A Stochastic Epidemiological Model and a Deterministic Limit for BitTorrent-Like Peer-to-Peer File-Sharing Networks

George Kesidis[1,⋆], Takis Konstantopoulos[2,⋆⋆], and Perla Sousi[3]

[1] CSE and EE Departments, Pennsylvania State University, University Park, PA 16802, USA
gik2@psu.edu

[2] School of Mathematical & Computer Sciences and the Maxwell Institute for Mathematical Sciences, Heriot-Watt University, Edinburgh EH14 4AS, UK
takis@ma.hw.ac.uk

[3] Statistical Laboratory, Centre for Mathematical Sciences, University of Cambridge, Wilberforce Road, Cambridge CB3 0WB, UK
P.Sousi@statslab.cam.ac.uk

Abstract. We propose a stochastic model for a file-sharing peer-to-peer network which resembles the popular BitTorrent system: large files are split into chunks and a peer can download or swap from another peer only one chunk at a time. We exhibit the fluid and diffusion limits of a scaled Markov model of this system and look at possible uses of them to draw practical conclusions.

1 Introduction

Peer-to-peer (p2p) activity continues to represent a very significant fraction of overall Internet traffic, 44% by one recent account [4]. BitTorrent [1,2,8,21,18,9,19] is a widely deployed p2p file-sharing network which has recently played a significant role in the network neutrality debate. Under BitTorrent, peers join "swarms" (or "torrents") where each swarm corresponds to a specific data object (file). The process of finding the peers in a given swarm to connect to is typically facilitated through a centralised "tracker". Recently, a trackerless BitTorrent client has been introduced that uses distributed hashing for query resolution [16].

For file sharing, a peer is typically uploads upload pieces ("chunks") of the file to other peers in the swarm while downloading his/her missing chunks from them. This chunk swapping constitutes a transaction-by-transaction incentive for peers to cooperate (i.e., trading rather than simply download) to disseminate data objects. Large files may be segmented into several hundred chunks, all of which the peers of the corresponding warm must collect, and in the process disseminate their own chunks before they can reconstitute the desired file and possibly leave the file's swarm.

⋆ Supported by an NSF grant.
⋆⋆ Corresponding author; supported by an EPSRC grant.

E. Altman and A. Chaintreau (Eds.): NET-COOP 2008, LNCS 5425, pp. 26–36, 2009.

In addition to the framework in which data objects are segmented into chunks to promote cooperation through swapping, there is a system whereby the rate at which chunks are uploaded is assessed for any given transaction, and peers that allocate inadequate bandwidth for uploading may be "choked" [14,17]. Choking may also be applied to peers who, by employing multiple identities (sybils), abuse BitTorrent's system of allowing newly arrived peers to a swarm to just download a few chunks (as they clearly cannot trade what they simply do not as yet possess). BitTorrent can also rehabilitate peers by (optimistically) unchoking them. In the following, we do not directly consider upload bandwidth and related choking issues.

In this paper, we motivate a deterministic epidemiological model of file dissemination for peer-to-peer file-sharing networks that employ BitTorrent-like incentives, a generalisation of that given in [10,11]. Our model is different from those explored in [15,21,18] for BitTorrent, and we compute different quantities of interest. Our epidemiological framework, similar to that we used for the spread of multi-stage worms [12], could also be adapted for network coding systems. In [9], the authors propose a "fluid" model of a single torrent/swarm (as we do in the following) and fit it to (transient) data drawn from aggregate swarms. The connection to branching process models [21,8] is simply that ours only tracks the number of active peers who possess or demand the file under consideration, i.e., a single swarm. Though our model is significantly simpler than that of prior work, it is derived directly from an intuitive transaction-by-transaction Markov process modelling file-dissemination of the p2p network and its numerical solutions clearly demonstrate the effectiveness of the aforementioned incentives. A basic assumption in the following is that peers do not distribute bogus files (or file chunks) [20].

2 The Stochastic Model

A file is represented as a set F of size n, the elements of which are called chunks. Consider a large networked "swarm" of N nodes called peers. Each peer possesses a certain (possibly empty) subset A of F. As time goes by, this peer interacts with other peers, the goal being to enlarge his set A until, eventually, the peer manages to collect all n chunks of F. The interaction between peers can either be a download or a swap; in both cases, chunks are being copied from peer to peer and are assumed never lost. A peer will stay in the network as long as he does not possess all chunks. After collecting everything, sooner or later a peer departs or switches off. By splitting the desired file into many chunks we give incentives to the peers to remain active in the swarm for long time during which other peers will take advantage of their possessions.

2.1 Possible Interactions

We here describe how two peers, labelled A, B, interact. The following types of interactions are possible:

1. **Download:** Peer A downloads a chunk i from B. This is possible only if A is a strict subset of B. If $i \in B$ then, after the downloading A becomes $A' = A \cup \{i\}$ and but B remains B because it since it gains nothing from A. Denote this interaction by: $(A \leftarrow B) \rightsquigarrow (A', B)$. The symbol on the left is supposed to show the type of interaction and the labels before it, while the symbol on the right shows the labels after the interaction.
2. **Swap:** Peer A swaps with peer B. In other words, A gets a chunk j from B and B gets a chunk i from A. It is required that j is not an element of A and i not an element of B. We denote this interaction by $(A \leftrightarrows B) \rightsquigarrow (A', B')$, where $A' = A \cup \{j\}$, $B' = B \cup \{i\}$. We thus need $A \setminus B \neq \varnothing$ and $B \setminus A \neq \varnothing$.

2.2 Notation

The set of all combinations of n chunks, which partition F, is denoted by $\mathscr{P}(F)$, where $|\mathscr{P}(F)| = 2^n$ and the empty set is included. We write $A \subset B$ (respectively, $A \subsetneq B$) when A is a subset (respectively, strict subset) of B. We (unconventionally) write

$$A \sqsubset A' \text{ when } A \subset A' \text{ and } |A' - A| = 1.$$

If $A \cap B = \varnothing$, we use $A + B$ instead of $A \cup B$; if $B = \{b\}$ is a singleton, we often write $A + b$ instead of $A + \{b\}$. If $A \subset B$ we use $B - A$ instead of $B \setminus A$. We say that

$$A \text{ relates to } B \text{ (and write } A \sim B) \text{ when } A \subset B \text{ or } B \subset A;$$

if this is not the case, we write $A \not\sim B$. Note that $A \not\sim B$ if and only if two peers labelled A, B can swap chunks. The space of functions (vectors) from $\mathscr{P}(F)$ into $\mathbb{Z}_+$ is denoted by $\mathbb{Z}_+^{\mathscr{P}(F)}$. The stochastic model will take values in this space. The deterministic model will evolve in $\mathbb{R}_+^{\mathscr{P}(F)}$. We let $e_A \in \mathbb{Z}_+^{\mathscr{P}(F)}$ be the vector with coordinates

$$e_A^B := \mathbf{1}(A = B), \quad B \in \mathscr{P}(F).$$

For $x \in \mathbb{Z}_+^{\mathscr{P}(F)}$ or $\mathbb{R}_+^{\mathscr{P}(F)}$ we let $|x| := \sum_{A \in \mathscr{P}(F)} |x^A|$.

2.3 Defining the Rates of Individual Interactions

We follow the logic of stochastic modelling of chemical reactions or epidemics and assume that the chance of a particular interaction occurring in a short interval of time is proportional to the number of ways of selecting the peers needed for this interaction [13]. Accordingly, the interaction rates *must* be given by the formulae described below.

Consider first finding the rate of a download $A \leftarrow B$, where $A \subsetneq B$, when the state of the system is $x \in \mathbb{Z}_+^{\mathscr{P}(F)}$. There are x^A peers labelled A and x^B labelled B. We can choose them in $x^A x^B$ ways. Thus the rate of a download $A \leftarrow B$ that results into A getting *some* chunk from B should be proportional to $x^A x^B$. However, we are interested in the rate of the *specific* interaction $(A \leftarrow B) \rightsquigarrow (A', B)$, that turns A into a specific set A' differing from A by one single chunk

($A \sqsubset A'$); there are $|B - A|$ chunks that A can download from B; the chance that picking one of them is $1/|B - A|$. Thus we have:

$$(DR) \quad \begin{cases} \textit{the rate of the download } (A \leftarrow B) \rightsquigarrow (A', B) \textit{ equals } \beta x^A \dfrac{x^B}{|B - A|}, \\ \textit{as long as } A \sqsubset A' \subset B, \end{cases}$$

where $\beta > 0$.

Consider next a swap $A \leftrightarrows B$ and assume the state is x. Picking two peers labelled A and B (provided that $A \not\sim B$) from the population is done in $x^A x^B$ ways. Thus the rate of a swap $A \leftrightarrows B$ is proportional to $x^A x^B$. So if we *fix* two chunks $i \in A \setminus B, j \in B \setminus A$ and specify that $A' = A + j, B' = B + i$, then the chance of picking i from $A \setminus B$ and j from $B \setminus A$ is $1/|A \setminus B||B \setminus A|$. Thus,

$$(SR) \quad \begin{cases} \textit{the rate of the swap } (A \leftrightarrows B) \rightsquigarrow (A', B') \textit{ equals } \gamma \dfrac{x^A x^B}{|A \setminus B||B \setminus A|}, \\ \textit{a long as } A \sqsubset A', \quad B \sqsubset B', \quad A' - A \subset B, \quad B' - B \subset A, \end{cases}$$

where $\gamma > 0$.

2.4 Deriving the Markov Chain Rates

Having defined the rates of each individual interaction we can easily define rates $q(x, y)$ of a Markov chain in continuous time and state space $\mathbb{Z}_+^{\mathscr{P}(F)}$ as follows.

Define functions $\lambda_{A,A'}, \mu_{A,B} : \mathbb{R}^{\mathscr{P}(F)} \to \mathbb{R}$ by:

$$\lambda_{A,A'}(x) := \left[\beta x^A \sum_{C: C \supset A'} \frac{x^C}{|C - A|} \right] \mathbf{1}(A \sqsubset A') \tag{1a}$$

$$\mu_{A,B}(x) := \gamma \frac{x^A x^B}{|A \setminus B||B \setminus A|} \mathbf{1}(A \not\sim B). \tag{1b}$$

Consider also constants $\delta \geq 0$ and $\alpha^A \geq 0$ for $A \in \mathscr{P}(F)$, i.e., $\alpha \in \mathbb{R}_+^{\mathscr{P}(F)}$. The transition rates of the closed conservative Markov chain are given by:

$$q(x, y) := \begin{cases} \lambda_{A,A'}(x), & \text{if } y = x - e_A + e_{A'} \\ \mu_{A,B}(x), & \text{if } \begin{cases} y = x - e_A - e_B + e_{A'} + e_{B'} \\ A \sqsubset A', B \sqsubset B', A' - A \subset B, B' - B \subset A, \end{cases} \\ \alpha^A & \text{if } y = x + e_A \\ \delta x^F & \text{if } y = x - e_F \\ 0, & \text{for any other value of } y \neq x, \end{cases} \tag{2}$$

where x ranges in $\mathbb{Z}_+^{\mathscr{P}(F)}$.

A little justification of the first two cases is needed: that $q(x, x - e_A - e_B + e_{A'} + e_{B'}) = \mu_{A,B}(x)$ is straightforward. It corresponds to a swap, which is only possible when $A \sqsubset A', B \sqsubset B', A' - A \subset B, B' - B \subset A$. The swap rate was defined by (SR). To see that $q(x, x - e_A + e_{A'}) = \lambda_{A,A'}(x)$ we observe that a peer labelled A can change its label to $A' \sqsupset A$ by downloading a chunk from some set C that contains A', so we sum the rates (DR) over all these possible individual interactions to obtain the first line in (2). We can think of having Poisson process of arrivals of new peers at rate $|\alpha|$, and that each arriving peer is labelled A with probability $\alpha^A/|\alpha|$. Peers can depart, by definition, only when they are labelled F and it takes an exponentially distributed amount of time (with mean $1/\delta$) for a departure to occur. Thus, $q(x, x - e_F) = \delta x^F$. We shall let Q denote the generator of the chain, i.e. $\mathsf{Q}f(x) = \sum_y (f(y) - f(x))q(x, y)$, when f is an appropriate functional of the state space.

Definition 1 ($\texttt{BITTORRENT}[x_0, n, \alpha, \beta, \gamma, \delta]$)**.** *Given* $x_0 \in \mathbb{Z}_+^{\mathscr{P}(F)}$ *(initial configuration),* $n = |F| \in \mathbb{N}$ *(number of chunks),* $\alpha \in \mathbb{R}_+^{\mathscr{P}(F)}$ *(arrival rates),* $\beta > 0$ *(download rate),* $\gamma \geq 0$ *(swap rate),* $\delta \geq 0$ *(departure rate) we let* $\texttt{BITTORRENT}[x_0, n, \alpha, \beta, \gamma, \delta]$ *be a Markov chain* $(X_t, t \geq 0)$ *with transition rates* (2) *and* $X_0 = x_0$*. We say that the chain (network) is* <u>*open*</u> *if* $\alpha^A > 0$ *for at least one* A *and* $\delta > 0$*; it is* <u>*closed*</u> *if* $\alpha^A = 0$ *for all* A*; it is* <u>*conservative*</u> *if it is closed and* $\delta = 0$*; it is* <u>*dissipative*</u> *if it is closed and* $\delta > 0$*.*

In a conservative network, we have $q(x, y) = 0$ if $|y| \neq |x|$ and so $|X_t| = |X_0|$ for all $t \geq 0$. Here, the actual state space is the simplex $\{x \in \mathbb{Z}_+^{\mathscr{P}(F)} : |x| = N\}$, where $N = |X_0|$. It is easy to see that the state e_F is reachable from any other state, but all rates out of e_F are zero. Hence a conservative network has e_F as a single absorbing state.

In a dissipative network, we have $|X_t| \leq |X_0|$ for all $t \geq 0$. Here the state space is $\{x \in \mathbb{Z}_+^{\mathscr{P}(F)} : |x| \leq N\}$, where $N = |X_0|$. It can be seen that a dissipative network has many absorbing points.

In an open network, there are no absorbing points. On the other hand, one may wonder if certain components can escape to infinity. This is not the case:

Lemma 1. *If* $\alpha^F > 0$ *then the open* $\texttt{BITTORRENT}[x, n, \beta, \gamma, \alpha, \delta]$ *is positive recurrent Markov chain.*

Proof. (sketch) If $\alpha^F > 0$, $\delta > 0$ the Markov chain is irreducible. The remainder of the proof is based on a the construction of a simple Lyapunov function: $V(x) := |x|$, for which it can be shown that there is a bounded set of states K such that $\sup_{x \notin K} (\mathsf{Q}V)(x) < 0$. Perhaps the easiest way to see this is by appealing to the stability of the corresponding ODE limit; see Theorem 1 below and [7]. □

3 Macroscopic Description: Fluid Limit and Diffusion Approximation

Analysing the Markov chain in its original form is complicated. We thus resort to a first-order approximation by an ordinary differential equation (ODE). Let $v(x)$ be the vector field on $\mathbb{R}_+^{\mathscr{P}(F)}$ with components $v^A(x)$ defined by

$$v^A(x) = \alpha^A - x^A\big(\beta\varphi_d^A(x) + \gamma\varphi_s^A(x)\big) + \beta \sum_{B:A\subset B} \frac{\psi_d^A(x)x^B}{1+|B\setminus A|} + \gamma \sum_{B:A\not\subset B} \frac{\psi_s^{A,B}(x)x^B}{1+|B\setminus A|} - \delta x^F \mathbf{1}(A=F), \quad (3)$$

where

$$\varphi_d^A(x) := \sum_{B\supset A} x^B, \quad \varphi_s^A(x) := \sum_{B\not\supset A} x^B$$
$$\psi_d^A(x) := \sum_{a\in A} x^{A-a}, \quad \psi_s^{A,B}(x) := \sum_{a\in A\cap B} x^{A-a} \quad (4)$$

Consider the differential equation

$$\dot{x} = v(x) \text{ with initial condition } x_0. \quad (5)$$

Consider the sequence of stochastic models `BITTORRENT[`$X_{N,0}, n, N\alpha, \frac{\beta}{N}, \frac{\gamma}{N}, \delta$`]` for $N \in \mathbb{N}$ and let $X_{N,t}$ be the corresponding jump Markov chain.

Theorem 1. *There is a has a unique smooth (analytic) solution to (5), denoted by x_t for $t \geq 0$. Also, if there is an $x_0 \in \mathbb{R}_+^{\mathscr{P}(F)}$ such that $X_{N,0}/N \to x_0$ as $N \to \infty$, then for any $T, \varepsilon > 0$,*

$$\lim_{N\to\infty} P\big(\sup_{0\le t\le T} |N^{-1}X_{N,t} - x_t| > \varepsilon\big) = 0.$$

Proof. See [11].

Next, let

$$Y_{N,t} := \sqrt{N}(X_{N,t}/N - x_t).$$

For each $y \in \mathbb{Z}_+^{\mathscr{P}(F)}$ let W_y be a standard one-dimensional Brownian motion; suppose that these Brownian motions are independent over y. Define the (time-inhomogeneous) Gaussian diffusion process Y by

$$dY(t) = \sum_y (y - x_t)\sqrt{q(x_t, x_t + y)}dW_y(t) + Dv(x_t)Y(t)dt,$$

where $Dv(x)$ is the matrix of partial derivatives of $v(x)$. Due to the form the rates (2), the first sum ranges over finitely many y and so only finitely many Brownian motions are needed.

Theorem 2. *If $\sqrt{N}(X_{N,t}/N - x_0) \to 0$ as $N \to \infty$, where $x_0 \in \mathbb{R}_+^{\mathscr{P}(F)}$, then the law of Y_N (as a sequence of probability measures in $D[0,\infty)$ with the topology of uniform convergence on compacta) converges weakly to the law of Y.*

Proof. We refer to [13] for the relevant arguments.

4 Examples

Suppose that F consists of $n = 2$ chunks. The limiting ODE is easily found to be:

$$\begin{aligned}
\dot{x}^{\varnothing} &= \alpha^{\varnothing} - \beta x^{\varnothing}(x^1 + x^2 + x^{12}) \\
\dot{x}^1 &= \alpha^1 - x^1(\beta x^{12} + \gamma x^2) + \beta x^{\varnothing}(x^1 + \tfrac{1}{2}x^{12}) \\
\dot{x}^2 &= \alpha^2 - x^2(\beta x^{12} + \gamma x^1) + \beta x^{\varnothing}(x^2 + \tfrac{1}{2}x^{12}) \\
\dot{x}^{12} &= \alpha^{12} + \beta(x^1 + x^2)x^{12} + 2\gamma x^1 x^2 - \delta x^{12}.
\end{aligned}$$

We look at its behaviour in three cases. To make things easier, assume that $\gamma = 0$.

4.1 Closed Conservative System: $\boldsymbol{\alpha^1 = \alpha^2 = \alpha^{12} = 0, \delta = 0}$

Letting $x = x^{\varnothing}, u = x^1 + x^2, w = x^{12}$, assuming that $x + u + w = 1$, and eliminating the variable w we obtain

$$\begin{aligned}
\dot{x} &= -\beta x(1 - x) \\
\dot{u} &= \beta u^2 - \beta u(1 - x) + \beta x(1 - x).
\end{aligned}$$

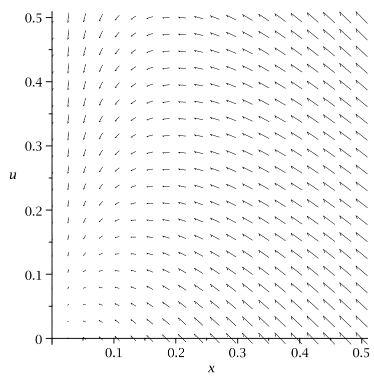

Fig. 1. Typical vector field plot for a closed conservative BitTorrent model

The vector field in the $x - u$ plane is depicted in Figure 1. The unique equilibrium point $x = 0, u = 0$ corresponds to $w = 1$, i.e. everybody possesses the full file. By solving the equation in x, substituting into the equation for u, we find the explicit solution

$$w_t = \frac{x_0 + (1 - x_0)e^{\beta t}}{x_0 + (1 - x_0)e^{\beta t} + x_0 \beta t + (1 - w_0)w_0^{-1}},$$

from which one can estimate the time required for w to reach an ε-neighbourhood of the equilibrium, which can be turned into an estimate for the original stochastic system.

4.2 Closed Dissipative System: $\alpha^1 = \alpha^2 = \alpha^{12} = 0, \delta > 0$

With $x = x^{\varnothing}, u = x^1 + x^2, w = x^{12}$ as before, change the time variable to $s = \beta t$, let $\rho = \delta/\beta$, and write x' for dx/ds, to obtain:

$$
\begin{aligned}
x' &= -x(u+w) \\
u' &= -uw + x(u+w) \\
w' &= uw - \rho w.
\end{aligned}
$$

Assume $x_0 + u_0 + w_0 = 1$, so that $x_t + u_t + w_t < 1$ for all $t > 0$. We cannot eliminate the variable w now since there is no obvious conserved quantity, but we can study the equilibria of the system. It is easily seen that the only equilibria are of the form $(0, u, 0)$ which are unstable if $u > \rho$ and stable if $u < \rho$. In terms of the original variables, the stable equilibria are $(x^{\varnothing}, x^1, x^2, x^{12}) = (0, x^1, x^2, 0)$, $0 \leq x^1 + x^2 < \rho$. This is as expected: since there is no swapping ($\gamma = 0$), the system eventually settles to a situation where there are peers with label 1 and peers with label 2. Had γ been positive, x^1, x^2 could not have simultaneously been positive in equilibrium.

4.3 Open System: $\alpha^1 = \alpha^2 = 0, \alpha^{12} = \lambda > 0, \delta > 0$

Choosing variables appropriately, we have

$$
\begin{aligned}
x' &= -x(u+w) \\
u' &= -uw + x(u+w) \\
w' &= \lambda + uw - \rho w.
\end{aligned}
$$

The system eventually settles to the unique stable equilibrium $(x, u, w) = (0, 0, \lambda/\rho)$. It is easily seen that the eigenvalues of the differential of the vector field at this point are $-\lambda/\rho$ and $-\rho$ (the first one has algebraic multiplicity 1 but geometric multiplicity 2), and so there is no possibility of spiralling.

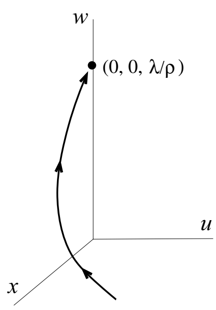

Fig. 2. Trajectory for open system

5 Performance Improvement in Presence of BitTorrent Incentives

To split or not to split? The answer is yes, but not always. We can address this question by looking at the behaviour of the ℓ_1-norm $|x^*|$ of the unique equilibrium point of an open system as the number of chunks (dimension) increases.

As an example, consider a situation where peers possessing nothing arrive at rate λ, download at rate β and depart at rate δ once they have the full file. This is described by

$$\dot{x}^{\varnothing} = \lambda - \beta x^{\varnothing} x^1$$
$$\dot{x}^1 = \beta x^{\varnothing} x^1 - \delta x^1.$$

The globally attracting stable equilibrium is given by $x^* = (\delta/\beta,\ \lambda/\delta)$.

Suppose now that we split into $n = 2$ chunks. Peers arrive and depart at the same rates but download at a rate $\widetilde{\beta} \geq \beta$ and swap at rate $\widetilde{\gamma}$. The new equilibrium is easily found to be

$$\widetilde{x}^* = \left(\frac{\delta}{\widetilde{\beta}}\Big(\frac{\delta}{\lambda}u + 1\Big)^{-1},\ \frac{u}{2},\ \frac{u}{2},\ \frac{\lambda}{\delta}\right),$$

where u is the positive number which solves

$$q(u) := u^2 + \frac{2\widetilde{\beta}\lambda}{\widetilde{\gamma}\delta}u - \frac{2\lambda}{\widetilde{\gamma}} = 0. \tag{6}$$

The following is shown in [11]:

Lemma 2

1. $\widetilde{x}^{*\varnothing} < x^{*\varnothing}$.
2. $|x^*| > |\widetilde{x}^*|$ *if and only if* $u < \widetilde{u}$, *where* $\widetilde{u}$ *is the unique positive number which satisfies*

$$\widetilde{q}(\widetilde{u}) := \widetilde{u}^2 - \Big(\frac{\delta}{\beta} - \frac{\lambda}{\delta}\Big)\widetilde{u} - \Big(\frac{\lambda}{\beta} - \frac{\lambda}{\widetilde{\beta}}\Big). \tag{7}$$

Furthermore, there exists a $\lambda_0 > 0$ *such that for all* $\lambda < \lambda_0$ (7) *holds.*

6 Conclusions and Open Problems

We proposed a stochastic model of a BitTorrent-like network and showed the existence of an ODE limit, along with a diffusion approximation. Several simulations [10] test the suitability of the model. Proofs of some of the results presented in this paper can be found in [11]. One can look at the ODE limit and, more specifically, its equilibria in order to obtain crude information about the stationary distribution of the original model. This requires proving a certain robustness result, as explained in [11, Sec. 5]. Rates of convergence to the ODE limit is another interesting open problem. Its solution requires estimating bounds on the vector field and its derivative. A criticism of the model is that the ODE lives in a high-dimensional space. One can reduce the dimension under some symmetry assumptions on the initial state and the parameters of the model [11, Sec. 7]. Another drawback is the Markovian assumption requires that times between transactions are exponential; this is clearly violated in practise (download times are typically heavy-tailed random variables). A quick fix of this problem is to represent the Markov process on a probability space supporting i.i.d. Poisson processes and then replace these by more general renewal processes.

References

1. BitTorrent, `http://www.bittorrent.com`
2. Cohen, B.: Incentives Build Robustness in BitTorrent. In: Workshop on Economics of Peer-to-Peer Systems, Berkeley, CA, USA (May 2003)
3. Daley, D.J., Gani, J.: Epidemic Modeling, an Introduction. Cambridge University Press, Cambridge (1999)
4. DC Info (Distributed Systems Newsletter) vol. XXII(8), June 30 (2008), `http://www.dcia.info/news/#Newsletters`
5. Ethier, S.N., Kurtz, T.G.: Markov Processes: Characterization and Convergence. Wiley, New York (1986)
6. Darling, R.W.R., Norris, J.R.: Differential equation approximations for Markov chains. Probability Surveys 5, 37–79 (2008)
7. Foss, S., Konstantopoulos, T.: An overview of some stochastic stability methods. Journal of the Operations Research Society of Japan 47, 275–303 (2004)
8. Ge, Z., Figueiredo, D.R., Jaiswal, S., Kurose, J., Towsley, D.: Modeling peer-to-peer file sharing systems. In: Proc. IEEE INFOCOM, San Francisco (April 2004)
9. Guo, L., Chen, S., Xiao, Z., Tan, E., Ding, X., Zhang, X.: Measurements, modeling and analysis of BitTorrent-like systems. In: Proc. Internet Measurement Conference (IMC) (October 2005)
10. Kesidis, G., Jin, Y., Mortazavi, B., Konstantopoulos, T.: An epidemiological model for file-sharing with BitTorrent-like incentives. In: Proc. IEEE GLOBECOM (November 2006)
11. Kesidis, G., Konstantopoulos, T., Sousi, P.: A stochastic epidemiological model and a deterministic limit for BitTorrent-like peer-to-peer file-sharing networks (November 2008), `http://arxiv.org/abs/0811.1003`
12. Kesidis, G., Vojnovic, M., Hamadeh, I., Jin, Y., Jiwasurat, S.: A Model of the Spread of Randomly Scanning Internet Worms that Saturate Access Links. In: ACM TOMACS (May 2008)
13. Kurtz, T.: Approximation of Population Processes. SIAM, Philadelphia (1981)
14. Legout, A., Liogkas, N., Kohler, E., Zhang, L.: Clustering and sharing incentives in BitTorrent systems. ACM SIGMETRICS Performance Evaluation Review 35(1) (June 2007)
15. Massoulié, L., Vojnovic, M.: Coupon replication systems. In: Proc. ACM SIGMETRICS, Banff, Alberta, Canada (2005)
16. Maymounkov, P., Mazieres, D.: Kademlia: a peer-to-peer information system based on the XOR metric. In: Proc. IPTPS, Cambridge, MA, USA (March 2002)
17. Mortazavi, B., Kesidis, G.: A peer-to-peer content-distribution game with a reputation-based incentive mechanism. In: Proc. IEEE Workshop on Information Theory and its Applications (ITA), UC San Diego (February 2006)

18. Qiu, D., Srikant, R.: Modeling and performance analysis of BitTorrent-like peer-to-peer networks. In: Proc. ACM SIGCOMM, Portland, Oregon (2004)
19. Turner, B.: Generalizing BitTorrent: how to build data exchange markets (and profit from them!) (January 2005), http://www.fractalscape.org/GeneralizingBitTorrent.htm
20. Walsh, K., Sirer, E.G.: Fighting peer-to-peer SPAM and decoys with object reputation. In: Proc. Workshop on Economics of Peer-to-Peer Systems (p2pecon), Philadelphia, PA (August 2005)
21. Yang, X., de Veciana, G.: Service capacity of peer to peer networks. In: Proc. IEEE INFOCOM, San Francisco (2004)