

Predicting the Impact of Measures Against P2P Networks on the Transient Behaviors

Eitan Altman, Philippe Nain, Adam Shwartz, Yuedong Xu

► **To cite this version:**

Eitan Altman, Philippe Nain, Adam Shwartz, Yuedong Xu. Predicting the Impact of Measures Against P2P Networks on the Transient Behaviors. IEEE Infocom, Apr 2011, Shangai, China. hal-00648867

HAL Id: hal-00648867

<https://hal.inria.fr/hal-00648867>

Submitted on 7 Dec 2011

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Predicting the Impact of Measures Against P2P Networks on the Transient Behaviors

Eitan Altman, Philippe Nain
INRIA Sophia Antipolis
2004 Route des Lucioles
{eitan.altman,philippe.nain}@inria.fr

Adam Shwartz
Technion
Haifa 32000, Israel
adam@ee.technion.ac.il

Yuedong Xu
INRIA Sophia Antipolis
2004 Route des Lucioles
yuedong.xu@inria.fr

Abstract—The paper has two objectives. The first is to study rigorously the transient behavior of some peer-to-peer (P2P) networks whenever information is replicated and disseminated according to epidemic-like dynamics. The second is to use the insight gained from the previous analysis in order to predict how efficient are measures taken against P2P networks. We first introduce a stochastic model which extends a classical epidemic model, and characterize the P2P swarm behavior in presence of free riding peers. We then study a second model in which a peer initiates a contact with another peer chosen randomly. In both cases the network is shown to exhibit phase transitions: a small change in the parameters causes a large change in the behavior of the network. We show, in particular, how phase transitions affect measures of content providers against P2P networks that distribute non-authorized music or books, and what is the efficiency of counter-measures.

I. INTRODUCTION

Along with the worldwide penetration of the Internet, a huge demand has appeared to copyrighted music and movies that have been accessible for free over the Internet. While benefiting a very large internaut community as well as potentially providing higher benefits for Internet access providers, it seems unclear whether the creators and the copyright owners have gained anything from this unregulated access. Two opposing approaches appeared, both proposing to protect the copyright owners. The first consists of fighting against non-authorized access whereas the second one, aims at finding cooperative solutions that would benefit both the Internauts as well as all other economic actors. An example of a cooperative solution is a flat taxation that would allow Internauts to pursue downloading freely music and films, and that would distribute the tax money between the copyright owners. This cooperative approach has several difficulties in its implementation; a major one is how to distribute the tax income fairly. A major drawback of the confrontation policy is the huge monitoring effort that it requires and that seems not to provide credible evidence for unauthorized downloads [1]. In order to assess the efficiency of non-cooperative measures against unauthorized downloads, the authors of [2] have analyzed the impact of the effort, of the authorities or of content provider companies, invested in (i) reducing file uploading in P2P networks and in (ii) reducing the demand for files, on the availability of files and, more generally, on the operation of the P2P networks. The stationary analysis there is based on a $M/G/\infty$ queuing model.

In this paper we are interested in predicting the impact

of measures as described in the previous paragraph, on the *transient* behavior of torrents. By how much should the request or departure rate in a P2P network be reduced in order to have a significant change in file availability? To achieve that, we consider abstract models of a torrent in simplified P2P networks, where a large number of peers are interested in a file which is initially available at a small fraction of the population.

Our models are formulated as epidemic type processes of file dissemination. We consider both *cooperative* peers, which are those that make a file available to other peers as soon as they obtain the file, and *free riders*, who leave the system immediately after obtaining the file. To understand the impact of measures against the cooperative sharing behavior, we parameterize the degree of free-riding in the system as well as the degree of cooperation.

The P2P dynamics is modeled by a Markov chain (Section II) which is approximated in two specific regimes: the first (Section III) is the early stage when a large fraction of the population does not yet have the file. The system is then well approximated by a branching process. In the case that there is a positive probability of not getting extinct in the first regime, the system is shown in Section IV to move with some non-zero probability to a second regime in which, for the case of a sufficiently large population size, its dynamics is close to the solution of a differential equation. A similar fluid limit is studied in VI for the case of limitation on uplink or downlink speed. We briefly state our contributions:

1. Modeling and approximating the transient behavior

Our first important contribution is to show in what sense each of the above two models approximates the original Markov chain, and how to use both in order to get the whole transient behavior of the P2P network. This is in contrast with all other models of P2P networks that we know of, which either use only a branching process approach [14] or which use only an epidemic mean-field approximation [11]. The latter approach (of using only the mean-field limit) is shown to provide a tight approximation when the initial number of peers with the file scales linearly with the total size N of the population of peers. With a fixed initial number of nodes that does not scale with N , there is a positive probability of early extinction (see Section VII for detail) for any set of system parameters, and this probability cannot be predicted by the mean-field limit alone.

2. Analysis and identifying phase transitions

We first

study a P2P model that corresponds to the epidemic-like file dissemination (Sections II-IV) We then study a second model (Section VI) in which, at random times, each peer contacts another peer randomly chosen within the set of existing peers. In both cases, we show the existence of phase transitions: a small change in the parameters causes a large change in the network behavior.

A phase transition occurs both in the branching model for the extinction time and in the epidemic model for the file availability. In the branching process, the existence of two phases was not known to Galton and Watson (considered as the founders of branching processes) and was only discovered and proved later in [3]. In the epidemiology community, the phase transition was already known in [4] for a model equivalent to our first model without the free riders. For the the second model [5], we show the existence of two phase transitions, one for the file availability and the other one for the maximum torrent size.

3. Application. In Section V, we present a counteraction against unauthorized file sharing in the presence of illegal publishers. We evaluate the impact of measures against Internet piracy on the performance of P2P systems in Section VII (see Figure 11).

The accuracy of the various approximations is investigated in Section VII, related studies are discussed in Section VIII, and concluding remarks are given in Section IX.

II. MODEL

A. Assumptions

Assume there is a population N of peers interested in a single file. Let $Y(t)$ be the number of peers that possess the file at time t . A peer acquires the file when it encounters another peer that has the file. We will consider two types of peers: *cooperative* and *non-cooperative* peers. Once a cooperative peer has acquired the file, it stays in the network for a random time distributed according to an exponential rv with parameter $1/\mu \geq 0$ and then leaves the network. During the lingering time of a cooperative peer with the file, it participates in the file dissemination. A non-cooperative peer, also called a *free-rider*, leaves the network at once when it receives the file. Note that “free riders” in our context is an abstract description of noncooperative behaviors, which is different from that in the current BitTorrent system.

Let $X_c(t)$ and $X_f(t)$ denote the number of cooperative peers without the file and the number of free-riders (necessarily without the file) at time t , respectively. Define the process $\mathbf{Y} := \{(Y(t), X_c(t), X_f(t)), t \geq 0\}$. Let $(Y(0), X_c(0), X_f(0))$ denote the initial state of \mathbf{Y} that has $Y(0) + X_c(0) + X_f(0) = N$. Let the ratio of various types of peers be $(y_0, x_{c,0}, x_{f,0}) := (\frac{Y(0)}{N}, \frac{X_c(0)}{N}, \frac{X_f(0)}{N})$. For simplicity, we introduce new variables $N_c = X_c(0)$ and $N_f = X_f(0)$.

We consider an abstract P2P network in which the file acquisition is via random contact between pair-wise peers. When two such peers meet, the cooperative peer transmits the file to the other peer. It is assumed that it takes an exponential time with rate $\lambda > 0$ for a peer without a file to encounter a cooperative peer with the file. The transmission of the file

is always supposed to be successful. This model describes a general P2P swarm without a tracker, and even the spreading of a file in current Internet. It is inspired by the contact process in [5] and [16]. One of the main difference lies in that a peer contacts all other connected peers in the system, instead of only one random peer periodically. We assume that the file transmission time is negligible compared to the time it takes for two peers to meet and therefore this time is taken to be zero.

All the random variables (rvs) introduced so far are assumed to be mutually independent. As a consequence, if $Y(t) = k$ then any peer without the file will meet a cooperative peer with the file after a time that is distributed according to the minimum of k independent and exponential rvs with rate λ , that is after a time distributed according to an exponential rv with rate λk .

Measures of the authorities or of content provider companies against file sharing systems may have an impact on the decrease in the population N interested in the file and an increase in the fraction of free riders among the population interested in the file. It can however have an impact also on the behavior of cooperative peers that would leave the system sooner (i.e. μ is expected to increase). Our model combines an epidemic type propagation of the file together with a description of the free riding behavior. Define $\rho := \lambda N_c / \mu$.

We first consider (Section II-B) the case where all peers are fully cooperative in the sense that $\mu = 0$ and $X_f(0) = 0$ (no free riders). $\mu = 0$ implies that cooperative peers do not leave the network after receiving the file. We then move to the general case where $\mu > 0$ and $X_f(0) \geq 0$ (Section II-C).

B. Fully cooperative network

When all peers are fully cooperative (i.e. $\mu = 0$ and $X_f(0) = 0$) the population of peers remains constant and equal to N , that is, $Y(t) + X_c(t) = N$ at any time t . The network dynamics can be represented by the process $\{Y(t), t \geq 0\}$.

This is a finite-state continuous-time Markov process with non-zero transitions given by

$$Y(t) \rightarrow Y(t) + 1 \quad \text{with rate} \quad \lambda Y(t)(N - Y(t)). \quad (1)$$

In other words the process $\{Y(t), t \geq 0\}$ is a pure birth Markov process on the state-space $\{y_0, \dots, N\}$, where state N is an absorbing state which is reached when all peers have the file.

Define $m(t) := E[Y(t)]$, the expected number of peers with the file at time t . Standard algebra shows that

$$\frac{dm(t)}{dt} = \lambda E[Y(t)(N - Y(t))], \quad t > 0. \quad (2)$$

Unfortunately, the right-hand side of (2) does not express as a function of $m(t)$, thereby ruling out the possibility of finding $m(t)$ in closed-form as the solution of an ODE.

Assume that λ is written as $\lambda = \beta/N$ and that $\lim_{N \rightarrow \infty} N^{-1}Y(0) = y_0 \in (0, 1]$. Then, for large N , $m(t)$ is well-approximated by $Ny(t)$ where $y(t)$ is obtained as the unique solution of the ODE [6, Thm 3.1]

$$\frac{dy(t)}{dt} = f(y(t)), \quad t > 0, \quad (3)$$

where $f(u) := \beta u(1-u)$ and $y(0) = y_0 \in (0, 1]$ (conditions (3.2)-(3.4) in [6, Thm 3.1] are clearly satisfied). It is found that

$$y(t) = \frac{y_0}{(y_0 + (1-y_0)\exp(-\beta t))}, \quad t \geq 0. \quad (4)$$

This is a well-known instance (see e.g. [12]) of what is known as mean-field approximation, a theory that focuses on the solution of ODEs obtained as limits of jump Markov processes [6]. The ODE (3) has been extensively used in epidemiology studies, where $y(t)$ represents the fraction of infected patients at time t when the population is of size N .

Proposition 1 below, whose proof can be found in [22], states that the mean-field approximation is an upper bound for $E[Y(t)]$.

Proposition 1: $E[Y(t)] \leq N y_0 / (y_0 + (1-y_0)e^{-\beta t}) \forall t \geq 0$.

C. General network

We consider the general network defined in Section II-A.

Define the vector $X(t) = \begin{pmatrix} X_c(t) \\ X_f(t) \end{pmatrix}$, where we recall that $X_c(t)$ is the number of cooperative nodes in the system who do not have the file at time t and $X_f(t)$ is the number of free-riders in the system at time t (by definition, none of these have the file at time t). Let $e_c = (1, 0)$ and $e_f = (0, 1)$. Under the statistical assumptions made in Section II-A it is seen that the process $\mathbf{Y} = \{(Y(t), X(t)), t \geq 0\}$ is a finite-state Markov process whose non-zero transitions are given by

$$\begin{pmatrix} Y(t) \\ X(t) \end{pmatrix} \rightarrow \begin{pmatrix} Y(t) + 1 \\ X(t) - e_c \end{pmatrix} \quad \text{with rate } \lambda Y(t) X_c(t), \quad (5)$$

$$\begin{pmatrix} Y(t) \\ X(t) \end{pmatrix} \rightarrow \begin{pmatrix} Y(t) - 1 \\ X(t) \end{pmatrix} \quad \text{with rate } \mu Y(t), \quad (6)$$

$$\begin{pmatrix} Y(t) \\ X(t) \end{pmatrix} \rightarrow \begin{pmatrix} Y(t) \\ X(t) - e_f \end{pmatrix} \quad \text{with rate } \lambda Y(t) X_f(t). \quad (7)$$

Throughout this paper we will assume that $\lambda > 0$ and $\mu > 0$.

The process \mathbf{Y} takes its values in the set $\mathcal{E} := \{(i, j, k), 0 \leq i \leq y_0 + N_c, 0 \leq j \leq N_c, 0 \leq j + k \leq N - y_0\}$. Furthermore, all states in \mathcal{E} of the form $(0, j, k)$ are absorbing states since there are no more transitions when the file has disappeared.

An explicit characterization of the transient behavior of the absorbing Markov process \mathbf{Y} is a difficult task due both to the presence of non-linear and non-homogeneous transition rates in the state variables and to the dimension of \mathbf{Y} . In this paper we will instead develop two approximations of the Markov process \mathbf{Y} . The first one, in Section III, will consist in replacing $X_c(t)$ by $N_c = X_c(0)$ in the transition rate (5), which will introduce a birth and death Markov branching process. As expected, this (so-called) branching approximation will lose its accuracy as the ratio $X_c(t)/N_c$ decreases.

The second approximation, in Section IV, will use an asymptotic argument as $N \rightarrow \infty$ based on a mean-field approximation of \mathbf{Y} . This approximation is justified if the initial state of \mathbf{Y} is of the order of N . Both the branching and the mean-field approximations approaches will allow us to approximate key characteristics of \mathbf{Y} such as the probability of disappearance of the file, the time before all files disappear, the maximum number of cooperative peers in the network and the fraction of peers that eventually receive the file.

III. BRANCHING APPROXIMATION

Let $\mathbf{Y}_b := \{Y_b(t), t \geq 0\}$ be a Markov process on $\mathbb{N} := \{0, 1, \dots\}$ (the subscript b refers to ‘‘branching’’) with non-zero transition rates given by

$$Y_b(t) \rightarrow Y_b(t) + 1 \quad \text{with rate } \lambda Y_b(t) N_c \quad (8)$$

$$Y_b(t) \rightarrow Y_b(t) - 1 \quad \text{with rate } \mu Y_b(t) \quad (9)$$

where we recall that N_c is the number of cooperative peers without the file at time $t = 0$.

Since $X_c(t)$, the number of cooperative peers without the file at time t , is non-increasing in t , a quick comparison between (5)-(7) and (8)-(9) indicates that the process \mathbf{Y}_b should dominate the process \mathbf{Y} . This bounding result is formalized and proved in the proposition below.

A word on the notation: a real-valued rv Z_1 is stochastically smaller than another real-valued rv Z_2 , denoted as $Z_1 \leq_{st} Z_2$, if $P(Z_1 > x) \leq P(Z_2 > x)$ for all x .

Proposition 2: If $Y(0) \leq Y_b(0)$ then $Y(t) \leq_{st} Y_b(t)$ for any $t > 0$.

The Markov process \mathbf{Y}_b is an absorbing continuous-time birth and death process on \mathbb{N} with absorbing state 0. Because its transition rates are linear functions of the system state, this is also a continuous-time Markov *branching process* [9], namely, a process in which at any time t each member of $Y_b(t)$ evolves independently of each other. The next section specializes known results of the theory of branching processes to the process \mathbf{Y}_b .

A. Extinction probability and extinction time

As previously observed the process \mathbf{Y}_b is a birth and death branching process [9, Chapter V]. Each object (peer) of this process has a probability of change in the interval $(t, t+h)$ given by $bh + o(h)$ with $u = \lambda N_c + \mu$; with probability $p_0 = \mu/u$ an object dies (a peer leaves) and with probability $p_2 = \lambda N_c/u$ an object is replaced by two objects (a peer receives the file).

Given $Y_b(0) = k$ the extinction time T_k is defined by

$$T_b(k) = \min\{t > 0 : Y_b(t) = 0\}$$

Let $G_k(t) := P(T_b(k) < t)$ be the CDF of $T_b(k)$. Given $Y_b(0) = k$, the extinction probability, q_k , is given $q_k = G_k(\infty)$. The CDF of $T_b(1)$ is obtained from [9, Eq. (7.3), p. 104] and is given by

$$G_1(t) = \frac{1 - e^{-\mu(1-\rho)t}}{1 - \rho e^{-\mu(1-\rho)t}}, \quad t \geq 0, \quad (10)$$

where we recall that $\rho = \lambda N_c / \mu$. From (10) we find

$$q_1 = \min\{1, 1/\rho\}. \quad (11)$$

In other words, the extinction will be certain iff $\rho \leq 1$. Since all objects behave independently of each other we have $q_k = q_1^k = \min\{1, 1/\rho^k\}$ and

$$G_k(t) = G_1(t)^k = \left(\frac{1 - e^{-\mu(1-\rho)t}}{1 - \rho e^{-\mu(1-\rho)t}} \right)^k, \quad t \geq 0. \quad (12)$$

In particular, if $\rho = 1$,

$$G_k(t) = \left(\frac{\mu t}{1 + \mu t} \right)^k, \quad t \geq 0. \quad (13)$$

B. Expected time to extinction

Assume that $\rho < 1$ (extinction is certain). The expected extinction time is equal to $E[T_b(k)] = \int_0^\infty (1 - G_k(t))dt$. In particular

$$E[T_b(1)] = -\frac{\log(1-\rho)}{\mu\rho}. \quad (14)$$

Let us now come back to the original process \mathbf{Y} . Define $T(y_0) := \inf\{t : Y(t) = 0\}$, the first time when the file has disappeared from the network given that $Y(0) = y_0$. When $Y(0) = Y_b(0) = y_0$, Proposition 2 implies that

$$P(T(y_0) > t) = P(Y(t) > 0) \leq P(Y_b(t) > 0) = G_{y_0}(t).$$

In particular $E[T(y_0)] \leq E[T_b(y_0)]$, so that $E[T(1)] \leq -\frac{\log(1-\rho)}{\mu\rho}$ from (14) for $\rho < 1$.

IV. MEAN-FIELD APPROXIMATION

In this section we investigate the behavior of the process \mathbf{Y} defined in Section II-C as N , the number of peers, gets large. We first show that this behavior (to be made more precise) is well approximated by a deterministic limit solution of an ODE, an approach known as mean-field approximation. See [6] for the theory and [10], [16], [12] for recent applications in the area of file sharing systems.

Like in Section II-B we assume that the pairwise contact rate, λ , is of the form $\lambda = \beta/N$ with $\beta > 0$. We recall that the initial state of \mathbf{Y} is given by

$$Y(0) = Ny_0, \quad X_c(0) = Nx_{c,0}, \quad X_f(0) = Nx_{f,0} \quad (15)$$

with $y_0 + x_{c,0} + x_{f,0} = 1$. [The analysis below holds under the weaker condition $\lim_{N \rightarrow \infty} N^{-1}(Y(0), X_c(0), X_f(0)) = (y_0, x_{c,0}, x_{f,0})$.]

Let $v_1 = (1, -1, 0)$, $v_2 = (-1, 0, 0)$ and $v_3 = (0, 0, -1)$. Denote by $g(Y, Y + v_i)$, $i = 1, 2, 3$, the non-zero transition rates of the process Markov process \mathbf{Y} out of state $Y = (Y_1, Y_2, Y_3)$. We have (cf. (5)-(7))

$$g(Y, Y + v_1) = \frac{\beta}{N} Y_1 Y_2, \quad g(Y, Y + v_2) = \mu Y_1, \quad g(Y, Y + v_3) = \frac{\beta}{N} Y_1 Y_3$$

which can be rewritten as

$$g(Y, Y + v_i) = Nf\left(\frac{Y}{N}, v_i\right), \quad i = 1, 2, 3 \quad (16)$$

where $f(u, v_1) = \beta u_1 u_2$, $f(u, v_2) = \mu u_1$ and $f(u, v_3) = \beta u_1 u_3$ for $u = (u_1, u_2, u_3)$.

We may therefore use Theorem 3.1 in [6] (it is easily that conditions (3.2)-(3.4) in [6] are satisfied) to obtain that the rescaled process $N^{-1}\mathbf{Y}$ converges in probability as $N \rightarrow \infty$, uniformly on all *finite* intervals $[0, T]$, to the solution (y, x_c, x_f) , $0 \leq y, x_c, x_f, y + x_c + x_f \leq 1$, of the system of ODEs

$$\frac{d}{dt} \begin{pmatrix} y \\ x_c \\ x_f \end{pmatrix} = \begin{pmatrix} y(\beta x_c - \mu) \\ -\beta y x_c \\ -\beta y x_f \end{pmatrix} \quad (17)$$

with initial condition $(y_0, x_{c,0}, x_{f,0})$.

In particular, for any finite t the solution y, x_c, x_f of (17) will approximate the fraction of peers with the file, the fraction of cooperative peers without the file and the fraction of free-riders, respectively, at time t . The accuracy of this approximation will increase with N , the total number of peers.

A. Peers that never receive the file: a phase transition

The fraction of cooperative peers x_c and the fraction of free-riders x_f that do not have the file monotonically decrease (this is true also for the original system) to some limit values. They can continue decreasing until there are no copies of the file in the system, namely until $y = 0$.

The first question we wish to address is whether these limits are close to 0 or are large. In other words, we wish to know whether all (or almost all) peers interested in the file are able to obtain it or not. If the answer is no, then we shall be interested in computing the fraction of peers that never receive the file.

Let $\theta := \beta/\mu$. From the first two equations in (17) we obtain x_c as

$$\frac{dy}{dx_c} = -1 + \frac{1}{\theta x_c}. \quad (18)$$

The solution of this differential equation is

$$x_c + y = \theta^{-1} \ln x_c + \phi(\theta) \quad (19)$$

where $\phi(\theta) := x_{c,0} + y_0 - \theta^{-1} \ln x_{c,0}$. Let y^{max} be the maximum ratio of cooperative peers with the file. According to the first equation in (17), y^{max} is reached when $x_c = \theta^{-1}$ if $\theta > 1$ and is expressed as

$$y^{max} = -\theta^{-1}(1 + \ln \theta) + \phi(\theta). \quad (20)$$

When $\theta \leq 1$, y^{max} is reached when $x_c = x_{c,0}$ (i.e. at time $t = 0$). On the other hand, as $t \rightarrow \infty$ y is approaching 0 (since we have assumed that $\mu > 0$) so that, from (19), $x_c(\infty)$ satisfies the equation

$$x_c(\infty) - \theta^{-1} \ln(x_c(\infty)) - \phi(\theta) = 0. \quad (21)$$

It is easily seen that this equation has a unique solution in $(0, x_{c,0})$ (note that $x_c(t) \leq x_{c,0}$ for any t since x_c is non-increasing from the second equation in (17)). From (17) we find that $x_f(t) = \frac{x_{f,0}}{x_{c,0}} x_c(t)$ for all t .

As recalled earlier the mean-field approximation only holds for finite t and there is therefore no guarantee that it will hold when $t = \infty$, namely, that $N\mathbf{Y}^{-1}$ will converge in probability to $(0, x_c(\infty), x_{f,0}x_c(\infty)/x_{c,0})$ as $N \rightarrow \infty$. However, due to the particular structure of the infinitesimal generator of \mathbf{Y} this convergence takes place as shown in [13, Sec. 5.2] (Hint: consider the rescaled Markov process $\tilde{\mathbf{Y}} := \{(\tilde{Y}(t), \tilde{X}_c(t), \tilde{X}_f(t)), t \geq 0\}$ with generator $\tilde{g}(\cdot, \cdot) = g(\cdot, \cdot)/Y_1$ and same state-space as \mathbf{Y} , so that starting from the same initial condition the terminal values of $X_c(t)$ and $\tilde{X}_c(t)$ (resp. $X_f(t)$ and $\tilde{X}_f(t)$) will have the same distribution. The mean-field approximation for $\tilde{\mathbf{Y}}$ shows that the solution of the associated ODE's is given by $(0, x_{c,0}e^{-\beta\tau}, x_{f,0}e^{-\beta\tau})$ for any $t \geq \tau$, with τ the unique solution in $(0, \infty)$ of $x_{c,0} + y_0 = x_{c,0}e^{-\beta\tau} + \mu\tau$, from which the result follows).

In summary, as N is large, the fraction of cooperative (resp. free riders) peers which will never receive the file is approximated by $x_c^{min} := x_c(\infty)$ (resp. $x_f^{min} := x_{f,0}x_c^{min}/x_{c,0}$) where $x_c(\infty)$ can be (numerically) calculated from (21).

We are interested in whether there is an abrupt change in content availability (i.e. $x_c(\infty)$) with the parameter θ . Obviously, if θ is 0, all the cooperative peers that do not have the file at time 0 will never receive it. To find a phase

transition, we approximate $\log(x_c(\infty))$ in (21) by using its Taylor extension at $x_{c,0}$ and obtain

$$x_c(\infty) \approx \left(\frac{1}{\theta} - x_{c,0} - y_0 \right) + \frac{1}{\theta} \left(\frac{x_c(\infty)}{x_{c,0}} - 1 \right)^2 / \left(\frac{1}{\theta x_{c,0}} - 1 \right).$$

Since the expression $\frac{1}{\theta} \left(\frac{x_c(\infty)}{x_{c,0}} - 1 \right)^2$ is bounded, the phase transition happens at $\theta = 1/x_{c,0}$.

Despite the similarity in the definitions of ρ in Section III and of $\theta x_{c,0}$ in the present section, the phase transition at $\rho = 1$ is different in nature from that at $\theta x_{c,0} = 1$. The former indicates whether or not the file will be extinct while the latter will drastically impact the final size of the torrent.

Figure 1 displays the mapping $\log_{10}(\theta x_{c,0}) \rightarrow x_c^{\min}$ for $x_{c,0} \in \{0.01, 0.1, 0.3, 0.5, 0.9\}$ and $y_0 = 0.05$. The curves for x_c^{\min} are monotonically decreasing in $x_{c,0}$ (the curve that intersects the vertical axis close to 1 is the one corresponding to $x_{c,0} = 0.01$, and so on.). For each curve we note the existence of a phase transition at $\theta x_{c,0} = 1$, which is more pronounced as the ratio of free riders increases.




Fig. 1. Ratio of cooperative peers (as N is large) that never receive the file as a function of $\log_{10}(\theta x_{c,0})$.

B. Combining the branching and the epidemic model

The mean-field approximation is accurate for large N if the initial state scales with N linearly. In the case that N is very large but the initial condition does not scale with N (e.g. $Y(0) = 1$, $X_c(0) + X_f(0) = N - 1$), we can do the following. Fix some N_0 much smaller than N but larger than 1. Use the branching process approximation until the number of peers with the file is N_0 . Then, switch to the epidemic model. (For the branching process, we recall that given that there is no extinction, the population size grows exponentially fast.)

V. CONTROL ACTIONS AGAINST P2P NETWORKS

In this section, we first investigate the major findings in the analysis of content availability. A set of control actions are proposed to protect copyrighted files against P2P file sharing.

A. Observations on file availability

Before proposing the counteractions against illegal P2P swarms, we investigate the impact of measures on file availability. The main question is how does a decrease or increase in one of the system parameters affect measures such as

- the size of the torrent: the fraction of those who are interested in the file and are able to get a copy of it. This can be seen as a global availability measure.
- the extinction probability or the expected extinction time,
- the maximum availability: the maximum number of copies that can be found simultaneously in the system. This can be viewed as an instantaneous availability measure.

According to the analysis in Sections III-IV, all above measures depend on the ratio $\frac{\lambda}{\mu}$ (or $\frac{\beta}{\mu}$ equivalently). A small ratio $\frac{\lambda}{\mu}$ means a poor availability of the file. However, the contact rate λ is an intrinsic parameter of P2P swarms that can hardly be changed technically. An even more challenging problem is that there usually exist several illegal publishers residing in the system for a very long time. They aim to spread the copyright protected file as wide as possible in the P2P swarm. To combat with undesirable file sharing, we present two methods, the cooperation control and the pollution attack. The former is to discourage the degree of cooperation of peers with the file. The latter introduces a number of polluters before the file dissemination begins, which can be found in [22] due to the page limit of this paper.

B. Control of cooperation

We introduce the cooperation control to prevent the dissemination of copyrighted files. We aim to reduce the degree of cooperation (i.e. increasing μ) so that the delay of obtaining the file is increased. To achieve this goal, the content owner can invest a certain amount of money in the very beginning to discourage the cooperation of peers. The cooperation control does not contradict with our opposition of collaborative solutions such as flat tax. In fact, we are focusing on this unilateral action of the content owner against unauthorized file dissemination.

We consider the same model as in Section II-C but we now assume that all peers are cooperative and that there is a number $Y_N^* > 0$ of permanent publishers, where the subscript N refers to the total number of peers in the system at time $t = 0$. The pairwise contact rate is $\lambda = \beta/N$. Denote by α the investment level of the content owner against P2P networks. The departure rate is an increasing function of α , denoted by $\mu(\alpha)$. We denote by $Y_N(t)$ the number of non-permanent publishers and by $X_N(t)$ the number of peers without the file at time t . Observe that $Y_N(0) + X_N(0) = N - Y_N^*$. If $\lim_{N \rightarrow \infty} Y_N(0)/N = y_0$ and $\lim_{N \rightarrow \infty} X_N(0)/N = x_0$, which implies that $\lim_{N \rightarrow \infty} Y^*/N = 1 - x_0 - y_0 := y^*$, then, by Kurtz's result [6], the rescaled process $N^{-1}(Y(t), X(t))$ converges in probability as $N \rightarrow \infty$, uniformly on all finite intervals $[0, T]$, to the solution of the ODEs

$$\frac{d}{dt} \begin{pmatrix} y \\ x \end{pmatrix} = \begin{pmatrix} \beta(y + y^*)x - \mu(\alpha)y \\ -\beta(y + y^*)x \end{pmatrix} \quad (22)$$

with initial state (y_0, x_0) . From now on we will assume that $y^* > 0$.

Consider an arbitrary peer without the file at time $t = 0$ and denote by T_N the time that elapses before it receives it. Let $P_N(t) := P(T_N < t)$. Similarly to [12, Page 6] we find

$$\frac{dP_N(t)}{dt} = \beta(1 - P_N(t)) \frac{E[Y_N(t)] + Y_N^*}{N}. \quad (23)$$

Solving for $P_N(t)$ gives

$$P_N(t) = 1 - e^{-\beta \int_0^t \frac{E[Y_N(s)] + Y_N^*}{N} ds}, \quad t \geq 0. \quad (24)$$

Hence

$$E[T_N] = \int_0^\infty (1 - P_N(t)) dt = \int_0^\infty e^{-\beta \int_0^t \frac{E[Y_N(s)] + Y_N^*}{N} ds} dt.$$

From the above we know that $(E[Y_N(t)] + Y_N^*)/N \rightarrow y + y^*$ as $N \rightarrow \infty$ for every $t > 0$, so that from (24)

$$\lim_{N \rightarrow \infty} (1 - P_N(t)) = e^{-\beta \int_0^t (y(s) + y^*) ds} \quad (25)$$

for every $t > 0$. On the other hand, $\lim_{N \rightarrow \infty} Y_N^*/N = y^*$ implies that for $0 < \epsilon < y^*$ there exists N_ϵ such that $Y_N^*/N > y^* - \epsilon$ for all $N > N_\epsilon$. Therefore, from (24),

$$1 - P_N(t) \leq e^{-\beta(y^* - \epsilon)t} \quad (26)$$

for $N > N_\epsilon$, $t > 0$. Since the r.h.s of (26) is integrable in $[0, \infty)$, (25) and (26) allow us to apply the bounded convergence theorem to conclude that

$$\bar{T}(\alpha) := \int_0^\infty \lim_{N \rightarrow \infty} (1 - P_N(t)) dt = \int_0^\infty e^{-\beta \int_0^t (y(s) + y^*) ds} dt.$$

The objective of the content owner is to choose an investment level $\alpha \geq 0$ which will maximize its utility

$$h(\alpha) := \bar{T}(\alpha) - \alpha. \quad (27)$$

To understand the impact of cooperation control on the delay, we present numerical studies in Section VII.

VI. P2P WITH A FIXED REQUEST RATE PER NODE

A. Model

In this section we will consider a slight variation of the model in [5]: there are N peers at time $t = 0$, at least one of them having a file. Each peer without the file sends a request for the file to another peer selected at random. These requests are initiated at Poisson rate $\lambda > 0$. It is assumed that a peer with the file leaves the system after an exponentially distributed random duration with rate $\mu > 0$. All these rvs are mutually independent. Let $Y(t)$ (resp. $X(t)$) be the number of peers with the file (resp. without the file) at time t . We have $Y(0) + X(0) = N$ with $Y(0) \geq 1$. Under the above assumptions $\mathbf{Z} := \{(Y(t), X(t)), t \geq 0\}$ is a Markov process on the set $\mathcal{E} := \{(y, x) \in \{0, 1, \dots, N\}^2 : 0 \leq y + x \leq N\}$. Let $q(z, z')$, $z = (y, x), z' = (y', x') \in \mathcal{E}$, denote its generator. Non-zero transition rates are given by

$$\begin{pmatrix} Y(t) \\ X(t) \end{pmatrix} \rightarrow \begin{pmatrix} Y(t) + 1 \\ X(t) - 1 \end{pmatrix} \text{ with rate } \frac{\lambda Y(t) X(t)}{Y(t) + X(t)}, \quad (28)$$

$$\begin{pmatrix} Y(t) \\ X(t) \end{pmatrix} \rightarrow \begin{pmatrix} Y(t) - 1 \\ X(t) \end{pmatrix} \text{ with rate } \mu Y(t). \quad (29)$$

This model differs from our previous model in that the rate of increase is normalized by the total number of peers in the system. More precisely, the rate in (28) follows from the fact that with probability $Y(t)/(Y(t) + X(t))$ a peer without the file will contact a peer with a file at time t (the latter implicitly assumes that a peer may contact itself as otherwise this probability would be $Y(t)/(Y(t) + X(t) - 1)$; the reason for doing this will next become apparent. Note that this assumption will have no effect when N gets large) so that the total rate of increase of the number of peers with the file is $\lambda Y(t) X(t)/(Y(t) + X(t))$ at time t .

The same model is considered in [5] with the difference that in [5] there is one permanent publisher, thereby implying that all peers will receive the file. These authors show that that the

mean broadcast time is $O(N)$ if $\lambda < \mu$ and is $O(\log(N))$ if $\lambda > \mu$. Thus, there is a phase transition at $\lambda = \mu$.

In this section we will instead focus on (i) the fraction of peers that will receive the file (in the absence of a permanent publisher this fraction is not always equal to 1) and on (ii) the maximum torrent size (maximum number of copies of the file at one time) as N is large. In both cases we will show the existence of phase transitions.

Our analysis will use Kurtz's theorem [6, Thm 3.1] like in Section IV. Note, however, that both metrics (i) and (ii) above require to use the mean-field limit as $t \rightarrow \infty$, something that Kurtz's result does not cover.

To overcome this difficulty, we will use the same argument as in [13] (see also Section IV where this argument was already used), taking advantage of the particular structure of the infinitesimal generator of the process \mathbf{Z} . More specifically, it is seen that the generator of \mathbf{Z} writes in the form $q(z, z') = y\tilde{q}(z, z')$ for $z = (y, x), z' = (y', x') \in \mathcal{E}$, where non-zero transition rates are given in (28)-(29).

Let $\tilde{\mathbf{Z}} = \{(\tilde{Y}(t), \tilde{X}(t)), t \geq 0\}$ be a Markov process with generator $\tilde{g}(z, z')$ and state-space \mathcal{E} (same state-space as \mathbf{Z}).

Since $\tilde{\mathbf{Z}}$ has been obtained by changing the time-scale of \mathbf{Z} , the final values of $\tilde{X}(t)$ and of $X(t)$ will have the same distribution (note that the final state of $\tilde{Y}(t)$ and $Y(t)$ is always zero since states $(0, \cdot)$ are all absorbing states) and so will have the maximum torrent size.

Since the generator $\tilde{g}(z, z')$ can be written as $\tilde{g}(z, z') = Nf(z/N, z'/N)$ (this is where the assumption that a peer may contact itself is useful) and since conditions (3.2)-(3.4) in [6, Thm 3.1] are clearly satisfied, we may apply [6, Thm 3.1] to obtain that, at any finite time t , $N^{-1}(\tilde{Y}(t), \tilde{X}(t))$ converges in probability as $N \rightarrow \infty$ to the solution (\tilde{y}, \tilde{x}) , $0 \leq \tilde{y}, \tilde{x} \leq 1$, $\tilde{y} + \tilde{x} \leq 1$, of the ODEs

$$\dot{\tilde{y}} = -\mu + \lambda\tilde{x}/(\tilde{x} + \tilde{y}), \quad \dot{\tilde{x}} = -\lambda\tilde{x}/(\tilde{x} + \tilde{y}) \quad (30)$$

given that $\lim_{N \rightarrow \infty} N^{-1}(\tilde{Y}(0), \tilde{X}(0)) = (\tilde{y}(0), \tilde{x}(0))$. Let $(y_0, x_0) := (\tilde{y}(0), \tilde{x}(0))$. We will assume that $0 < y_0 < 1$ and $y_0 + x_0 = 1$ (the case $y_0 = 0$ (resp. $y_0 = 1$) has no interest since it corresponds to a P2P network with no file at any time (resp. where all peers have the file at time $t = 0$)).

B. Phase transitions

Define $\xi := \lambda/\mu$. Adding both ODEs in (30) yields $\tilde{y}(t) + \tilde{x}(t) = -\mu t + 1$. Plugging this value back into (30) gives $\tilde{x}(t) = x_0(1 - \mu t)^\xi$ for $0 \leq t < 1/\mu$ and, by continuity, $\tilde{x}(t) = x_0(1 - \mu t)^\xi$ for $0 \leq t \leq 1/\mu$ with $\tilde{x}(1/\mu) = 0$.

In order to approximate the fraction of peers which will never receive the file as N is large, one needs to find the first time $\tau > 0$ where either $\tilde{x}(\tau) = 0$ or $\tilde{y}(\tau) = 0$. This time τ is easy to find as shown below.

We already know that $\tilde{x}(t) > 0$ for $0 \leq t < 1/\mu$ and $\tilde{x}(1/\mu) = 0$ so that we only need to focus on the zeros of $\tilde{y}(t)$ in $[0, 1/\mu]$. By writing \tilde{y} as $\tilde{y}(t) = (1 - \mu t)(1 - x_0(1 - \mu t)^{\xi-1})$ we conclude that the smallest zero of \tilde{y} in $[0, 1/\mu]$ is $(1 - x_0^{1/(1-\xi)})/\mu$ if $\xi < 1$ and is $1/\mu$ if $\xi \geq 1$. Therefore, $\tau = (1 - x_0^{1/(1-\xi)})/\mu > 0$ if $\xi < 1$ and $\tau = 1/\mu$ if $\xi \geq 1$. Introducing this value of τ in $\tilde{x}(t)$ yields $\tilde{x}(\tau) = x_0^{1/(1-\xi)}$ if $\xi < 1$ and $\tilde{x}(\tau) = 0$ if $\xi \geq 1$. In other words, as N is large,

all peers will get the file if $\xi \geq 1$ and a fraction $x_0^{1/(1-\xi)}$ of them will not if $\xi < 1$. In other words we observe a phase transition at $\xi = 1$: all peers will get the file if $\xi \geq 1$ and a fraction $\xi \rightarrow x_0^{1/(1-\xi)}$ will not if $\xi < 1$.

Let us now turn to the maximum torrent size. As N is large it will be approximated by the maximum of \tilde{y} over the interval $[0, \tau]$. A straightforward analysis of the mapping $t \rightarrow \tilde{y}(t)$ in $[0, \tau]$ shows that

- it is decreasing if $\xi \leq 1$ or if $\xi > 1$ and $\xi x_0 \leq 1$ – these conditions can be merged into the single condition $\xi \leq 1/x_0$ – so that its maximum, y^{\max} , is given by $y^{\max} = y_0 = 1 - x_0$,
- it is unimodal (first increasing then decreasing) if $\xi > 1/x_0$, with its maximum reached at $t_1 := (1 - (\xi x_0)^{1-\xi})/\mu$ and given by

$$y^{\max} = (x_0)^{1/(1-\xi)} \xi^{\xi/(1-\xi)} (\xi - 1) > 0. \quad (31)$$

In summary, as N is large, the maximum torrent size is approximated by Ny^{\max} with y^{\max} given in (31) if $\xi > 1/x_0$ and $y^{\max} = 1 - x_0$ if $\xi \leq 1/x_0$. This shows another phase transition (see Fig. 2) at $\xi = 1/x_0$ (i.e. at $\lambda x_0 = \mu$) in the sense that the torrent is maximum at $t = 0$ if $\xi \leq 1/x_0$ and is maximum at a later time if $\xi > 1/x_0$.




Fig. 2. Maximum torrent size over N (as N is large) as a function of ξ for $x_0 = 0.95$ (left figure) and $x_0 = 0.8$ (right figure).

VII. NUMERICAL RESULTS

This section has two goals: to investigate the accuracy of the approximations developed in the previous sections (to be made more precise) and to study the impact of measures against non-authorized uploading or downloading on the file availability in P2P swarming systems. Due to lack of space, we will not report any numerical result for the P2P model considered in Section VI; we will instead focus on the P2P model introduced in Section II-C and on its branching and mean-field approximations developed in Sections III and IV, respectively (Fig. 3-10), as well as on the optimization problem set in Section V (Fig. 11).

For each set of parameters, between 200 and 1000 discrete-event simulations of the Markov model in Section II-C have been run. In each figure (except in Fig. 5-6 where only simulation results are displayed and in Fig. 11 where only mean-field results are shown) both simulation and approximation results are reported for the sake of comparison. Let $r := (Y(0) + X_c(0))/N$ be the ratio of cooperative peers at time $t = 0$ and recall that $N_c = X_c(0)$ (see Section II). The total number of peers, N , at time $t = 0$ is equal to 400 in Fig. 3-6, to 300 in Fig. 7-8 and to 500 in Fig. 11.

A. File extinction time and the branching approximation

In this section we focus on Fig. 3-8. Fig. 3 (resp. Fig 4) compares the CDF of the extinction time obtained by simulation and by the branching approximation in (12) when $Y(0) = 1$ (resp. $Y(0) = 3$), $\lambda = 6 \cdot 10^{-3}$, $\mu = 1$, and for two values of r ($r = 0.6$ implies that $N_c = 239$ and $X_f(0) = 160$, $r = 1$ implies that there are no free riders ($X_f(0) = 0$) and $N_c = 399$). Note that $\rho = \lambda N_c / \mu$ is close to 2.4 when $r = 1$ and is close to 1.43 when $r = 0.6$. In all cases, the simulation and the branching approximation are in close agreement up to a certain time (time T_B in Fig. 3) which, interestingly, corresponds to the extinction time in the branching model. After this time, the extinction of the file in the Markov model increases sharply (the larger r the larger the increase). In other words, the extinction of the file in the original Markov model has two modes, an *early* extinction mode and a *late* extinction mode. The former occurs when the file disappears before the dissemination has reached its peak value (i.e. most peers do not get the file) and the latter when most peers leave the network with the file. One may also check that the branching approximation provides an upper bound for the CDF of the extinction time, as predicted by Proposition 2. Last, we note that when there are less cooperative peers ($r = 0.6$) the file lifetime is prolonged (see e.g. point D in Fig. 3 where simulation curves for $r = 1$ and $r = 0.6$ cross each other); this can be explained by the fact that there are less contact opportunities between cooperative peers when $r = 0.6$. The main difference between Fig. 3 and Fig 4 lies in the increase of the probability of the late extinction that is steeper with three initial seeds ($Y(0) = 3$) than with one initial seed ($Y(0) = 1$).

Simulation results in Fig. 5-6 exhibit the same *early-late* extinction pattern as in Fig. 3-4; they have been obtained for $\lambda = 25 \cdot 10^{-4}$ and for two different values of μ , r and $Y(0)$.

Fig. 7-8 show the expected time to extinction as a function of the pairwise contact rate λ , in the case of an early extinction (i.e. for small values of λ), for $\mu = 1$ and for two values of r . The curves "Model" display the mapping $\lambda \rightarrow E[T_b(k)]$, with $E[T_b(k)]$ the expected extinction time in the branching process given $Y(0) = k$ (see Section III). We observe an excellent match between the simulation and the branching approximation thereby showing that the latter works well for early file extinction. Also note that having three seeds instead of one greatly extends the expected extinction time.

B. File availability and the mean-field approximation

We now look at the fraction of peers that will not acquire the file. We assume that $Y(0) = 10$ and we recall that $N = 300$. Fig. 9 (resp. Fig.10) displays this fraction as a function of λ (resp. μ) for two different values of r ($r = 1$ corresponding to $X_c(0) = 290$ and $r = 0.5$ corresponding to $X_c(0) = 140$ and $X_f(0) = 150$). In each figure, both simulation and mean-field approximation results are reported. The fraction of peers without the file is a decreasing function of the pairwise contact rate λ and an increasing function of the cooperation degree μ . The mean-field approximation is obtained as the unique solution $x_c(\infty)$ in $(0, x_c(0))$ of equation (21) where the initial condition of the ODEs (17) is given by $(y_0, x_c(0), x_f(0)) =$




Fig. 3. CDF of extinction time for $Y(0) = 1$, $\lambda = 0.006$, $\mu = 1$.




Fig. 4. CDF of extinction time for $Y(0) = 3$, $\lambda = 0.006$, $\mu = 1$.




Fig. 5. CDF of extinction time for $Y(0) = 1$ and different μ .




Fig. 6. CDF of extinction time for $Y(0) = 3$ and different μ .




Fig. 7. Early extinction time as a function of λ with $Y(0) = 1$.




Fig. 8. Early extinction time as a function of λ with $Y(0) = 3$.




Fig. 9. Fraction of peers without the file as a function of λ .




Fig. 10. Fraction of peers without the file as a function of μ .




Fig. 11. Investment vs. utility of content owner.

$(Y(0)/N, X_c(0)/N, X_f(0)/N)$. In both figures we observe a remarkable agreement between the simulation and the mean-field results (relative errors never exceed 2% when all peers are cooperative ($r = 1$) and never exceed 7% when half of the peers are free riders ($r = 0.5$)). We also note that the fraction of peers without the file considered as a function of λ (resp. μ) is larger (resp. smaller) when $r = 0.5$ than when $r = 1$; this is of course not surprising since, unlike cooperative peers, free riders do not contribute to the file dissemination.

C. Action against unauthorized file downloading

We now evaluate the impact of actions against unauthorized file downloading. For that, we use the framework developed in Section V. Since the simulations in [22] show that, for large N , $\bar{T}(\alpha)$ in Section V is a good approximation of the expected time, T_N , needed for an arbitrary peer to get the

file we only consider the utility function $h(\alpha)$ (see (27)). We assume that the cooperation degree $\mu(\alpha)$ is given by $\mu(\alpha) = \mu\alpha$. There are 500 peers ($N = 500$) at time $t = 0$ including two persistent publishers ($Y^* = 2$). We assume that $Y(0) = 0$ so that $X(0) = N - Y^* = 498$. The initial condition of the ODE (22) is $(y_0, x_0) = (0, X_N(0)/N)$ with $y^* = 2/N$. Fig. 11 displays the mapping $\alpha \rightarrow h(\alpha)$, for two different values of β and μ . We observe that a small investment cannot obviously postpone the expected delivery delay of the file, resulting in a decreased utility. As the investment grows, the utility of the content owner increases significantly. The curves in Fig. 11 also show how large an investment has to be to counteract P2P illegal downloading. Note the content owner can still have an increased utility when the ratio $\beta/\mu(\alpha)$ is greater than one, as the utility is maximized across all curves when the ratio $\beta/\mu(\alpha)$ lies between two and three.

VIII. RELATED WORK

There has been a number of work on the mathematical studies of structured and unstructured P2P-based content distribution. A seminal work can be found in [14]. The authors propose a continuous-time branching process to analyze service capacity (i.e. maximum rate of downloading) and a coarse-grain Markov model to characterize the steady state of downloading rate. In [11], Qiu and Srikant propose a fluid model composed of ordinary differential equations to describe the dynamics of BitTorrent systems. Authors in [15] further propose a novel fluid model based on stochastic differential equations. This new model also extends [11] to multi-classes system and is able to describe chunk availability. Mundinger et al. [17] propose a deterministic scheduling algorithm to achieve the optimal makespan for a structured system which requires global knowledge. A coupon model is put forward in [16] to investigate the effectiveness of a generic P2P file sharing system.

Recently, the process of file dissemination has attracted a lot of attentions. Clévenot et al. adopt a hybrid approach (fluid and stochastic) to analyze Squirrel, a P2P cooperative web cache in [18]. In [5] Queija et al. study the scaling law of mean broadcasting time in a closed P2P swarm with constant request rate. Authors in [21] formulate a ball-and-urn model to characterize the “flash crowd” effect in a closed P2P networks. The content provided by P2P networks such as music, movies and software are usually unauthorized. Content provided are therefore inclined to combat illegal downloading/uploading via technical solutions. Authors of [2] and [19] propose a M/G/ ∞ queueing model to access the efficiency of non-cooperative measures against unauthorized downloading. Authors in [20] propose a similar but elegant queueing model to study the impact of bundling strategy of file availability and downloading rate. Our general model is inspired by the one in [5]. However, it differs from [5], [11], [14], [19] in four ways: 1) we are studying the transient behavior; 2) a peer can initiate a number of random contacts, instead of one, with other peers; 3) we observe several phase transitions in response to system parameters; 4) we adopt Markov branching process and mean-field approaches to characterize the file dissemination model comprehensively.

IX. CONCLUSION

In this paper we have proposed to use the theory of continuous time branching process as well as of the dynamics of epidemics in order to study the transient behavior of torrents that occur in P2P systems. The use of these tools allowed us to compute the probability of early extinction of the torrent as well as the expected time until that extinction, the availability of a file in the system, the maximum availability and when it occurs, and the size of the torrent. This is used for analyzing the impact of measures to decrease non-authorized Internet access to copyrighted files. We identify regimes in which the performance measures are quite sensitive to such measures and others in which the measures have very limited impact. In particular, we present two counteractions against unauthorized file sharing in the presence of illegal publishers.

Our methodology can be extended to analyze file bundling that serves as a positive action of file dissemination.

REFERENCES

- [1] M. Piatek, T. Kohno and A. Krishnamurthy, “Challenges and Directions for Monitoring P2P File Sharing Networks - or - Why My Printer Received a DMCA Takedown Notice,” *USENIX HotSec’08*, July 29, 2008, San Jose, CA, USA.
- [2] S. Wong, E. Altman and M. Ibrahim, “P2P Networks: The Interplay between Legislation and Information Technology,” *Proc. Networking and Electronic Commerce Research Conf.* (NAEC) Italy, October 8-11, 2009, Riva del Garda, Italy. See also INRIA Research Report RR-6889, version 2, June 2009.
- [3] J. G. Steffensen. Deux problèmes du calcul des probabilités. *Annales de l’Institute Henri Poincaré*, 1933.
- [4] D. Mollison, *Epidemic Models: their Structure and Relation to Data*. Ed. D. Mollison, Cambridge University Press, Cambridge, 1995.
- [5] R. Nunez-Queija and B. J. Prabhu. “Scaling Laws for File Dissemination in P2P Networks with Random Contacts,” *Proc. 16th Int. Workshop on Quality of Service (IWQoS 2008)*, June 2-4, 2008, Univ. Twente, Enschede, The Netherlands.
- [6] T. G. Kurtz, “Solutions of Ordinary Differential Equations as Limits of Pure Jump Markov Processes,” *Journal of Applied Probability*, Vol. 7, No. 1, pp. 49-58, April 1970.
- [7] J. K. Hal, *Ordinary Differential Equations*, Pure and Applied Mathematics XXI, Robert E. Krieger Publishing Company, Huntington, New York, 1980.
- [8] R. Groenevelt, P. Nain and G. Koole, “The Message Delay in Mobile Ad Hoc Networks,” *Performance Evaluation*, Vol. 62, Issues 1-4, pp. 210-228, October 2005.
- [9] T. E. Harris, *The theory of Branching Processes*, Dover Publications, 1989 Printing, 1963 Edition.
- [10] A. Chaintreau, J.-Y. Le Boudec and N. Ristanovic, “The Age of Gossip: Spatial Mean Field Regime,” *Proc. ACM Sigmetrics 2009 Conf.*, June 15-19, 2009, Seattle, WA, USA.
- [11] D. Qiu and R. Srikant, “Modeling and Performance Analysis of BitTorrent-like Peer-to-Peer Networks”, *Proc. of ACM Sigcomm 2004 Conf.*, Aug. 30 - Sep. 3, 2004, Portland, OR, USA.
- [12] X. Zhang, G. Neglia, J. Kurose and D. Towsley, “Performance Modeling of Epidemic Routing,” *Computer Networks*, Vol. 51, No. 10, pp. 2859-2891, 2007.
- [13] R. Darling and J. Norris. “Differential Equation Approximations for Markov Chains,” *Probability Surveys*, Vol. 5, pp. 37-79, 2008.
- [14] X. Yang and G.D. Veciana, “Service Capacity of Peer to Peer Networks”, *Proc. of IEEE Infocom 2004 Conf.*, March 7-11, 2004, Hong Kong, China.
- [15] B. Fan, D.M. Chiu and J.C.S. Lui, “On the Performance and Availability of BitTorrent-like P2P File Sharing Systems”, *Proc. 14th Int. Workshop on Quality of Service (IWQoS’06)*, June 19-21, 2006, New Haven, CT, USA.
- [16] L. Massoulié and M. Vojnović, “Coupon Replication Systems,” *Proc. ACM Sigmetrics 2005 Conf.*, June 6-10, 2005, Banff, Alberta, Canada.
- [17] J. Mundinger, R. Weber and G. Weiss, “Optimal scheduling of peer-to-peer file dissemination”, *Journal of Scheduling*, Vol.11, No.2, 2008.
- [18] F. Clévenot and P. Nain, “A Simplified Model for the Analysis of the Squirrel Peer-to-Peer Caching System”, *Proc. of IEEE Infocom 2004 Conf.*, March 7-11, 2004, Hong Kong, China.
- [19] E. Altman, S. Wong and J. Rojas-Mora, “P2P Business and Legal Models for Increasing Accessibility to Popular Culture,” *Proc. of Int. Conference on Digital Business*, June 17 - 19, 2009, London, UK.
- [20] D.S. Menasche, A.A. Rocha, B. Li and D. Towsley, “Bundling Builds Availability in BitTorrent”, *Proc. of ACM CoNext 2009*, December 1-4, 2009, Rome, Italy.
- [21] F. Simatos, P. Robert and F. Guillemin, “A Queueing System for Modeling a File Sharing Principle”, *Proc. of ACM Sigmetrics 2008 Conf.*, June 2-6, 2008, Annapolis, MD, USA.
- [22] E. Altman, P. Nain, A. Shwartz and Y.D. Xu, “Predicting the Impact of Measures Against File Sharing Networks on the Transient Behaviors”, *Technical Report*, INRIA, 2010. <http://arxiv.org/abs/1009.2230>