IPTV

Tobias Hoßfeld† ††

† University of Würzburg, Insitute of Computer Science
Am Hubland, 97074 Würzburg, Germany
††

-

E-mail: †hossfeld@informatik.uni-wuerzburg.de, ††leibnitz@ist.osaka-u.ac.jp

IPTV

2

IPTV

IPTV,

Performance Evaluation of an IPTV Recording Service

Tobias HOSSFELD† and Kenji LEIBNITZ††

† University of Würzburg, Insitute of Computer Science
Am Hubland, 97074 Würzburg, Germany
†† Osaka University, Graduate School of Information Science and Technology,
1-5 Yamadaoka, Suita, Osaka 565-0871, Japan
E-mail: †hossfeld@informatik.uni-wuerzburg.de, ††leibnitz@ist.osaka-u.ac.jp

Abstract In this paper we study the performance of an IPTV recording service which permits subscribers to record their favorite TV shows over the Internet and download them for offline viewing. The request arrivals for file downloads play an important role and the system reacts differently when the arrivals are time-dependent. We will compare two analytical modeling approaches: a steady state Markov chain analysis with constant arrivals and a fluid model to capture non-stationary flash crowd effects. Furthermore, our approach also takes a realistic distribution of the offered files of an existing IPTV recording service into account, as well as a user model with impatience which leads to aborted download attempts.

Key words IPTV, content delivery networks, user impatience

1. Introduction

Recently, a multitude of new services have emerged which utilize the Internet as a delivery mechanism for multimedia content. With the advent of broadband accesses and fiber-to-the-home technology, a whole new array of services has become available complementing the traditional Internet applications such as web browsing. Currently, the trend shows a great demand for transferring large volume multimedia content, such as VoIP (Voice-over-IP) or IPTV which is the Internet-based delivery of television programs.

Internet users are no longer satisfied with services such as YouTube[1], where they can upload and share short video clips using Adobe Flash technology, but more commercial and non-commercial services emerge that offer entire TV shows for download (e.g. Apple's iTunes Store[2]). New services include Zattoo[3] and Joost[4] which use architectures based on peer-to-peer (P2P) for multi-cast streaming of live TV programs or Video-on-Demand (VOD). A large problem that these services face is that due to license restrictions not all TV programs are available in all countries. The distinction of which content can be accessed from which country is

usually performed based on the user's IP address.

Another mechanism for IPTV content delivery is the so-called *network based video recording*, where the live TV program is recorded at a server and users can download their favorite shows and later view them offline on a PC. An example for such a video recording service is OnlineTVRecorder (OTR [5]) in Germany. The video file sizes offered there can consist from several hundred megabytes up to 1 GB or more depending on the length of the TV show, as well as on the encoding format, e.g. high quality H.264, standard quality DivX, or MPEG-4 for portable devices (iPod, PSP, etc.). Thus, OTR can be regarded basically as an example for a server-based content delivery system with large data files corresponding to long service times. In the case of OTR, the content can be either downloaded directly from the OTR server, from user-created mirror sites, or alternatively via P2P file-sharing network (e.g. BitTorrent).

However, as the OTR server farms are often overloaded, new requests are queued when the number of provided download slots is full. The restriction to a maximum number of simultaneous downloads guarantees a minimal download bandwidth for each user. Additionally, the service offers premium users prioritized access to downloading. The download duration itself depends on the total capacity of the server and the number of users sharing this capacity. On the other hand, users who might encounter slow downloads may abort their downloading attempt if their patience is exceeded.

In this paper, we extend our previous work in [6] and analytically investigate the performance of such an OTR server with different file size distributions and model the user behavior by using an impatience threshold. The paper is organized as follows. After describing the analytical models for both constant and time-dependent arrivals, we provide numerical results and compare their performance in terms of download duration and success ratio. Especially, we deal with the question of how to properly dimension the number of simultaneous downloads at a server in order to optimize the performance of the system and to maximize the user's satisfaction.

2. Analytical Modeling of OTR

Let us consider the following system. User requests arrive at the server with an arrival rate λ. While we will at first consider a fixed arrival rate in order to evaluate a steady state Markov model, we will also consider later a non-stationary arrival rate $\lambda(t)$. This is a more realistic scenario when looking at individual files, since the popularity of a TV show highly depends on the date it was recorded. Once a show becomes outdated, the interest for this file decreases. This phenomenon is usually referred to as *flash crowd arrivals* [7].
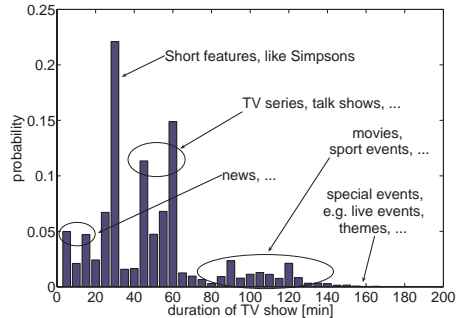

Fig. 1   Probability distribution of TV show durations

However, since a server may offer several different files, the overall arrival rate registered at the server may remain nearly constant. The superposition of time-dependent arrival processes with different starting points can be modeled as stationary Poisson process for a sufficiently large number of offered files per server.

When a request arrives and there are free download slots, the client may proceed with the download. We assume that the server system has a total fixed capacity C which is shared among all simultaneously downloading clients $D(t)$ at time t. The maximum number of users served in parallel is restricted to n. Thus, the time-dependent download rate $\mu(t)$ for a file size f_s is

$$\mu(t) = \frac{1}{f_s} \min \left\{ \frac{C}{\min \{D(t), n\}}, R \right\} \qquad (1)$$

and the maximum download rate is limited by the maximum physical rate R of each client.

2.1   Measurements of Offered OTR Files

As we need the distribution of the file sizes to compute the download rate $\mu(t)$, we investigated the actual file sizes of video files offered at OTR. The measurements which were made in April 2007 consist of 11563 file samples from 19 different TV channels. According to the information provided by OTR [8], standard video files are encoded at a resolution of 512×384 pixels at a video bitrate of about 750 kbps and an audio bitrate of 128 kbps. The measured data contains only standard quality video files and consist of approximately 80% encoded in the DivX format and 20% in Windows Media Video (WMV) format.

Fig. 1 shows the probability distribution of the TV show durations in minutes. The majority of the files (95%) are discretized in units of 5 min. We can distinguish 4 different categories of TV shows. Most files are short features (e.g. animation series) of about 30 min and shorter files may be for instance news shows. Another peak can be found between 45-60 min which is the usual duration of TV dramas or other periodical shows. Movies usually have the duration between 90-120 min and very few larger recordings of special events
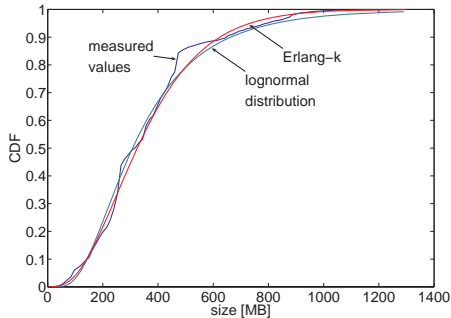
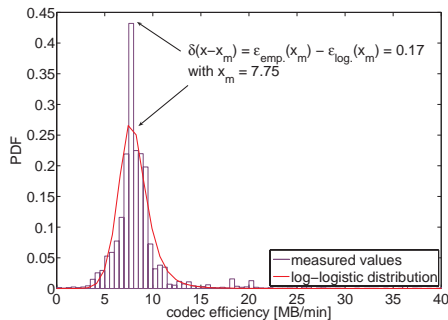Fig. 2   CDF of measured file size

Fig. 3   Codec efficiency as ratio of file size and duration of the TV show

exist, like the broadcasts of live sports events.

However, we are more interested in the file size distribution in order to approximate the download time than the duration of the shows themselves. Fig. 2 shows that the actual file size distribution has a mean of 368.31 MB and standard deviation of 196.82 MB. It can be well fitted by a lognormal or an Erlang-k distribution with $k = 3.34$ phases and an average volume of $B =$107.67 MB per phase, i.e., it is the sum of $\lfloor k \rfloor$ independent identically distributed random variables each having an exponential distribution with mean B and an exponential distribution with mean $(k - \lfloor k \rfloor) B$.

Fig. 3 shows the codec efficiency as ratio of the file size over the duration of the TV show. The PDF has a distinct peak at about 7.75 MB which corresponds to a total encoding bitrate of about 1 Mbps. This matches the description of the provider and is comparable to other standard quality formats, such as VCD or SVCD. The measured values could be well fitted with a log-logistic distribution superimposed with a Dirac function at the peak value x_m to match the measurements. The normalized height added by the Dirac peak is approximately 0.17.

2.2   Discussion of Related Modeling Approaches

In general, with a slight abuse of the Kendall notation for queuing systems, the model as described above can be expressed as M(t)/GI/1^n-PS with user impatience θ, an unlimited waiting queue, and a server capacity which is shared among n users at maximum. Thus, the service rate is influenced by μ and θ and depends on the number of currently served users.

Admission control to the system can be taken into account by restricting the size of the waiting queue. However, in this paper we use the number of download slots n to guarantee the bandwidth per user and only investigate the impact of the user's impatience on the system's performance. While *reneging* is considered with an i.i.d. random variable θ, *balking*, i.e., taking back the download request if the waiting queue is too long, is neglected in this paper. We focus on the effect of wasted capacity due to users' impatience regardless of whether they are being served or not, and the impact of variability of the file size distribution, which is expressed by the service rate. Our findings show that the ratio of successful downloads increases with the variability of service time.

Basically, there are several approaches on how to analytically evaluate such a system depending on the number of available download slots n. If $n \leq \lfloor \frac{C}{R} \rfloor$, the user's access bandwidth limits the download rate. This effectively results in a M(t)/GI/n-FCFS system with independent service rates, since θ is an i.i.d. r.v. and μ is constant. An analytical evaluation is provided in [9]. For $n > \lfloor \frac{C}{R} \rfloor$, the download rate and therefore the service rates depend on the current state of the system. On the other hand, if the downlink of a user is not the limiting factor, i.e., a user can always utilize the offered bandwidth of the server $(C < R)$, the system approaches a real processor sharing system with increasing n, which is investigated in [10], [11]. In the past, a lot of research has been dedicated to the analysis of queuing models with impatience. Barrer [12] was among the first to analyze an M/M/1 system with deterministic impatience thresholds. In the following, more sophisticated FCFS models with Markovian arrival and service processes were investigated in [13]   [15].

In this paper only very simple models are considered which are easily analytically tractable. As we will see later, the general service time has a great impact on the performance, however, it is well known that for such systems only approximative evaluations can be performed for metrics of interest [13]. As we consider time-dependent flash crowds arrivals a transient analysis as described later in Section 2.4 is required.

2.3   Steady State Analysis with Markov Model

We now consider a steady state analysis for evaluating the performance of the server system with aborted downloads due to impatience. Maximally n users are served at a time and they share the server's total capacity C with the processor sharing (PS) discipline. We assume homogeneous users with equal access bandwidths R and generally independent
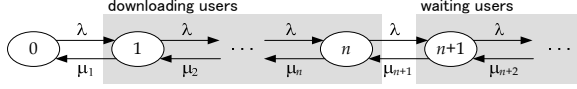
Fig. 4 State space of Markov chain model

patience time θ. In our model, θ is the time threshold after which a user aborts his download attempt if the download time, i.e., the sum of the waiting and the service time, takes longer than that. However, this GI assumption is not an accurate model for the actual user's behavior. In reality, a user will have a state-dependent patience, since he is more willing to wait if the file is nearly completed. However, in order to make the model analytically tractable, we consider an exponentially distributed θ. Some research exists on impatience in queuing systems, e.g. [16]. The model will be denoted $M/M/1^n$-PS with a slight abuse of the Kendall notation for queuing systems. Thus, we have a homogeneous Poisson arrival process, exponential service time, a single server unit which services up to n clients and operates with the processor sharing regime. The queue length for waiting users is assumed to be infinite.

The infinite state space for the $M/M/1^n$-PS model is shown in Fig. 4. Each state is identified by the number of clients currently in the system. The model itself is a simple birth-death process where only transitions between neighboring states are possible. The service rates μ_i are dependent on state $i = 1, 2, \ldots$ and are expressed as follows.

$$\mu_i = \frac{i}{\theta} + \frac{\min\{i, n\}}{f_s} \min\left\{\frac{C}{\min\{i, n\}}, R\right\} \qquad (2)$$

2.4 Time-Dynamic Fluid Model

The Markov model described in the previous section only allows to investigate the steady state behavior. In order to also consider the flash crowd arrivals mentioned above, we use a fluid analysis technique. The state space of transitions is shown in Fig. 5 and the differential equation system is given in (6).
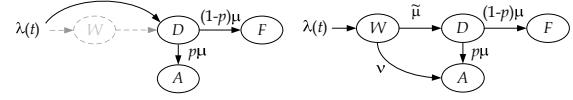
$$\dot{W} = \begin{cases} 0 & \text{if } D < n \\ \lambda - D\,\mu - \nu W & \text{otherwise} \end{cases} \qquad (3)$$

$$\dot{D} = \begin{cases} \lambda - D\,\mu & \text{if } D < n \\ 0 & \text{otherwise} \end{cases} \qquad (4)$$
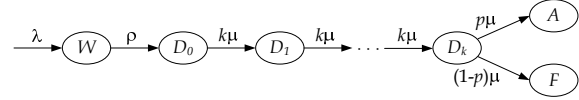
$$\dot{A} = D\,p\,\mu + \nu W \qquad (5)$$

$$\dot{F} = D\,(1 - p)\,\mu \qquad (6)$$

Arrivals enter the waiting population W with rate λ or directly the downloading population D, if the number of slots n is not full, see Fig. 5(a). If the slots are full (Fig. 5(b)), waiting users simply proceed to the downloading state with

(a) $D < n$

(b) $D = n$

(c) State D split up into $D_0, \ldots, D_k$ with $k \to \infty$

Fig. 5 Fluid model state transitions

rate $\tilde{\mu} = \mu D$, which does not depend on W. After entering state D, the client remains in this state until he either fully downloads the file and enters the finished state F or he aborts the download when the download duration exceeds his patience threshold θ. The latter is expressed by entering abort state A. In both cases the transitions are performed at rate μ multiplied with a probability p (when the download fails) or $1 - p$ in the case of success. The probability p can be interpreted in the following way. An abort occurs when the patience of the downloading user is exceeded either during downloading or waiting. The patience in this model is characterized by the exponential random variable θ with rate $\nu = 1/E[\theta]$ and the downloading time is exponentially distributed as well with rate $\psi = C(t)/E[f_s]$. The variable $C(t)$ denotes the time-dependent capacity per user, i.e., $C(t) = C/D(t)$ and $E[f_s]$ is the mean file size. Thus, the probability that the patience is exceeded at time t can be expressed as

$$p(t) = \frac{\nu}{\nu + \psi} = \frac{D(t)\,E[f_s]}{D(t)\,E[f_s] + C\,E[\theta]}. \qquad (7)$$

Note that in the case of a single downloading state D, exponential file sizes f_s and thus exponentially distributed rates μ are assumed. If we consider Erlang-k distributed file sizes as obtained in our measurements, the state D must be expanded to several intermediate states $D_0, D_1, \ldots, D_k$. For $k \to \infty$ this approaches deterministic values, see Fig. 5(c).

3. Numerical Results

We now provide some numerical results obtained by the analytical models and simulations. A good match for both the steady state analysis with the Markov model and the time-dynamic evaluation with fluid model can be found in Fig. 6 and Fig. 7. Fig. 6 shows the cumulative distribution function (CDF) of the steady state population sizes for different values of n. From this we can derive the download time of a user and the success ratio to obtain a file.
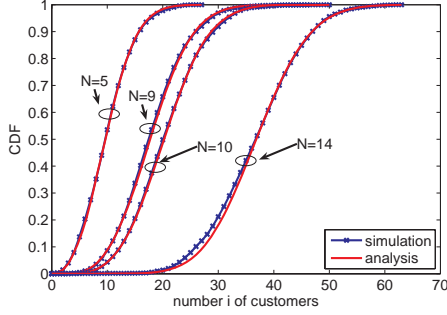
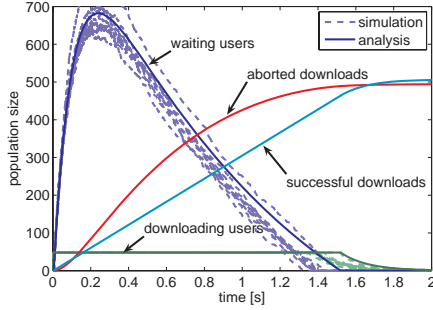Fig. 6 Steady state population distributions

Fig. 7 Population changes with fluid model

Next, we consider the flash crowd scenario with the fluid model. We assume an exponentially decreasing arrival rate

$$\lambda(t) = \beta\, e^{-\alpha t}$$

with $\beta = 1$ and $\alpha = 10^{-3}$. Thus, the total number of arriving users in the system is limited to

$$\lim_{t \to \infty} \lambda(t)\, dt = \frac{\beta}{\alpha} = 1000.$$

Fig. 7 shows the time-dynamic evolution of the population size in the flash crowd scenario. We compare the population sizes from several simulation runs with the numerical solution of the differential equation system (6) which is obtained by the ODE solver of Matlab using the Runge-Kutta method [17].

In the following we look at the simulated behavior of the system when there are constant and flash crowd arrivals. In order to compare systems with both types of arrivals, we matched the arrival rate for the constant case to get the same number of arrivals as in the case of flash crowds, where the arrival rate $\lambda(t)$ exponentially decreases as $\dot{\lambda} = -\alpha\lambda$ and $\lambda(0) = \beta$. Here, we used the parameters $\beta = 1$, $\alpha = 10^{-4}$, as well as the server capacity $C = 100\,\mathrm{Mbps}$, user bandwidth $R = 2\,\mathrm{Mbps}$ and patience threshold $\theta = 200\,\mathrm{min}$, and the file size distribution was taken from measurement values. The number of total download slots $n = \lfloor C/R \rfloor$ represents the optimally dimensioned case as will be shown later.

Fig. 8 and Fig. 9 depict the two measures of interest to
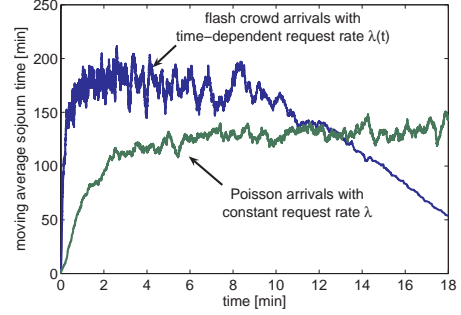

Fig. 8 Sojourn times for flash crowd / Poisson arrivals
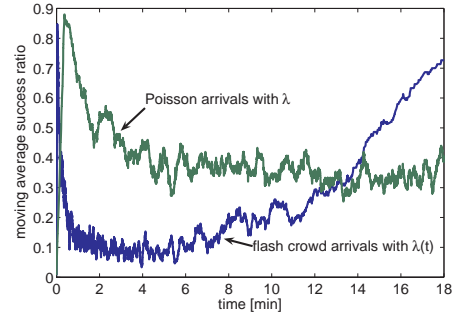

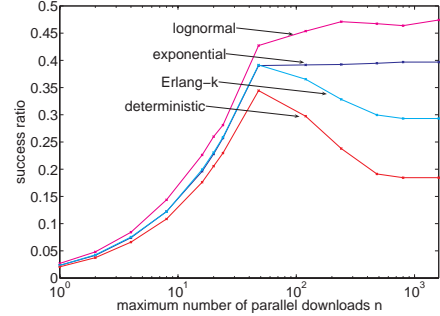
Fig. 9 Moving average of success ratio

Fig. 10 Success ratio for different file size distributions

us, the download time and success ratio. We take a look at the temporal evolution using a moving average with a window size of 100. Both figures show that there is a significant difference when constant or time-dependent arrivals are considered. With constant arrival rate, after an initial transient phase, both the download duration and the success ratio become rather constant, whereas with flash crowds, there is a higher variation of both values as the arrivals rapidly decrease over time from which later arrivals benefit. The figures show that it is very important to consider if the arrivals are time-dependent or not, as they yield quite different results.

The next investigation aims at the optimal dimensioning of the number of download slots n for different file size distributions. We focus on the flash crowd scenario with the same parameters as above except for $\alpha = 10^{-3}$. While Fig. 10 shows the success ratio when the file size is distributed ei-
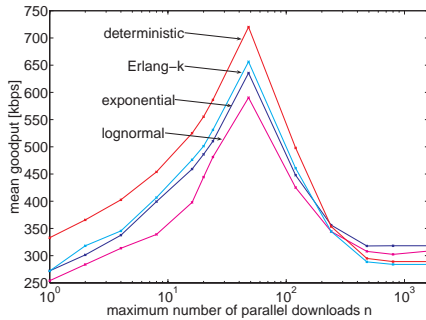
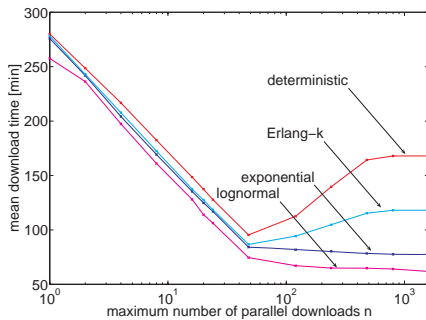Fig. 11 Goodput depending on file size variation

Fig. 12 Download time depending on file size variation

ther deterministic, exponential, Erlang, or lognormal, Fig. 11 depicts the average goodput in kbps depending on the maximum number n of simultaneously served users.

Both figures illustrate the influence of the skewness on the system behavior. It is remarkable that for deterministic and Erlang-distributed file sizes a maximum success ratio exists, whereas for exponential and lognormal the success ratio remains nearly constant when $n > \lfloor C/R \rfloor$. However, this is caused by the fact that with a higher skewness, smaller files are downloaded more often. In all four cases the goodput is highest at $n = \lfloor C/R \rfloor$, as can be seen from Fig. 11. The goodput is defined as the ratio of the file size and the download time for successful downloads. For larger n the system capacity is wasted due to longer download times caused by the processor sharing discipline and the aborts due to the user's impatience.

Finally, the download time reflects also the behavior seen for the success ratio, see Fig. 12. Since mostly the download of large files are aborted when the downlink bottleneck is exceeded and the lognormal distribution has a larger weight on small values, it achieves the shortest download time.

4.　Conclusion and Outlook

In this paper we modeled and investigated the performance of an online TV recording service for distributing large-volume video files. The user behavior was characterized with an impatience threshold after which the client aborts the download. We derived stationary and a transient fluid flow model for analysis and compared their performance in terms of the mean download duration and success ratio.

In the future, we wish to perform a more detailed comparison with content distribution methods using peer-to-peer networks [18]. By utilizing the benefits of distributed serving nodes as in P2P with optimal strategies for caching contents, our goal is to design better content delivery networks with a higher reliability and scalability.

Acknowledgment

References

[1] http://www.youtube.com/.

[2] http://www.apple.com/itunes/.

[3] http://www.zattoo.com/.

[4] http://www.joost.com/.

[5] http://www.onlinetvrecorder.com/.

[6] T. Hoßfeld, K. Leibnitz, and M.-A. Remiche, "Modeling of an online TV recording service," *ACM SIGMETRICS Performance Evaluation Review, Special Issue on the MAMA 2007 Workshop*, vol. 35, pp. 15–17, Sept. 2007.

[7] T. Hoßfeld, K. Leibnitz, R. Pries, K. Tutschku, P. Tran-Gia, and K. Pawlikowski, "Information diffusion in eDonkey-like P2P networks," in *Proc. ATNAC 2004*, (Bondi Beach, Australia), 2004.

[8] http://wiki.onlinetvrecorder.com/index.php/Sender.

[9] N. Gans, G. Koole, and A. Mandelbaum, "Commissioned paper: Telephone call centers: Tutorial, review, and research prospects," *Manufacturing & Service Operations Management*, vol. 5, no. 2, pp. 79–141, 2003.

[10] E. Coffman, A. Puhalskii, M. Reiman, and P. Wright, "Processor-shared buffers with reneging," *Perform. Eval.*, vol. 19, no. 1, pp. 25–46, 1994.

[11] H. C. Gromoll, P. Robert, B. Zwart, and R. Bakker, "The impact of reneging in processor sharing queues," in *Proc. of SIGMETRICS '06/Performance '06*, (New York, NY, USA), pp. 87–96, ACM Press, 2006.

[12] D. Y. Barrer, "Queueing with impatient customers and ordered service," *Oper. Res.*, no. 5, 1957.

[13] B. Gnedenko and D. König, *Handbuch der Bedienungstheorie II*. Berlin: Akademie-Verlag, 1984.

[14] F. Baccelli and G. Hebuterne, "On queues with impatient customers," Tech. Rep. 94, INRIA - Centre de Rocquencourt, 1981.

[15] A. Brandt and M. Brandt, "On the M(n)/M(n)/s queue with impatient calls," *Perform. Eval.*, vol. 35, no. 1-2, 1999.

[16] S.-C. Yang and G. de Veciana, "Bandwidth sharing: The role of user impatience," in *Proc. IEEE GLOBECOM*, (San Antonio, TX), pp. 2258–2262, Dec. 2001.

[17] J. R. Dormand and P. J. Prince, "A family of embedded Runge-Kutta formulae," *J. Comp. Appl. Math.*, vol. 6, pp. 19–26, 1980.

[18] K. Leibnitz, T. Hoßfeld, N. Wakamiya, and M. Murata, "Peer-to-peer vs. client/server: Reliability and efficiency of a content distribution service," in *20th International Teletraffic Congress (ITC-20)*, (Ottawa, Canada), June 2007.