P2P Video-on-Demand: Steady State and Scalability

Samuli Aalto, Pasi Lassila
Aalto University School of Science and Technology
Email: firstname.lastname@tkk.fi

Niklas Raatikainen, Petri Savolainen, Sasu Tarkoma
Helsinki Institute for Information Technology
Email: firstname.lastname@hiit.fi

***Abstract*—The fundamental P2P principle that downloading peers help other peers can be applied in the context of video-on-demand. This represents a demanding application combining aspects of other well-known P2P applications, i.e., live streaming and traditional file sharing. We seek to provide insight on fundamental questions about the performance and scalability of the system. A deterministic fluid model is derived that explicitly takes into account the video transfer and playback phases. The analytical results are complemented with extensive simulations from the corresponding stochastic model, as well as traces from a more realistic BitTorrent simulator.**

I. Introduction

Modern peer-to-peer (P2P) content distribution networks such as BitTorrent [1] achieve scalability by utilizing the fundamental idea that the peers not only act as clients downloading content from other peers but also serve the other peers by uploading onwards the downloaded content. While P2P content distribution technologies are widely used for file sharing and live media streaming, they are less common in the closely related, but yet distinct case of *on-demand streaming of stored media* [2], of which a well-known example is *video-on-demand* (VoD). However, the popularity of client-server based VoD services such as YouTube raises the question, whether, and under which conditions, P2P technologies could be used to alleviate the high cost of building and maintaining the infrastructures of such systems.

In P2P VoD, the whole stored media file needs to be retrieved (cf., file sharing) at such a rate which allows, as soon as possible, the pieces to be played back in sequential order at the media playback rate (cf., live streaming). Thus, there are two phases, the *transfer* and the *playback* phases, which are partly overlapping. If the retrieval rate is sufficient, the playback phase extends beyond the transfer phase.

In this paper we analyze the performance and scalability of a P2P VoD system. More specifically, our contributions include:

- Presenting a fluid model that allows both download and upload constrained solutions and also explicitly takes into account the transfer and the playback phases.
- Validating the fluid model by comparing it to results from simulating the corresponding stochastic model and to traces from a more realistic BitTorrent simulator.
- Utilizing the analytical and the numerical results to give closed-form conditions (i) under which QoS is acceptable and (ii) under which the system is scalable.

Notably, our results show that to achieve a piece retrieval rate exceeding the video viewing rate, the efficiency parameter is critical. On the other hand, the better-performing, download-constrained steady state can always be achieved with a high enough fraction of altruistic seeds or a sufficient number of permanent seeds.

The paper is organized as follows. Section II discusses the related work. The model and the analytical results are presented in Section III, followed by a synthesis of the results given in Section IV. Section V includes the numerical validations, and conclusions are given in Section VI.

II. Related work

P2P live streaming has been analyzed recently in many papers, see, e.g., [3] analyzing the impact of various peer/piece selection policies or [4], where a stochastic model is developed for analyzing the probability that peers can download the stream at a sufficiently high rate. However, our objective is to consider the case of on-demand streaming.

Qiu and Srikant [5] develop a simple fluid model to analyze the performance of a BitTorrent-like P2P *file sharing* system under a steady flow arrival scenario. The key parameters of the model are the arrival rate of new peers, λ, the efficiency of P2P file sharing, η, and the departure rate of seeds, γ, which reflects their selfishness aspect. The efficiency parameter η comprises the effect of the piece selection policy, the number of downloading connections, and the number of pieces. The model also assumes a homogeneous peer population with joint download and upload rates, c and μ, respectively.

Inspired by [5], Parvez et al. [2] develop fluid models to analyze the performance of BitTorrent-like P2P *on-demand streaming*. They start from the premise that the system is upload-constrained leading finally to the following fluid model for the number of leechers and seeds in the system at time t, denoted by $x(t)$ and $y(t)$, respectively, for their variant of the in-order piece selection policy:

$$\begin{cases} x'(t) = \lambda - \mu(\eta x(t) + y(t)), \\ y'(t) = \mu(\eta x(t) + y(t)) - \gamma y(t). \end{cases} \tag{1}$$

Like Qiu and Srikant, Parvez et al. conclude that $\eta \approx 1$ at least for "most scenarios of interest". However, while the assumption that $\eta \approx 1$ is plausible for file sharing systems, it may not be the case for on-demand-streaming. For example, simulations with realistic P2P VoD systems in [6] have shown that the efficiency of the piece exchange is limited by the used windowing mechanism.

Our fluid model for an on-demand streaming system is also inspired by [5]. The essential difference between our model and that of [2] is in the handling of the playback phase and the modeling of selfishness. In fact, model (1) totally ignores the playback phase and simply assumes that any seed, whether it has played back the media file or not, departs with rate γ. Instead of that, we include the playback phase explicitly in our model. Nor do we restrict ourselves to the upload-constrained case. We also generalize the previous models by adding a finite number of original seeds that stay in the system permanently.

III. System modeling and analysis

Consider a P2P VoD application. We model the system dynamics related to sharing of a single video file assuming a homogeneous peer population and a steady flow arrival scenario, which is a reasonable assumption for YouTube [7].

A. Modeling framework

Let m denote the size of the video file to be shared (in bits). The video is played back at a constant rate w (bits per time unit). New peers arrive at rate λ (arrivals per time unit on average). Each peer is connected to the network over an asymmetric access link (e.g., an ADSL link) with a download capacity of d (bits per time unit) and upload capacity u (bits per time unit). The corresponding download and upload rates are: $c = d/m$ and $\mu = u/m$ (file transfers per time unit).

The life span of a peer consists of two phases, the file *transfer phase* and the *playback phase*, which are partly overlapping. The first part of the transfer phase (called the startup delay) lasts until there are sufficiently pieces in the playback buffer to start the playback phase. After that, the video transfer and playback proceed in parallel until the entire video is transferred. It is desirable that the transfer rate is greater than the playback rate so that the video can be played back without any breaks or delays. In such a case, the playback phase extends beyond the transfer phase.

We assume that the startup delay is negligible. This is possible to achieve with a sufficiently small window size. In addition, we assume that

$$d > w, \tag{2}$$

which is a necessary requirement for a VoD system in order that the video can be played back with good quality, since the transfer rate for a downloading peer is always upper bounded by its own access link rate. If the video is played back without any breaks or delays, the length of the playback phase equals $z = m/w$. It follows from (2) that

$$z > 1/c. \tag{3}$$

During the transfer phase, the peer is called a *leecher*. According to the fundamental P2P principle, leechers help each other. Let $\eta \in [0, 1]$ denote the efficiency of this operation. In the case of windowed BitTorrent, parameters that affect η are the basic efficiency of BitTorrent, the size of the window used, the number of downloading connections, and the number of pieces. A more detailed examination of efficiency in the windowed BitTorrrent case can be found in [6].

An altruistic leecher becomes a *seed* as soon as its own video file download is completed. Let $\zeta \in [0, 1]$ denote the fraction of altruistic peers, which continue to upload to leechers even after their own transfer phase. Non-altruistic peers are assumed to leave the system immediately after the transfer phase, while altruistic peers stay in the system until the end of the playback phase. In addition to the non-permanent seeds, we allow a number of original seeds, say k, that stay in the system permanently.

B. Fluid model

Let $x(t)$ and $y(t)$ denote the number of leechers and seeds, respectively, at time t. The number of leechers, $x(t)$, increases with rate λ. On the other hand, it is decreased with the rate at which file transfers complete. Assuming that there are no bottlenecks in the core network, this completion rate $\phi(t)$ is determined by the minimum of the total download rate and the total upload rate (cf. [5], [2]),

$$\phi(t) = \min\{cx(t), \mu(\eta x(t) + y(t) + k)\}. \tag{4}$$

Thus, we have $x'(t) = \lambda - \phi(t)$.

As mentioned above, an altruistic leecher becomes a seed after the transfer phase. Thus, the number of seeds, $y(t)$, increases with rate $\zeta\phi(t)$. Characterizing the rate at which seeds depart from the system is more difficult and requires some approximations given below.

We base our model on *initially* making the optimistic assumption that the transfer rate is sufficient so that the video can be played back without any breaks or delays. Since we also assumed that the startup delay is negligible and altruistic peers leave the system as soon as they complete the playback phase, this fixes the total time that such a peer spends in the system to be equal to z. Now we make an *approximative* assumption that the system is quasi-stationary at every point in time t. Then by Little's result, the average time that an altruistic peer stays as a seed equals $z - x(t)/\lambda$. Thus, the total departure rate of seeds becomes $y(t)/(z - x(t)/\lambda)$ implying that $y'(t) = \zeta\phi(t) - y(t)/(z - x(t)/\lambda)$.

All in all, our fluid model is as follows:

$$\begin{cases} x'(t) = \lambda - \phi(t), \\ y'(t) = \zeta\phi(t) - \frac{y(t)}{z - x(t)/\lambda}. \end{cases} \tag{5}$$

Note that there are no a priori guarantees that the difference $z - x(t)/\lambda$ stays positive. If this is not the case, the differential equation system behaves in an unstable manner, which can be interpreted as a sign of problems in the playback quality of the P2P VoD application.

C. Steady-state analysis

Model (5) allows us to explicitly solve the equilibrium of the system by setting $x'(t) = y'(t) = 0$ and solving for the corresponding values of x and y. However, due to the min-operation in (4), we need to separately consider whether the system is download or upload constrained at the equilibrium, which results in two different equilibria.

Let us denote by $\bar{x}$ and $\bar{y}$ the equilibrium values of (5). The system is *download-constrained* at the equilibrium if $c\bar{x} \leq \mu(\eta\bar{x} + \bar{y} + k)$, implying that

$$\begin{cases} \bar{x}_{\rm d} &= \frac{\lambda}{c}, \\ \bar{y}_{\rm d} &= \zeta\lambda(z - \frac{1}{c}), \end{cases} \quad (6)$$

where the subscript (d) refers to the download-constrained solution. Now the constraint $c\bar{x}_{\rm d} \leq \mu(\eta\bar{x}_{\rm d} + \bar{y}_{\rm d} + k)$ is equivalent to

$$\frac{1}{\mu} \leq \frac{\eta}{c} + \zeta(z - \frac{1}{c}) + \frac{k}{\lambda}. \quad (7)$$

We see that $\bar{x}_{\rm d} > 0$ for sure and $\bar{y}_{\rm d} > 0$ by (3) so that the solutions are meaningful in the download-constrained case and the equilibrium transfer rate for a leecher exceeds the watching rate.

On the other hand, if $c\bar{x} > \mu(\eta\bar{x} + \bar{y} + k)$, then the system is *upload-constrained* at the equilibrium and we have

$$\begin{cases} \bar{x}_{\rm u} &= \frac{\lambda}{\eta-\zeta}(\frac{1}{\mu} - \zeta z - \frac{k}{\lambda}), \\ \bar{y}_{\rm u} &= \frac{\zeta\lambda}{\eta-\zeta}(-\frac{1}{\mu} + \eta z + \frac{k}{\lambda}), \end{cases} \quad (8)$$

The constraint $c\bar{x}_{\rm u} > \mu(\eta\bar{x}_{\rm u} + \bar{y}_{\rm u} + k)$ is equivalent to

$$\begin{cases} \frac{1}{\mu} > \frac{\eta}{c} + \zeta(z - \frac{1}{c}) + \frac{k}{\lambda}, & \text{if } \eta > \zeta, \\ \frac{1}{\mu} < \frac{\eta}{c} + \zeta(z - \frac{1}{c}) + \frac{k}{\lambda}, & \text{if } \eta < \zeta. \end{cases} \quad (9)$$

Additionally, for the solution to be meaningful, we require that $\bar{x}_{\rm u} > 0$ and $\bar{y}_{\rm u} > 0$. This leads to the following additional constraint:

$$\begin{cases} \zeta < \frac{1}{z}(\frac{1}{\mu} - \frac{k}{\lambda}) < \eta, & \text{if } \eta > \zeta, \\ \eta < \frac{1}{z}(\frac{1}{\mu} - \frac{k}{\lambda}) < \zeta, & \text{if } \eta < \zeta. \end{cases} \quad (10)$$

The equilibrium transfer rate for a leecher exceeds the watching rate if and only if $\bar{y}_{\rm u} > 0$.

Figure 1 (left panel) illustrates the different solution areas in the (η, ζ)-space.

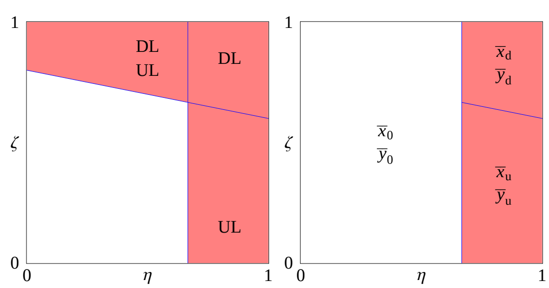

Fig. 1. *Left panel*: Solution areas, where DL [UL] refers to a positive download [upload] constrained solution. The horizontal bordering line satisfies $\frac{1}{\mu} = \frac{\eta}{c} + \zeta(z - \frac{1}{c}) + \frac{k}{\lambda}$ and the vertical bordering line satisfies $\eta = \frac{1}{z}(\frac{1}{\mu} - \frac{k}{\lambda})$. *Right panel*: Steady state synthesis.

D. Local stability analysis

Analyzing the global stability of (5) is difficult due to the nonlinear and switched nature of the model, see [8]. Instead we focus on the local stability of the system. In particular, our approach is to linearize (5) around its equilibrium, and apply the control theory of linear systems.

The local stability properties of the steady-state solutions in the three different areas in Figure 1 (left) can be summarized as follows (see [9] for details). The system has a unique equilibrium (either DL or UL) which is locally stable if $\eta > \frac{1}{z}\left(\frac{1}{\mu} - \frac{k}{\lambda}\right)$. In the area UL/DL, there are two equilibria of which the DL solution is locally stable and the UL solution is unstable.

E. Stochastic model

In addition to a fluid model, we developed a more detailed stochastic model to test the accuracy of the approximative assumptions made for the fluid model. The stochastic model *does not* utilize the quasi-stationarity approximation to estimate the departure rate of altruistic seeds. We refer to [9] for details.

IV. Steady state and scalability

By experimenting with the two models (fluid/stochastic), we obtained that the fluid model closely approximates the mean values of the stochastic model when η is high enough but shows a qualitatively different behavior when η is below a certain threshold. Based on these experiment, we provide below a meaningful synthesis of the system performance as well as an interpretation of the results from the point of view of system scalability.

A. Steady-state synthesis

(i) If

$$\eta > \frac{1}{z}(\frac{1}{\mu} - \frac{k}{\lambda}), \quad (11)$$

the equilibrium transfer rate for a leecher exceeds the playback rate w (so that the the playback quality is sufficient for the VoD application). If additionally (7) is satisfied, then the system is download-constrained, and the numbers of leechers and seeds are well estimated by $\bar{x}_{\rm d}$ and $\bar{y}_{\rm d}$ given in (6). Otherwise, the system is upload-constrained, and the numbers of leechers and seeds are well estimated by $\bar{x}_{\rm u}$ and $\bar{y}_{\rm u}$ given in (8).

(ii) If

$$\eta \leq \frac{1}{z}(\frac{1}{\mu} - \frac{k}{\lambda}), \quad (12)$$

then the equilibrium transfer rate for a leecher stays below the playback rate w resulting in playback quality problems. In this case, in the stochastic model peers have to stay longer in the system than the actual viewing time. In our fluid model this situation can be enforced by evaluating the steady-state solution with the condition $\zeta = 0$. Then, the system is always upload-constrained (see Figure 1) and the numbers of leechers and seeds are well estimated by formula (8) applied with $\zeta = 0$ leading to estimates

$$\begin{cases} \bar{x}_0 &= x_{\rm u}|_{\zeta=0} &= \frac{\lambda}{\eta}(\frac{1}{\mu} - \frac{k}{\lambda}), \\ \bar{y}_0 &= y_{\rm u}|_{\zeta=0} &= 0. \end{cases} \quad (13)$$

These conclusions are summarized in Figure 1 (right panel).

Note that the altruism parameter ζ has no effect on the threshold (11). However, the per-leecher service rate increases with ζ (up to the download rate limitation d) implying that the system offers better viewing quality. In addition, we see from (11) that $\frac{1}{\mu} - \frac{k}{\lambda} < z$ is a necessary condition and $\frac{1}{\mu} - \frac{k}{\lambda} < 0$ is a sufficient condition for good playback quality.

B. Scalability and the effect of permanent seeds and altruism

As explained in the previous section IV-A, the system seems to have a sufficient playback quality if and only if condition (11) is satisfied. We say that the system is *scalable* if (11) is satisfied for all λ. Thus, scalability is equivalent to the requirement that

$$\eta > \frac{1}{z\mu}. \tag{14}$$

A necessary condition for scalability is clearly that $z > 1/\mu$, which is equivalent to $u > w$. On the other hand, we see from (11) that if the system is *not* scalable (so that $1/\mu \geq \eta z$), good playback quality is achieved up to

$$\lambda < \frac{k}{\frac{1}{\mu} - \eta z}. \tag{15}$$

In general, it would seem intuitive that increasing the number of permanent seeds, k, would somehow help the system to perform better. However, as can be seen from (14), to achieve scalability does not depend on k. But in the non-scalable case, by (15), the limit for λ up to which the playback quality is sufficient increases linearly with k.

While we have shown that neither the number of permanent seeds k nor the level of altruism ζ had any effect on the scalability of the system, it does not follow that these are unnecessary parameters. In fact, ζ is a determining factor in whether the system reaches a download or upload constrained steady state. Given an efficiency η satisfying (11), one can always find a needed level of altruism $\zeta < 1$ so that the steady state is download-constrained. In a download-constrained steady state, the service provider could drop some of its permanent seeds without repercussions. In the case of a YouTube-like VoD service provider, the required upload bandwidth per video is a significant portion of costs [10]. Thus even a slight reduction in needed upload bandwidth can be significant.

V. Numerical results and validation

In this section, we validate the accuracy of the fluid model against simulations from the stochastic model of Section III-E and traces from a BitTorrent simulator implementing a windowing algorithm. We first study the system dynamics and then consider the steady-state performance. Similarly as in the steady-state analysis, we focus on the behavior of the system in the (η, ζ) parameter space, cf., Figure 1.

A. Youtube scenario

In the following tests we consider a viewing scenario where the parameters correspond to a typical YouTube setting. The users are viewing a video file consisting of 800 pieces each 32kB in size. The video coding rate is $w = 300$ kbit/s, and thus the viewing time is $z = 682$ s. The upload and download bandwidths of the users are $u = 512$ kbits/s and $d = 1024$ kbit/s (typical asymmetric ADSL subscriber rates). Also, we assume that $k = 1$ (one permanent seed). In the simulations, new users arrive according to a Poisson process with rate $\lambda = 0.2$ peers/s.

B. Validation of system dynamics

Here we give results to validate the accuracy of the fluid model dynamics (5) and the use of the quasi-stationarity approximation. The idea is to compare the dynamic behavior of (5) with the results both from the stochastic model and the BitTorrent simulations.

We first study two examples where the users are first altruistic with $\zeta = 0.9$ and then non-altruistic with $\zeta = 0.3$. The parameters correspond to the situations where condition (11) holds. The results are shown in Figure 2. The figure shows the time evolution of the fluid model (solid smooth line), stochastic model (dashed line) and the BitTorrent simulation (solid jagged line) for both the number leechers $x(t)$ and the number of seeds $y(t)$. The simulations represent averages over 20 sample paths. In the figure, we have numerically determined the value of the efficiency parameter η so that it matches best the equilibrium values of the actual BitTorrent simulations. This resulted in $\eta = 0.85$ for $\zeta = \{0.9, 0.3\}$.

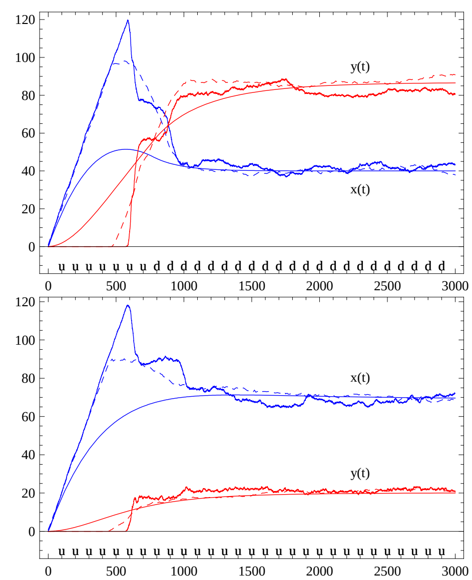

Fig. 2. Comparison of the fluid model (solid smooth line) against the stochastic model (dashed line) and the BitTorrent simulation (solid jagged line) with $\zeta = 0.9$ (upper panel) and $\zeta = 0.3$ (lower panel).

When $\zeta = 0.9$ (see upper panel of Figure 2), one can observe that the system is initially upload-constrained. However, after reaching equilibrium it is download-constrained. As ζ is decreased the number of seeds begins to decrease, which decreases the service capacity available for leechers, and with $\zeta = 0.3$ (see lower panel of Figure 2) the available capacity is now limited by the uploading bandwidth of the system.

With regard to the accuracy of the fluid model we can observe that the equilibrium values of the system are very accurately the same between the stochastic model and the fluid model, and also the BitTorrent simulation. However, the dynamics of the fluid model (solid smooth lines) are somewhat smoother than of the corresponding discrete time simulation of the stochastic model (dashed lines). This can be partly explained by the fact that the model assumes that seeds are continuously generated (recall the quasi-stationary approximation in the use of Little's result), while in the simulated system seeds start accumulating only after the leechers have received enough service to complete the download.

Additional validation results related to the dynamics and the steady state estimates of the fluid model and the simulated systems can be found in [9].

C. Validation of scalability

Next we validate the scalability results in Section IV-B and study the ratio of the number of leechers to the number of seeds when λ is increased. The results are shown in Figure 3, where the solid lines correspond to simulated estimates from the stochastic model and the dashed lines correspond to the analytical estimates of the fluid model. Note that the leecher-to-seed ratio is shown on logarithmic scale to highlight the differences between the cases.

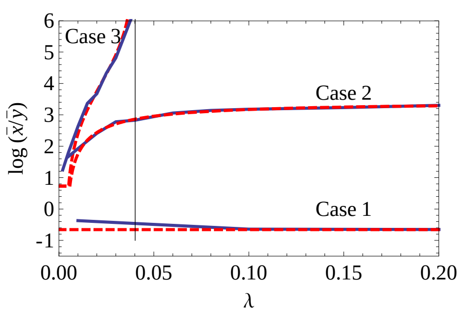

Fig. 3. Validation of the scalability as a function of λ. Cases 1 and 2 are scalable, while Case 3 is not.

By (14), the system is scalable if $\eta > 0.59$ with the present parameter values. This is illustrated in the figure by Cases 1 and 2, where we have used $\eta = 0.65$. In Case 1, $\zeta = 0.8$ and the system stays download-constrained independent of λ and by (6) the leecher-to-seed ratio is also constant. In Case 2, $\zeta = 0.2$ and the steady state evolves so that for small λ the system is download-constrained but then it switches to being upload-constrained. However, the leecher-to-seed ratio tends to a constant, as predicted by (8). On the other hand, in Case 3 we have $\eta = 0.55$ and by (12) the system is stable only up to $\lambda = 0.043$ (illustrated by the vertical line in the figure). As can be seen, the above behavior, as predicted by the analytical results of the fluid model, matches very well the simulated results from the stochastic model. The numerical accuracy of the leecher-to-seed ratio is also well approximated by the results from the fluid model, especially for higher values of λ.

VI. Conclusions

We have derived a fluid model for P2P on-demand streaming of stored content, where the coupled transfer and playback phases are explicitly taken into account. Our model also makes what we feel is a more realistic assumption about peer altruism, where altruistic peers stay on as seeds only as long as they are watching the video. The model consists of a system of differential equations for characterizing the time-evolution of the mean number of leechers and seeds in the system.

Our conclusions are as follows. The efficiency parameter η plays a central role both (i) to guarantee the QoS and (ii) to have a scalable system. Also, a necessary condition for scalability is that both the download rate d and the upload rate u must be greater than the viewing rate w. If the system is not scalable, then acceptable viewing quality is achieved only up to a given λ. This bound can be made higher by, e.g., increasing the number of permanent seeds k. We also showed that a sufficient level of playback phase altruism ζ implies that the system reaches a download-constrained steady state rather than an upload-constrained one, which we believe would result in reducing the bandwidth needed by a VoD service provider, as well as the associated costs.

The potential for future research in this area is rich, including relaxation of some of the modeling assumptions (non-homogeneous upload/dowload rates, finite playback buffer). The actual P2P systems useful for VoD, such as windowing BitTorrent, remain open topics, especially making them as ISP-friendly as possible.

Acknowledgement

This work was supported by TEKES as part of the Future Internet program of TIVIT (Finnish Strategic Centre for Science, Technology and Innovation in the field of ICT).

References

[1] B. Cohen, "Incentives build robustness in bittorrent," in *Proc. of the 1st Workshop on Economics of Peer-to-Peer Systems*, 2003.

[2] K. Parvez, C. Williamson, A. Mahanti, and N. Carlsson, "Analysis of BitTorrent-like protocols for on-demand stored media streaming," in *Proc. of ACM SIGMETRICS*, 2008, pp. 301–312.

[3] T. Bonald, L. Massoulié, F. Mathieu, D. Perino, and A. Twigg, "Epidemic live streaming: Optimal performance trade-offs," in *Proc. of ACM SIGMETRICS*, 2008, pp. 325–336.

[4] R. Kumar, Y. Liu, and K. Ross, "Stochastic fluid theory for p2p streaming systems," in *Proc. of IEEE INFOCOM*, 2007, pp. 919–927.

[5] D. Qiu and R. Srikant, "Modeling and performance analysis of BitTorrent-like peer-to-peer networks," in *Proc. of ACM SIGCOMM*, 2004, pp. 367–378.

[6] P. Savolainen, N. Raatikainen, and S. Tarkoma, "Windowing BitTorrent for video-on-demand: Not all is lost with tit-for-tat," in *Proc. of IEEE GLOBECOM*, 2008.

[7] M. Cha, H. Kwak, P. Rodriguez, Y.-Y. Ahn, and S. Moon, "I tube, you tube, everybody tubes: analyzing the world's largest user generated content video system," in *Proc. of USENIX IMC*, 2007, pp. 1–14.

[8] D. Liberzon and A. Morse, "Basic problems in stability and design of switched systems," *IEEE Control Systems Magazine*, vol. 19, no. 5, pp. 59–70, Oct 1999.

[9] N. Raatikainen, S. Tarkoma, P. Savolainen, S. Aalto, and P. Lassila, "P2p video-on-demand: Steady state and scalability," TKK Helsinki University of Technology, Tech. Rep., 2009, available at: http://www.netlab.hut.fi/tutkimus/shok-fi/publ/p2pvod.pdf.

[10] C. Huang, J. Li, and K. Ross, "Can internet video-on-demand be profitable?" in *Proc. of ACM SIGCOMM*, 2007, pp. 133–144.