Modeling file-sharing with BitTorrent-like incentives

GEORGE KESIDIS
Pennsylvania State University
kesidis@engr.psu.edu

TAKIS KONSTANTOPOULOS
Heriot-Watt University
takis@ma.hw.ac.uk

PERLA SOUSI
Ecole Normale Supérieure
perlasousi@gmail.com

***Abstract*— We propose a new model for file-sharing peer-to-peer (P2P) networks that mimics the incentives provided by the popular BitTorrent system. In it, larger files are split into chunks and a peer can download or swap only one chunk at a time. We propose a Markov chain model in continuous time that resembles a stochastic epidemic/coagulation model. We prove that the Markov chain is approximated by a differential equation which, by itself, can give some rough information about the performance of the system. Finally, using this model, we explore the performance of BitTorrent-like incentives for an open system with peer departures and arrivals and a single file (torrent) with two chunks.**

Keywords: Internet, Internetworking

I. INTRODUCTION

BitTorrent [1], [4], [7], [19], [14], [8], [15] is a widely deployed peer-to-peer (P2P) file-sharing network wherein a peer is typically required to upload portions of the file to a set of peers while downloading missing pieces (or "chunks") from others. This requirement is a "memoryless" transaction-by-transaction incentive for peers to cooperate (i.e., upload not just download) to disseminate files. Large files may be segmented into several chunks thereby requiring peers to collect all chunks, and in the process disseminate their own, before they can reconstitute the desired file and possibly "leave" the system (torrent or swarm). The process of finding the peers to connect to is facilitated through a *centralized* "tracker." Recently, a trackerless BitTorrent client has been introduced that uses distributed hashing for query resolution [12].

In this paper, we motivate a deterministic epidemiological/coagulation model of file dissemination for peer-to-peer file-sharing networks that employ BitTorrent-like incentives, a generalization of that given in [9]. In [8], 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 [16], [19], [7] 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 modeling 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) [17].

This paper is organized as follows. In Section II, we provide a general Markov epidemiological model for a single swarm. The limiting deterministic ODE obtained by Kurtz's theorem is given in Section III. The benefits of BitTorrent-like incentives is then explored in Section IV for the case of a two-chunk torrent with exogenous arrivals in steady-state, i.e., between an initial phase where the torrent experiences potentially exponential growth in the peer population and the final phase where interest in the file dwindles and the peer population of the torrent drifts to zero. Finally, we conclude with a summary in Section V. Proofs, omitted for brevity in the following, can be found in [10].

II. THE MARKOVIAN MODEL

Suppose there are N peers in the system which can communicate with one another in all possible ways. There is a single file F split into n chunks. At each point of time, each peer possesses some (or none) of the chunks. Think of F as a set with n elements (chunks), and let the label of a peer be a subset of set F. We let $A, B, C, \dots$ denote such labels which range over the set $\mathscr{P}_n(F)$ of all subsets of F. We use the notation $A \subset B$, to mean that A is a subset of B (i.e. $A \subseteq B$) and $A \neq B$. The symbol $|A|$ denotes the cardinality of A. We write $\mathscr{P}_k(F)$ for the collection of sets of cardinality k. We use $A \sqsubset B$ to mean $A \subset B$ and $|B - A| = 1$; that is, $A \sqsubset B$ is the strict subset relation between two sets A, B with $A \in \mathscr{P}_k(F)$, $B \in \mathscr{P}_{k+1}(F)$ for some $0 \leq k < n$. We write $A + B$ for the disjoint union between two sets and $A - B$ for the difference ($A \setminus B$) of two sets when $B \subseteq A$. If B contains only one element b we write $A + b$, $A - b$ in lieu of $A + \{b\}$, $A - \{b\}$, respectively. Finally, we $A \not\sim B$ to denote the fact that $A \setminus B \neq \varnothing$ and $B \setminus A \neq \varnothing$.

A. Possible transactions

At each point of time one of the following transactions can take place:

Download: Peer A downloads a chunk c from B. This is possible only if $A \subset B$. If $c \in B$ then, after the downloading A becomes $A' = A + c$ and but B remains B because it has nothing to gain from A. We denote this transaction by $(A \leftarrow B) \rightsquigarrow (A', B)$. The symbol on the left is supposed to show what the labels are before the transaction, what type the transaction is, and the symbol on the right shows the labels after the transaction.

Swap: Peer A swaps with peer B. In other words, A gets a chunk from B and B gets a chunk from A. For this to happen we need that B have a chunk b that A does not have, and, vice versa, A has a chunk a that B does not possess. In other words we need $A \not\sim B$. After the transaction, A becomes $A+b$, and B becomes $B+a$. We denote this by $(A \leftrightarrows B) \rightsquigarrow (A', B')$.

Full swap: This is a special case of a swap that makes the two peers become identical after the transaction. For this to happen we need $|A \setminus B| = |B \setminus A| = 1$. After the transaction both peers attain the same labels: $A' = B' = A + (B \setminus A) = B + (A \setminus B)$. Thus, a full swap is denoted by $(A \leftrightarrows B) \rightsquigarrow (A', A')$.

B. The stochastic model

To describe a stochastic model, we let x^A denote the number of peers labelled A (at some point of time) and set

$$x = (x^A, A \in \mathscr{P}_n(F)) \in \mathbb{Z}_+^{2^n}.$$

We follow the logic of stochastic modelling of chemical reactions or epidemics and assume that the chance of a particular transaction occurring in a short interval of time is proportional to the number of ways of selecting the peers needed for this transaction [11]. According to this logic, the rates of the three types of transactions *must* be given by the formulae described below. **This said, note that we would not allow for the number of "seeds" x^F to transit to zero in the Markov chain as this is an absorbing state.**

a) Rates of individual transactions: Consider first finding the rate of a download $A \leftarrow B$, where $A \subset B$, when the state of the system is x. 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* transaction $(A \leftarrow B) \rightsquigarrow (A', B)$, 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: *the rate of the download* $(A \leftarrow B) \rightsquigarrow (A', B)$ *equals* $\beta_d x^A x^B / |B - A|$*, a long as* $A \subseteq B, \quad A \sqsubset A', \quad A' - A \subseteq B,$ where $\beta_d > 0$.

Consider next a swap $A \leftrightarrows B$ and assume the state is x. Picking two peers labelled A and B, $A \not\sim B$, 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 $a \in A \setminus B, b \in B \setminus A$ and specify that $A' = A + b, B' = B + a$, then the chance of picking a from $A \setminus B$ and b from $B \setminus A$ is $1/|A \setminus B||B \setminus A|$. Thus, *the rate of the swap* $(A \leftrightarrows B) \rightsquigarrow (A', B')$ *equals* $\beta_s x^A x^B / (|A \setminus B||B \setminus A|)$*, a long as* $A \sqsubset A', \quad B \sqsubset B', \quad A' - A \subseteq B, \quad B' - B \subseteq A,$ where $\beta_d > 0$.

This applies equally well to the full swap case: *the rate of the full swap* $(A \leftrightarrows B) \rightsquigarrow (A', A')$ *equals* $\beta_s x^A x^B$*, a long as* $A' - A = \{b\} \subseteq B, \quad A' - B = \{a\} \subseteq A, \quad a \neq b.$

b) Markov chain rates: Having defined the rates of each individual transaction we can easily define rates $q(x, y)$ of a Markov chain in continuous time and state space $\mathbb{Z}_+^{2^n}$ as follows. We let $e_A \in \mathbb{Z}_+^{2^n}$ be the vector with coordinates $e_A^B := \mathbf{1}(A = B)$, $B \in \mathscr{P}_n(F)$.

Let

$$\lambda_{A,A'}(x) := \beta_d x^A \sum_{C: C \supseteq A'} \frac{x^C}{|C - A|} \mathbf{1}(A \sqsubset A') \tag{1a}$$

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

Lemma II.1. *The transition rates of the closed conservative Markov chain are given by:*

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

We will let $(X_t, t \geq 0)$ denote a Markov chain with transition rates as above. Let $|x| := \sum_{A \in \mathscr{P}_n(F)} x^A$. Observe that, by definition, $|X_t| = |X_0| = N$ for all t, and that is why we refer to it as a closed conservative system: peers do not arrive or depart; they simply download or swap chunks and the number of peers is always N. Thus, the actual state space is the simplex

$$S_N^{2^n} := \{x \in \mathbb{Z}_+^{2^n} : |x| = N\}.$$

Note that the state e_F is reachable from any other state, but all rates out of e_F are zero. Hence the chain has e_F as a single absorbing state.

To include the phenomenon of "free-riding", i.e. of peers acting selfishly by departing once they obtain all the n chunks of the file, we introduce an additional transition with rate

$$q(x, x - e_F) = \delta x^F,$$

as long as $x^F > 0$, where $\delta \geq 0$ is the departure rate. Note that the system now is closed but not conservative. Indeed, $|X_t| \leq |X_0| = N$ for all t. Here the state space is

$$\{x \in \mathbb{Z}_+^{2^n} : |x| \leq N\}.$$

It can be seem that are many absorbing points and they always lie at the faces of T_N.

Let us take the special case where the file consists of a single chunk ($n = 1$), i.e., $x^1 \equiv x^F$ so that the state here is $x = (x^\varnothing, x^F)$. There is only one type of transaction possible: $(\varnothing \leftarrow 1) \rightsquigarrow (1, 1)$. Hence the rates are:

$$q\big((x^\varnothing, x^F), (x^\varnothing - 1, x^F + 1)\big) = \beta_d x^\varnothing x^F$$
$$q\big((x^\varnothing, x^F), (x^\varnothing, x^F - 1)\big) = \delta x^F.$$

This is the classical stochastic version of the Kermack-McKendrick model for a simple epidemic process [5]. Its absorbing points are states of the form $(x^\varnothing, 0)$. In epidemiological terminology, x^F is the number of (fully) infected individuals, whereas $x^\varnothing$ is the number of susceptible ones. Contrary to the epidemiological interpretation, infection *is* desirable, for infection is tantamount to downloading the file. Consider this simple system with parameter β_d/N instead of β_d, where N is the initial number of peers, and let $X_t^{(N)}$ be the corresponding Markov chain. We then have [5] that $N^{-1}X_t^{(N)}$ converges, as $N \to \infty$, in a strong sense to a smooth function $(x^\varnothing(t), x^F(t))$ which satisfies the differential equations:

$$\frac{dx^\varnothing}{dt} = -\beta_d x^\varnothing x^F \quad \text{and} \quad \frac{dx^F}{dt} = \beta_d x^\varnothing x^F - \delta x^\varnothing \quad (3)$$

This is a rough deterministic model (the fluid limit of the Markov chain) that captures some crude performance measures of the stochastic system.

III. Macroscopic description

We first derive an expression for the drift of the process. The drift is defined as (the vector field)

$$v(x) := \sum_y (y-x)q(x,y). \quad (4)$$

From the formula (2) for $q(x,y)$, we see that, unless $y - x = -e_A + e_{A'}$ or $y = -e_A - e_B + e_{A'} + e_{B'}$, we have $q(x,y) = 0$. Let us consider the second of the rates and rewrite it as

$$q(x, x - e_A, x - e_B + e_{A'} + e_B) = \mu_{A,B}(x)\delta_{A,A',B,B'},$$

where

$$\delta_{A,A',B,B'} := \mathbf{1}(A \sqsubset A', A' - A \subseteq B, B \sqsubset B', B' - B \subseteq A). \quad (5)$$

Notice that swapping A with B or A' with B' will not change the value of $x - e_A, x - e_B + e_{A'} + e_B$, so we need to make sure to take into account this change only once in the summation (4). If we *simultaneously* swap A with B *and* A' with B' then neither $x - e_A, x - e_B + e_{A'} + e_B$ nor the value $\mu_{A,B}(x)\delta_{A,A',B,B'}$ of the rate change because, obviously,

$$\mu_{A,B}(x)\delta_{A,A',B,B'} = \mu_{B,A}(x)\delta_{B,B',A,A'},$$

as readily follows from (1b) and (5). We now see that to swap A with B *without* swapping A' with B' is impossible (unless $A' = B'$). Indeed, it is an easy exercise that

$$\delta_{A,B,A',B'} = \delta_{B,A,A',B'} \;\Rightarrow\; A' = B'.$$

Taking into account this, we write

$$\begin{aligned} v(x) \;&=\; \sum_{A,A'} (-e_A + e_{A'})\lambda_{A,A'}(x) + \quad (6)\\ &\frac{1}{2} \sum_{A,B,A',B'} (-e_A - e_B + e_{A'} + e_{B'})\mu_{A,B}(x)\delta_{A,B,A',B'}, \end{aligned}$$

where the $1/2$ appears because each term must be counted exactly once. The variables A, A', B, B' in both summations are free to move over $\mathscr{P}_n(F)$ (but notice that restrictions have effectively been pushed in the definitions of $\lambda_{A,A'}$, $\mu_{A,B,}$ and $\delta_{A,B,A',B'}$).

Notice also that it is not necessary to specify that all components of x are nonzero: the definitions of $\lambda_{A,A'}(x)$, $\mu_{A,B}(x)$ automatically take care of the rates $q(x,y)$ when x is at the boundary; indeed, if $x^A = 0$ then $\lambda_{A,A'}(x) = 0$; similarly, if $x^A = 0$ or $x^B = 0$ then $\mu_{A,B}(x) = 0$.

Let $\varphi_d^A(x) := \sum_{B \supset A} x^B = \sharp$ peers from which an A-peer can download from, $\varphi_s^A(x) := \sum_{B \not\sim A} x^B = \sharp$ peers an A-peer can swap with, $\psi_d^A(x) := \sum_{a \in A} x^{A-a} = \sharp$ peers which can assume label A after a download, and $\psi_s^{A,B}(x) := \sum_{a \in A \cap B} x^{A-a} = \sharp$ peers which can assume label A after a swap a B-peer.

Lemma III.1. *The drift $v(x) = (v^A(x), \in \mathscr{P}_n(F))$, is given by*

$$\begin{aligned} v^A(x) \;&=\; -x^A\big(\beta_d\varphi_d^A(x) + \beta_s\varphi_s^A(x)\big) + \quad (7)\\ &\beta_d \sum_{B: A\subseteq B} \frac{\psi_d^A(x)x^B}{1+|B\setminus A|} + \beta_s \sum_{B:A\not\subseteq B} \frac{\psi_s^{A,B}(x)x^B}{1+|B\setminus A|}, \end{aligned}$$

for all $A \in \mathscr{P}_n(F)$ and all x in the state space.

Theorem III.1. *Consider the Markov chain $(X_t^{(N)}, t \geq 0)$ corresponding to the closed conservative model with rates as in* (2) *and parameters β_d/N and β_s/N instead of β_d and β_s, respectively. Let x_0 be an arbitrary point in $\mathbb{R}_+^{2^n}$ with L_1 norm $|x_0| = 1$, and assume that $X_0^{(N)}/N \to x_0$, a.s. as $N \to \infty$. Let $(x_t, t \geq 0)$ be the solution to the ODE*

$$\dot{x} = v(x)$$

starting from x_0. Then for all $t > 0$, and all $\varepsilon > 0$, $\lim_{N\to\infty} P(\sup_{0\leq s\leq t} |X_s^{(N)} - x_s| > \varepsilon) = 0$.

Note how x takes on real components in the theorem statement and that, by construction, we have $\langle v(x), 1\rangle = 0$, i.e., $v(x)$ is orthogonal to the vector whose components are all 1.

Example ODE for $n = 2$: Here A can take 4 values: $\varnothing$, $\{1\}$, $\{2\}$, $\{1,2\}$. According to Theorem III.1,

$$\begin{aligned} \dot{x}^\varnothing &= -\beta_d x^\varnothing(x^1 + x^2 + x^{12})\\ \dot{x}^1 &= -x^1(\beta_d x^{12} + \beta_s x^2) + \beta_d x^\varnothing(x^1 + \tfrac{1}{2}x^{12}) \quad (8)\\ \dot{x}^2 &= -x^2(\beta_d x^{12} + \beta_s x^1) + \beta_d x^\varnothing(x^2 + \tfrac{1}{2}x^{12})\\ \dot{x}^{12} &= \beta_d(x^1 + x^2)x^{12} + 2\beta_s x^1 x^2. \end{aligned}$$

IV. Evaluating BitTorrent-like incentives for a two-chunk torrent/swarm with exogenous peer arrivals

We can accommodate arrivals of new peers demanding the file in our model for a *one*-chunk torrent *without* BitTorrent-like incentives by using

$$\frac{dx_0^\varnothing}{dt} \;=\; -\beta x_0^F x_0^\varnothing + \lambda \quad (9)$$

instead of (3), where the additional term $\lambda > 0$, the (mean) arrival rate of new peers, is due to the transition from $(X^F, X^\varnothing)$ to $(X^F, X^\varnothing + 1)$ at rate $\Lambda := \lambda \mathsf{E}N$, i.e., using the steady-state mean number of peers in the torrent. Note that we have dropped the subscript "*d*" from the β parameter in this case of a one-chunk torrent for the purposes of subsequent comparison.

Now consider a two-chunk torrent. By defining $w = x^1 + x^2 + x^F$ note that we can write

$$\frac{ds}{dt} \;=\; -\beta_d ws + \lambda, \quad (10)$$

and adding the deterministic ODEs (8) gives

$$\frac{dw}{dt} \;=\; \beta_d ws - \delta x^F \;\geq\; \beta_d ws - \delta w \tag{11}$$

where we have also added a peer departure term with parameter δ. Since

$$\begin{aligned} 0 &= -\beta x_0^{\varnothing}(\infty) x_0^F(\infty) + \lambda \;\text{ and} \\ 0 &= \beta x_0^{\varnothing}(\infty) x_0^F(\infty) - \delta x_0^F(\infty), \end{aligned}$$

we get that $x_0^{\varnothing}(\infty) = \delta/\beta \equiv \rho$ and $x_0^F(\infty) = \lambda/\delta$.

Directly from (8) and (11), we can similarly get that

$$\begin{aligned} 0 &= -\beta_d w(\infty) s(\infty) + \lambda \;\text{ and} \\ 0 &\geq \beta_d w(\infty) s(\infty) - \delta w(\infty). \end{aligned}$$

Thus,

$$s(\infty) \;\leq\; \frac{\delta}{\beta_d} \;\leq\; \frac{\delta}{\beta} \;=\; x_0^{\varnothing}(\infty)$$

when $\beta_d \geq \beta$.

To show $x^F(\infty) = \lambda/\delta$ we will simply invoke Little's theorem [18] here, i.e., the steady-state time a peer is a seed is $1/\delta$ and λ is the net arrival rate of new peers or seeds.

Again according to Little's theorem, the steady-state time a peer remains a seed prior to departure without BitTorrent-like incentives is

$$\frac{1}{\beta x_0^F(\infty) x_0^{\varnothing}(\infty)} x_0^F(\infty) \;=\; \frac{1}{\delta} \tag{12}$$

i.e., the departure rate δ as defined.

So, by (12) and Little's theorem, the steady-state mean time that a peer waits to acquire its first chunk (from the time s/he arrives) is

$$\frac{1}{\lambda} s(\infty) \;\leq\; \frac{1}{\lambda} x_0^{\varnothing}(\infty) \;=\; \frac{\delta}{\lambda\beta}$$

where $\lambda\delta/\beta$ is the mean time a peer waits to acquire the entire file without BitTorrent-like incentives.

Theorem IV.1. *If $\beta_d = \beta$ and*

$$y \;\equiv\; \frac{\delta}{\beta} - \frac{\lambda}{\delta} \;\leq\; 0, \tag{13}$$

then the steady-state mean time to become a seed (acquire the entire file) without BitTorrent-like dynamics is not shorter than that of BitTorrent-like networks.

Note that the assumption $y \leq 0$ is equivalent to $x_0^F(\infty) \leq x_0^{\varnothing}(\infty)$. If we want to consider the case where $y \geq 0$, we need to restrict $x^1(\infty) \geq y/2$. We can interpret this as bounds on β_s:

Theorem IV.2. *If $\beta_d = \beta$ and*

$$0 \;\leq\; \beta_s \;\leq\; \frac{2\beta_d \lambda^2}{(\delta y)^2} - \beta_d \tag{14}$$

then the steady-state mean time to become a seed (acquire the entire file) under the dynamics without BitTorrent-like incentives is not shorter than that of BitTorrent-like networks.

V. Summary

In this paper, we developed Markovian stratified epidemiological/coagulation models of the dissemination of a single popular file (torrent) by peer-to-peer file-sharing networks that employ BitTorrent-like incentives. The limiting deterministic ODE, by Kurtz's theorem, was then given. We then used this simple model in steady-state with a fixed peer arrival rate to evaluate the effect of such incentives by comparison with a system that does not segment files and always involves only a single file transfer per transaction, i.e., involves a client-peer and server-peer. That is, conditions were given for a two-chunk swarm under which steady-state complete download times were shorter using BitTorrent-like incentives.

References

[1] BitTorrent. www.bittorrent.com

[2] P. Antoniadis, C. Courcoubetis, and R. Mason. Comparing economic incentives in peer-to-peer networks. *Computer Networks*, **46**(1):133–146, 2004.

[3] Z. Chen, L. Gao, and K. Kwait. Modeling the spread of active worms. *Proc. IEEE INFOCOM*, 2003.

[4] B. Cohen. Incentives Build Robustness in BitTorrent. *Workshop on Economics of Peer-to-Peer Systems*, Berkeley, CA, USA, May 2003.

[5] D.J. Daley and J. Gani. *Epidemic Modeling, an Introduction*. Cambridge University Press, 1999.

[6] S.N. Ethier and T.G. Kurtz. *Markov Processes: Characterization and Convergence*. Wiley, New, York, 1986.

[7] Z. Ge, D. R. Figueiredo, S. Jaiswal, J. Kurose, and D. Towsley. Modeling peer-to-peer file sharing systems. *Proc. IEEE INFOCOM*, April 2004.

[8] L. Guo, S. Chen, Z. Xiao, E. Tan, X. Ding and X. Zhang, Measurements, modeling and analysis of BitTorrent-like systems. *Proc. Internet Measurement Conference (IMC)*, Oct. 2005.

[9] G. Kesidis, Y. Jin, B. Mortazavi, and T. Konstantopoulos. An epidemiological model for file-sharing with BitTorrent-like incentives. *Proc. IEEE GLOBECOM*, Nov. 2006.

[10] G. Kesidis, T. Konstantopoulos and P. Sousi. Modelling file sharing with BitTorrent-like incentives. *Technical Report*, Sept. 2006, //www.cse.psu.edu/∼kesidis/bittor.pdf

[11] T. Kurtz. *Approximation of Population Processes*. SIAM, Philadelphia, 1981.

[12] P. Maymounkov and D. Mazieres. Kademlia: A peer-to-peer information system based on the XOR metric. *Proc. IPTPS*, Cambridge, USA, March 2002.

[13] J.R. Norris. *Markov Chains*. Cambridge Univ. Press, 1997.

[14] D. Qiu and R. Srikant. Modeling and performance analysis of BitTorrent-like peer-to-peer networks. *Proc. ACM SIGCOMM*, Portland, Oregon, 2004.

[15] B. Turner. Generalizing BitTorrent: how to build data exchange markets (and profit from them!), Jan. 2005, www.fractalscape.org/GeneralizingBitTorrent.htm

[16] G. de Veciana and X. Yang. Fairness, incentives and performance in peer-to-peer networks. *Allerton Conference on Communication, Control and Computing*, 2003.

[17] K. Walsh and E.G. Sirer. Fighting peer-to-peer SPAM and decoys with object reputation. In *Proceedings of the Third Workshop on the Economics of Peer-to-Peer Systems (p2pecon)*, Philadelphia, PA, 2005.

[18] R.W. Wolff. *Stochastic Modeling and the Theory of Queues*. Prentice-Hall, Englewood Cliffs, NJ, 1989.

[19] X. Yang and G. de Veciana. Service capacity of peer to peer networks. In *Proc. IEEE INFOCOM*, 2004.

[20] B. Yu, M.P. Singh, and K. Sycara. Developing trust in large-scale peer-to-peer systems. *Proc. of First IEEE Symposium on Multi-Agent Security and Survivability*, 2004.