamma Y^N(t)$.

$(\widetilde{X}^N(t), \widetilde{Y}^N(t))$ is a jump Markov process with transition rates of the form

$$N\beta_l\left[\left(X^N(t), Y^N(t)\right) + O\left(\frac{1}{N}\right)\right]$$

for $l \in \mathbb{Z}^2$ (l represents a possible transition). As β_l is bounded and Lipschitz on compact subsets, result follows directly from Kurtz's Theorem (Theorem 2.1, p. 456) in [2]. ■

The previous proposition is illustrated in Figure 3. In the left we show the simulation of one trajectory of the scaled Markov chain (number of leechers and seeds) for large N and the trajectory of the ODE. In the right we show for the same simulation the evolution on the plane of the Markov chain and the ODE. We can see from that picture that, for large time values, the number of leechers is around the ODE's fixed point. We analyze this in Theorem 1.

The proof of Kurtz's Theorem relies on a characterization on the process (X^N, Y^N) as a sum of independent Poisson processes (one for each direction of possible transitions) evaluated in a random time change. Under this characterization the theorem follows from Gronwall's inequality and the law of large numbers for the Poisson process.

The result of Proposition 1 is also valid for the stochastic model associated with the system in [13], as it verifies the same hypotheses and it is also valid for the priority scheme proposed in the same paper [13]. The Markov chain that
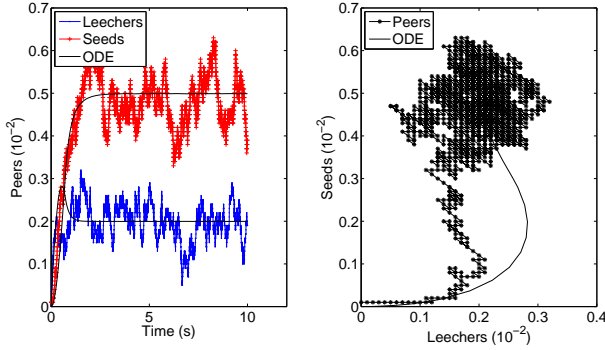
Figure 3. Theorem 1. (Model description in II-B and parameters in Table I.)

represents the priority scheme is of the same type as the previous ones, so the fluid approximation also holds.

We briefly analyze the model in [11], as it is intrinsically different. As before, let $(\widetilde{X}^N(t), \widetilde{Y}^N(t))$ be the number of leechers and seeds at time t and $(X^N(t), Y^N(t)) = \frac{1}{N}(\widetilde{X}^N(t), \widetilde{Y}^N(t))$. Consider the following transitions:

- a leecher arrives with rate $N\lambda$,
- a leecher becomes seed with rate $N\mu\left(X^N(t) + Y^N(t)\right)$,
- a seed leaves the system with rate $N\gamma Y^N(t)$.

The convergence stated in Proposition 1 relies on the fact that transition rates from state $(\widetilde{X}^N(t), \widetilde{Y}^N(t))$ are of the form $N\beta_l\left[\left(X^N(t), Y^N(t)\right) + O(1/N)\right]$, with β a Lipschitz function. This assumption does not hold for the model in [11], as transition rates are discontinuous in the boundary $x = 0$. This corresponds to a class of jump Markov processes studied in [3], called *flat boundary processes*. From [3] (Chap. 8), in this case there is an analogous of Kurtz's Theorem, and the ODE that approximates the scaled process $(X^N(t), Y^N(t))$ is the following.

If $x > 0$ or $\lambda - \mu(\eta x + y) \geq 0$,

$$\begin{cases} x' = \lambda - \mu(x + y), \\ y' = \mu(x + y) - \gamma y, \end{cases}$$

and if $x = 0$ and $\lambda - \mu(x + y) < 0$,

$$\begin{cases} x' = \pi_0\lambda + (1 - \pi_0)(\lambda - \mu(x + y)), \\ y' = -\pi_0\gamma y + (1 - \pi_0)(\mu(x + y) - \gamma y), \end{cases}$$

with

$$\pi_0 = \begin{cases} \dfrac{\mu(x + y) - \lambda}{\mu(x + y)} & \text{if } \lambda - \mu(x + y) < 0, \\ 0 & \text{if } \lambda - \mu(x + y) \geq 0. \end{cases}$$

The above equations show that it is possible to obtain fluid limits for this kind of models, despite discontinuities in transition rates.

Now we turn our attention to the ergodicity of $(\widetilde{X}^N(t), \widetilde{Y}^N(t))$. It seems not simple to find the stationary distribution explicitly. Classical sufficient conditions as reversibility are not verified, so we cannot assume local balance equations. We prove ergodicity by using a Lyapunov function. The ergodicity result is also stated in [13], as there is a Markov

model that is compared with the fluid one by simulations, using queuing arguments that allow to reduce the analysis of the existence of a stationary regime to the study of a Jackson network. However, the proof presented here is simpler and more detailed.

Proposition 2. *The process* $(\widetilde{X}^N(t), \widetilde{Y}^N(t))$ *is ergodic for each* N.

Proof: The proof is based on [20] (Proposition 8.14, p. 225). Function $f(x, y) = x + y$ is a Lyapunov function for $(\widetilde{X}^N(t), \widetilde{Y}^N(t))$. We must verify that there exists K and h such that the following conditions hold:

1) for $f(x, y) > K$, $Q(f)(x, y) \leq -h$, with $Q(f)(x, y) = \sum_{l \in \mathbb{Z}^2, l \neq 0} q\left((x, y), (x, y) + l\right)\left[f((x, y) + l) - f(x, y)\right]$ ($q\left((x, y), (x, y) + l\right)$ is the transition rate from (x, y) to $(x, y) + l$);
2) the random variables

$$\sup\{f(\widetilde{X}^N(s), \widetilde{Y}^N(s)) : s \leq 1\},$$

$$\int_0^1 |Q(f)(\widetilde{X}^N(s), \widetilde{Y}^N(s))|\, ds$$

are integrable;
3) $F = \{(x, y) : f(x, y) \leq K\}$ is finite.

These assumptions imply that the process is ergodic. Let us verify each one of them:

1) $Q(f)(x, y) = \lambda N - \theta x - \gamma y \leq -h$ for $x + y > K$; it suffices then to take $K \geq (\lambda N + h)/\min(\theta, \gamma)$.
2) The Poisson process $Z(s)$ with rate $N\lambda$ is an upper bound of $f(\widetilde{X}^N(s), \widetilde{Y}^N(s)) = \widetilde{X}^N(s) + \widetilde{Y}^N(s)$, and it is integrable on each bounded interval. Analogously $\lambda N + \max(\theta, \gamma)Z(s)$ is an upper bound of $|Q(f)(\widetilde{X}^N(s), \widetilde{Y}^N(s))|$; the integral of the former is thus bounded by $\lambda N + \max(\theta, \gamma)\int_0^1 Z(s)\, ds$ and it is then integrable.
3) Immediate.

From the previous assumptions $(\widetilde{X}^N(t), \widetilde{Y}^N(t))$ is ergodic for each N. ∎

The same result holds for the stochastic model considering two classes of leechers described above. In that case a Lyapunov function is $f(x_a, x_b, y) = x_a + x_b + y$. The assumptions are verified as above. It follows from noticing again that the only possible transition away from a region $\{(x_a, x_b, y) : x_a + x_b + y \leq K\}$ is when a new peer arrives. As arrivals follow Poisson processes, the hypotheses about finite expectation hold.

In what follows we prove the convergence of the stationary regime to the ODE's fixed point. As the process $(X^N(t), Y^N(t))$ converges in bounded intervals and has a stationary distribution, one can expect that the stationary distribution converges to the ODE's fixed point. This result is used for the analysis in [12], and in different contexts in other works (see for example [16]), sometimes without a detailed proof. In [12] it is proven that the ODE has an unique fixed
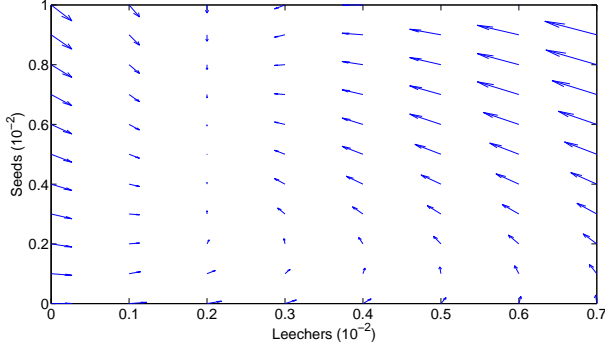
Figure 4. Vector field for equation (1) (parameter set in Table I.).

point

$$(x^*, y^*) = \left(\frac{\lambda}{\beta \left(1 + \frac{\theta}{\beta}\right)}, \frac{\lambda}{\gamma \left(1 + \frac{\theta}{\beta}\right)} \right),$$

with

$$\frac{1}{\beta} = \max\left\{ \frac{1}{c}, \frac{1}{\mu} - \frac{1}{\gamma} \right\}.$$

and that the system is locally stable. The work from Qiu and Sang [21] is devoted to the analysis of equation (1), and they prove that the unique fixed point is a global attractor. We show in Figure 4 the vector field associated with equation (1).

Theorem 1. *Let $\left(X^N(\infty), Y^N(\infty) \right)$ be the scaled number of leechers and seeds in stationary regime. Let (x^*, y^*) be the fixed point in* (1)*. Then*

$$\lim_{N \to \infty} \left(X^N(\infty), Y^N(\infty) \right) = (x^*, y^*)$$

in probability.

Proof: Let $\mu^N(t)$ be the distribution of $\left(X^N(t), Y^N(t) \right)$ and let $\pi^N(\infty)$ be the stationary distribution of the process (we know from Proposition 2 that there exists a unique stationary distribution for each N). We will use for our proof Theorem 6.89, p. 165 in [3]. This theorem assures that under a set of hypotheses that will be verified, if (x^*, y^*) is a global attractor then $\lim_{N \to \infty} \int_{B_\varepsilon(q)} d\pi^N(\infty) = 1$, with $(x^*, y^*) = q$ and $B_\varepsilon(q) = \left\{ y \in \mathbb{R}^2 : \|y - q\| < \varepsilon \right\}$, which implies that

$$\lim_{N \to \infty} \left(X^N(\infty), Y^N(\infty) \right) = (x^*, y^*)$$

in probability. To apply the result we must verify that:
1) the jumps of the Markov process take integer values in each direction,
2) the rates β_l are uniformly Lipschitz continuous in a neighborhood of (x^*, y^*),
3) the process is positive recurrent,
4) if $\tau_\varepsilon(N) = \inf \left\{ t : \left\| \left(X^N(t), Y^N(t) \right) - (x^*, y^*) \right\| < \varepsilon \right\}$, then for each K, ε and for all N, there exists a constant $C_{\varepsilon,M}$ (that depends on ε and M) such that

$$\sup_{\|p-q\| \leq M} E_p\left[\tau_\varepsilon(N) \right] \leq C_{\varepsilon,M} < \infty,$$

with $p = (x(0), y(0))$ and E_p the expected value starting from p.

The first three assumptions are immediately verified (the third one arises from Proposition 2). So, we focus on fourth assumption. Also from [3] (Lemma 6.32 p. 143) the distribution of $\tau_\varepsilon(N)$ has geometric tails for large N, that is, there is a $T(\varepsilon) < \infty$ and a constant $C_0(\varepsilon)$ such that

$$P_p(\tau_\varepsilon(N) > kT) \leq e^{-NC_0(\varepsilon)k}.$$

This implies the bound for $E_p\left[\tau_\varepsilon(N) \right]$ and thus completes the proof. ∎

The convergence for the stationary distribution to the ODE's fixed points is a widely discussed topic. The authors of [4] prove this convergence in the case of the occupation measure of a system with N individuals and a finite state space. Our problem differs from that situation because of the compactness of the state space. However, our proof and the proof in [4] rely in large deviations results. The proof in [4] is based on [22], where a very general result (considering the case with multiple invariant distributions and a much more complex asymptotic behavior for the ODE) is proven using large deviations arguments together with dynamical systems ones. In [4] it is also discussed why the existence of a unique fixed point does not guarantee the convergence of a sequence of invariant distributions. It shows examples where there is only one fixed point but the support of accumulation points of invariant distributions lies on set that is a limit cycle for the ODE. In order to avoid the problem of proving asymptotic stability, [10] presents a very general result of convergence for the stationary distribution when there is a unique fixed point in case of reversible processes, a strong assumption that is not valid in our model. The convergence for the stationary distribution of the occupation measure is also discussed in [23], in a more general framework (denumerable state spaces). The proof there is strongly related with the mean field decoupling assumption (the asymptotic independence and the convergence of the stationary distributions are proved together).

For the model in [13] we do not have a proof of global stability of ODE's fixed point, so we cannot yet extend the previous theorem for that case. It is only observed in simulations in [13] that the stationary regime converges to the fixed point of the associated ODEs, both in the priority and non priority schemes.

B. Gaussian approximation

Here we derive a Gaussian approximation for the distribution of the difference between the stochastic and the deterministic processes. This approximation describes in a precise way the system behavior for large values of N simultaneously for all t providing confidence intervals for the number of leechers and seeds.

Theorem 2. *Consider $\left(X^N(t), Y^N(t) \right)$ and let (x, y) be the solution to equation* (1) *with initial condition $(x(0), y(0))$. If*

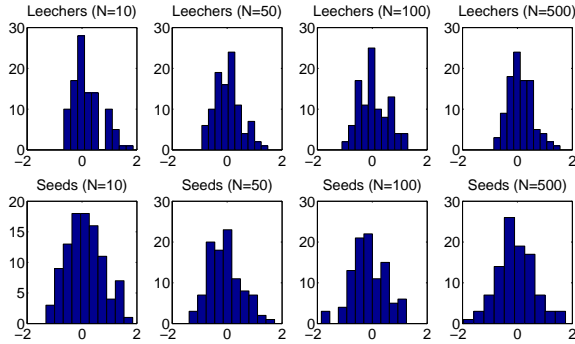$$\lim_{N \to \infty} \sqrt{N} \left[\left(X^N(0), Y^N(0) \right) - (x(0), y(0)) \right] = V(0)$$

Figure 5. Histograms of 100 independent samples of $\sqrt{N}\left(X^N(1) - x(1)\right)$ and $\sqrt{N}\left(Y^N(1) - y(1)\right)$ for different values of N.
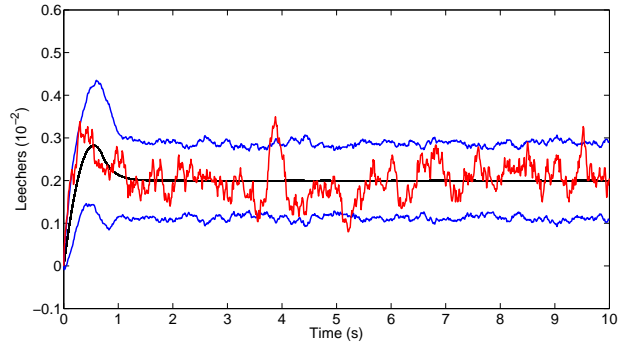

Figure 6. ODE trajectory, confidence interval for each t and one trajectory for the scaled number of leechers $X^N(t)$. The variance is computed from 100 independent replications of the experiment (parameters in Table I).

in probability, with $V(0)$ deterministic, then,

$$\sqrt{N}\left[\left(X^N(t), Y^N(t)\right) - (x(t), y(t))\right] \Rightarrow_N V(t),$$

where $\Rightarrow$ means convergence in distribution. $V(t)$ is a Gaussian process with covariance matrix

$$Cov\left(V(t), V(r)\right)$$
$$= e^{M(t)t + M(r)^T r} \int_0^{t \wedge r} e^{-\left(M(t) + M(r)^T\right)s} G(s)\, ds,$$

$$G(s) = \lambda + \theta x(s) + \gamma y(s) + 2\min\left(cx(s),\, \mu(\eta x(s) + y(s))\right),$$

$$M(t) = \begin{pmatrix} -(c+\theta) & 0 \\ c & -\gamma \end{pmatrix} \text{ if } cx(t) < \mu(\eta x(t) + y(t)),$$

$$M(t) = \begin{pmatrix} -(\mu\eta + \theta) & -\mu \\ \mu\eta & \mu - \gamma \end{pmatrix} \text{ if } cx(t) > \mu(\eta x(t) + y(t)),$$

$M(t)^T$ *denote the transposed of $M(t)$.*

Proof: Result follows as a consequence of Kurtz's Theorem (see Theorem 2.3, p.458, in [2]). We use the explicit form of the covariance matrix provided there. The proof of that theorem relies on a representation of $V^N(t)$ and $V(t)$ by an integral involving the differential $dF(x(t), y(t))$, so the original theorem assumes that the transition rates $\beta_l(x, y)$ are C^1 functions. This assumption is not valid in our case, but there is only one t where $\beta_l(x(t), y(t))$ is not differentiable (that is when $cx(t) = \mu(\eta x(t) + y(t))$). As this happens at only one point, it does not affect the integral representation. The justification that there is only one t for which $cx(t) = \mu(\eta x(t) + y(t))$ follows from the fact that the fixed point is a global attractor, so the trajectories $(x(t), y(t))$ hit $\{(x, y) : cx = \mu(\eta x + y)\}$ only a finite number of times (this is sufficient for the validity of the integral representation of $V^N(t)$ as in Theorem 2.3, p. 458, [2]). In Figure 4 it can be seen that there is at most one hitting point in our case. ∎

In Figure 5 we show, for different values of N, histograms of 100 independent samples of $\sqrt{N}\left(X^N(t) - x(t)\right)$ and $\sqrt{N}\left(Y^N(t) - y(t)\right)$ for a fixed t and in Figure 6 we show the 95% confidence interval for the scaled number of leechers.

Paper [12] describes without a detailed proof that the variability around the fluid limit (the solution to equation (1)). For a large arrival rate λ, the number of leechers and seeds are approximately $x(t) + \sqrt{\lambda}\hat{x}(t)$ and $y(t) + \sqrt{\lambda}\hat{y}(t)$, with $\hat{x}(t)$ and $\hat{y}(t)$ gaussian processes (Ornstein-Uhlenbeck). In our framework we have that the arrival rate is $N\lambda$ and the number of leechers and seeds are characterized by $\widetilde{X}^N(t) \approx Nx(t) + \sqrt{N}V_1(t)$ and $\widetilde{Y}^N(t) \approx Ny(t) + \sqrt{N}V_2(t)$, with $V = (V_1, V_2)$ the gaussian process described in Theorem 2. We observe that the limit process $V(t)$ verifies a stochastic differential equation (see equation (2.18), p. 458 in [2]). The gaussian process stated in [12] can be obtained from this stochastic differential equation replacing $x(t)$ and $y(t)$ by its respective limits x^* and y^*.

In [2] we can find two approaches in order to characterize the variability of the stochastic process (X^N, Y^N) around the deterministic process (x, y). The first one is the approximation using the Central Limit Theorem that we use here, and the second one is the *diffusion approximation*. These approximations are equivalent in bounded time intervals for large N [2]. In both cases certain regularity of the transition rates is assumed. In Theorem 2 we have weakened this regularity assumption in the context of the Central Limit Theorem (see Theorem 2.3, p. 458, in [2]). Concerning the diffusion approximation for non-regular transition rates we refer to [24].

IV. CONCLUSIONS AND FUTURE WORK

In this paper we provide new elements to the understanding of the well known fluid models of BitTorrent systems in the spirit of [12] and extensions such as [13]. We consider sequences of stochastic models (Markov chains) representing the population of different types of peers in the network. Leechers arrive following a Poisson process, with its rate increasing with an integer parameter N. We scale the number of leechers and seeds by that same factor N. We then provide rigorous justifications for the passage from the sequences of stochastic models to deterministic limits when N goes to infinity. This convergence basically follows from Kurtz's Theorem. We prove the existence of a stationary regime for the processes in the sequence by constructing a Lyapunov function for each

one of them and the convergence of the stationary distribution to the ODE's fixed point. This result involves some arguments from large deviations' theory. We also prove, as a consequence of Kurtz's Theorem, that asymptotically the scaled number of leechers and seeds follows a gaussian process. These are topics to be further analyzed in future work, together with some stability issues. We will also analyze the assumptions about Poisson arrivals and exponentially distributed times, in order to consider more general models.

REFERENCES

[1] R. W. R. Darling and J. R. Norris, "Differential equation approximations for Markov chains," *Probab. Surv.*, vol. 5, pp. 37–79, 2008.

[2] S. N. Ethier and T. G. Kurtz, *Markov processes: Characterization and convergence*, ser. Wiley Series in Probability and Statistics. New York: John Wiley & Sons Inc., 1986.

[3] A. Shwartz and A. Weiss, *Large deviations for performance analysis*, ser. Stochastic Modeling Series. London: Chapman & Hall, 1995, queues, communications, and computing, With an appendix by Robert J. Vanderbei.

[4] M. Benaïm and J.-Y. Le Boudec, "A class of mean field interaction models for computer and communication systems," *Perform. Eval.*, vol. 65, no. 11-12, pp. 823–838, 2008.

[5] C. Bordenave, D. McDonald, and A. Proutiere, "A particle system in interaction with a rapidly varying environment: Mean field limits and applications," *Networks and Heterogeneous Media*, vol. 5, no. 1, pp. 31–62, 2010.

[6] A. Chaintreau, J.-Y. Le Boudec, and N. Ristanovic, "The age of gossip: Spatial mean-field regime," in *ACM Sigmetrics*, , Ed., ACM Sigmetrics, 2009, best paper award.

[7] J.-Y. Le Boudec, D. McDonald, and J. Mundinger, "A generic mean field convergence result for systems of interacting objects," in *QEST '07, Proceedings of the Fourth International Conference on Quantitative Evaluation of Systems*. Washington, DC, USA: IEEE Computer Society, 2007, pp. 3–18.

[8] N. Gast, B. Gaujal, and J.-Y. Le Boudec, "Mean field for Markov Decision Processes: from Discrete to Continuous Optimization," Tech. Rep., 2010.

[9] H. Tembine, J.-Y. Le Boudec, R. El-Azouzi, and E. Altman, "Mean Field Asymptotic of Markov Decision Evolutionary Games and Teams," in *Gamenets 2009*, 2009, invited Paper.

[10] J.-Y. Le Boudec, "The Stationary Behaviour of Fluid Limits of Reversible Processes is Concentrated on Stationary Points," Tech. Rep., 2010.

[11] X. Yang and G. de Veciana, "Service capacity of peer to peer networks," in *IEEE INFOCOM*, 2004, pp. 1–11. [Online]. Available: citeseer.ist.psu.edu/yang04service.html

[12] D. Qiu and R. Srikant, "Modeling and performance analysis of bittorrent-like peer-to-peer networks," in *SIGCOMM '04: Proceedings of the 2004 conference on Applications, technologies, architectures, and protocols for computer communications*. New York, NY, USA: ACM, 2004, pp. 367–378.

[13] M. E. Rivero-Angeles and G. Rubino, "Priority-based scheme for file distribution in peer-to-peer networks," in *IEEE International Conference on Communications (ICC), 2010*, 2010, pp. 1–6.

[14] G. Kesidis, T. Konstantopoulos, and P. Sousi, "Modeling file-sharing with bittorrent-like incentives," in *IEEE International Conference on Acoustics, Speech and Signal Processing, ICASSP 2007*, vol. 4, apr. 2007, pp. IV–1333 –IV–1336.

[15] ——, "A stochastic epidemiological model and a deterministic limit for bittorrent-like peer-to-peer file-sharing networks," in *Network Control and Optimization*, ser. Lecture Notes in Computer Science, E. Altman and A. Chaintreau, Eds. Springer Berlin / Heidelberg, vol. 5425, pp. 26–36.

[16] L. Massoulié and M. Vojnović, "Coupon replication systems," *SIGMETRICS Perform. Eval. Rev.*, vol. 33, no. 1, pp. 2–13, 2005.

[17] L. Guo, S. Chen, Z. Xiao, E. Tan, X. Ding, and X. Zhang, "Measurements, analysis, and modeling of bittorrent-like systems," in *IMC '05: Proceedings of the 5th ACM SIGCOMM conference on Internet Measurement*. Berkeley, CA, USA: USENIX Association, 2005, pp. 4–4.

[18] M. Izal, G. Urvoy-Keller, E. W. Biersack, P. Felber, A. A. Hamra, and L. Garcés-Erice, "Dissecting bittorrent: Five months in a torrent's lifetime," in *PAM*, 2004, pp. 1–11.

[19] J. A. Pouwelse, P. Garbacki, D. H. J. Epema, and H. J. Sips, "The bittorrent p2p file-sharing system: Measurements and analysis," in *IPTPS*, 2005, pp. 205–216.

[20] P. Robert, *Stochastic Networks and Queues*, ser. Stochastic Modelling and Applied Probability Series. New York: Springer-Verlag, 2003.

[21] D. Qiu and W. Sang, "Global stability of peer-to-peer file sharing systems," *Comput. Commun.*, vol. 31, no. 2, pp. 212–219, 2008.

[22] M. Benaïm, "Dynamics of stochastic approximation algorithms," in *Séminaire de Probabilités, XXXIII*, ser. Lecture Notes in Math. Berlin: Springer, 1999, vol. 1709, pp. 1–68.

[23] C. Bordenave, D. McDonald, and A. Proutiere, "Performance of random medium access control, an asymptotic approach," in *SIGMETRICS '08: Proceedings of the 2008 ACM SIGMETRICS international conference on Measurement and modeling of computer systems*. New York, NY, USA: ACM, 2008, pp. 1–12.

[24] R. Hayden and J. T. Bradley, "A functional central limit theorem for PEPA," in *8th Workshop on Process Algebra and Stochastically Timed Activities*, August 2009, pp. 13–23.