

Transfer Rate Models for Gnutella Signaling Traffic

Dragos Ilie

*Dept. of Telecommunication Systems
School of Engineering
Blekinge Institute of Technology
371 79 Karlskrona, Sweden
dragos.ilie@bth.se*

David Erman

*Dept. of Telecommunication Systems
School of Engineering
Blekinge Institute of Technology
371 79 Karlskrona, Sweden
david.erman@bth.se*

Adrian Popescu

*Dept. of Telecommunication Systems
School of Engineering
Blekinge Institute of Technology
371 79 Karlskrona, Sweden
adrian.popescu@bth.se*

Abstract

This paper reports on transfer rate models for the Gnutella signaling protocol. New results on message-level and IP-level rates are presented. The models are based on traffic captured at the Blekinge Institute of Technology (BTH) campus in Sweden and offer several levels of granularity: message type, application layer and network layer. The aim is to obtain parsimonious models suitable for analysis and simulation of P2P workload.

1. Introduction

The last few years has seen a large increase of Peer-to-Peer (P2P) research in various forms. This activity has been motivated by the dramatic increase of P2P traffic in the Internet. According to recent research, P2P traffic represents 60% of all traffic in the Internet [1].

Though there is a large amount of research on P2P, little of it presents detailed models of the actual signaling messages in the networks. There are several publications dealing with system-level and IP-level characteristics [2, 10, 17], but to the best of our knowledge, this paper represents the first publication reporting on message-specific models for Gnutella.

This paper presents detailed Gnutella traffic characteristics. The aim has been to keep the model complexity at a level suitable for analysis and simulation. Consequently, the mixture models used in this paper employ no more than two

components in order to prevent an increase in the number of parameters.

The rest of the paper is as follows. Section 2 briefly discusses the Gnutella protocol messages, followed by a description of the modeling methodology employed in obtaining the models presented in sections 5–7. Section 4 gives a short overview of the measurements performed for this work. Section 7 concludes the paper.

2. The Gnutella Protocol

Gnutella is a decentralized P2P system [5, 13]. Participants can share any type of resources, although the currently available specification covers only file resources. The protocol is easily extensible, which has led to the adoption of a variety of proprietary and non-proprietary extensions (e.g., ultrapeers [19] and the Query Routing Protocol [16]).

Gnutella signaling activities are concerned with discovering the network topology (with the help of PING and PONG messages) and locating resources of interest (with the help of QUERY and QUERY_HIT messages).

All nodes in the Gnutella overlay regularly send PING messages to their neighbours, which forward the messages to their neighbours, essentially flooding the overlay. The extent of the flooding is limited by a TTL variable in the message header. Hosts that receive a PING message reply with one or more PONG messages which are sent in the reverse direction on the route taken by the corresponding PING message, eventually reaching the initial sender. Each PONG message contains the IP address and port where a

peer, different from the one sending the PONG message, can be reached.

Peers search for resources using QUERY messages, which flood the network in a manner similar to how PING messages do. Hosts that share resources requested by the query respond with one or more QUERY_HIT messages, depending on how many resources match and how large the corresponding information to be transmitted becomes. QUERY_HIT messages are routed in a similar way as PONG messages.

3. Modeling Methodology

This work focuses on analytical models for Gnutella signaling traffic. Our goal is to use the models developed here in future work that requires simulation of realistic node workloads caused by Gnutella signaling traffic. Due to the complexity of the Gnutella protocol single Probability Density Functions (PDFs) cannot accurately describe all random variables of interest; at least two probability distributions must be mixed. Typically, one distribution accounts for the main body of the empirical distribution and a different one describes the behaviour in the tail. The crux of the problem becomes to find a way to combine the two distributions in a meaningful way.

The method used here is based on finite mixture distributions as described in [20]. A mixture distribution $F_X(x)$ with two components has the following density and distribution function, respectively:

$$f_X(x) = \pi g_1(x) + (1 - \pi)g_2(x) \quad (1)$$

$$F_X(x) = \pi G_1(x) + (1 - \pi)G_2(x) \quad (2)$$

where $g_1(x)$ and $g_2(x)$ are PDFs, $G_1(x)$ and $G_2(x)$ are Cumulative Distribution Functions (CDFs) and π is a constant called *mixing weight*. The mixing weight π is selected such that $0 < \pi < 1$ and decides how much each component is allowed to influence the distribution $F_X(x)$.

Sometime a mixture distribution (in particular one with only two components) cannot describe the data accurately enough. This may be improved by increasing the number of components in the mixture distribution at the expense of an increase in the number of parameters. For this work a different approach was used. Typically, the major discrepancies between the estimated distribution and the true one appear either in the body or in the tail. If, for example, the discrepancies appear in the tail, one can attempt to improve the model accuracy by adjusting the values of the distribution parameters. However, this is likely to decrease the accuracy of the model in the body. Similarly, attempts to increase the accuracy in the body may lead to (higher) discrepancies in the tail. Thus, a trade-off is required: accuracy in the body versus accuracy in the tail. A decision must be taken on

which part of the distribution (body or tail) is more important to model accurately. In the case of transfer rates the tail of the distribution models high rates of traffic (bursts) that occur rarely. On the other hand, the body of the distribution models the “average” size of transfer rates. Because of our interest in simulating workload, the distribution tail is deemed more important than the body. This is because bursts of traffic have a greater impact on the workload than the typical traffic.

In order to keep the mathematical formulas brief we use the following conventions. CDFs are denoted by capital letters and PDFs by lower case letters, as in Table 1. The parameters are as follows: μ and σ are the mean and standard deviation, while α , β and κ are the shape, scale and location parameters.

Table 1. Model notation

Normal	$n_X(x; \mu, \sigma)$	$N_X(x; \mu, \sigma)$
Log-normal	$ln_X(x; \alpha, \beta)$	$LN_X(x; \alpha, \beta)$
Pareto	$pa_X(x; \alpha, \kappa, \beta)$	$PA_X(x; \alpha, \kappa, \beta)$

3.1. Distribution Fitting

It is assumed that a random sample, X_1, X_2, \dots, X_n of the data has been collected. The first step is to construct an empirical distribution, denoted by F_n , from the random sample. Following [6], this can be done by first obtaining the order statistics $X_{(1)} < X_{(2)} < \dots < X_{(n)}$ and then applying the formula:

$$F_n(x) = \begin{cases} 0 & x < X_{(1)} \\ \frac{i}{n} & X_{(i)} < x < X_{(i+1)} \\ 1 & X_{(n)} \leq x \end{cases} \quad (3)$$

The empirical distribution $F_n(x)$ is plotted and a distribution family $G_1(\cdot)$ (e.g., the Log-normal family) is identified that matches the body of the data, preferably the tail as well. The parameters of the distribution are estimated, yielding a specific distribution function $\hat{G}_1(x; \hat{\theta}_1)$, where $\hat{\theta}_1$ represents the estimated distribution parameters. The $\hat{G}_1(x; \hat{\theta}_1)$ is visually compared to the true distribution to assess the fit in the tail. If the fit is good, then a goodness-of-fit measure $E\%$ is computed as explained in Section 3.2. Otherwise, it is necessary to find the *cutoff point* x_c and the corresponding *cutoff quantile*, q_c , where \hat{G}_1 diverges from the true distribution. The probability mass between q_c and 1 is used to identify the distribution family, $\hat{G}_2(\cdot)$ that matches the tail. The parameters of the new distribution must be estimated as well, yielding $\hat{\theta}_2$. Then, a finite mixture distribution is assembled:

$$\hat{F}(x; \hat{\theta}) = \pi \hat{G}_1(x; \hat{\theta}_1) + (1 - \pi) \hat{G}_2(x; \hat{\theta}_2) \quad (4)$$

Since the single distributions, \hat{G}_1 and \hat{G}_2 , are now combined in a finite mixture, the parameters $\hat{\theta}_1$ and $\hat{\theta}_2$ must be recomputed. Their original values may be used as a starting point. An optimal value for π must be computed as well. The parameter set $\hat{\theta}$ in F is the set containing the parameters for both distributions and π , i.e., $\hat{\theta} = \{\hat{\theta}_1, \hat{\theta}_2, \pi\}$.

3.2. Fitness Assessment

The quality of a fitted model should not be estimated solely by visual means. Therefore more formal methods should be employed. Currently, various goodness-of-fit tests exist, e.g., the χ^2 , Kolmogorov-Smirnov and Anderson-Darling tests. The common denominator for these is the test of the null hypothesis:

$$H_0 : \text{The } X_i\text{'s have distribution } \hat{F}(x, \hat{\theta}) \quad (5)$$

Unfortunately, all these tests suffer from the problem of erroneously rejecting the null hypothesis when the number of samples is large [4, 14]. In order to circumvent these problems a goodness-of-fit measure, referred to as $E\%$, based on the Probability Integral Transform (PIT) [14] and [3], is used. The method was introduced in [7] and works as follows:

Algorithm 1 Calculate Error Percentage

Fit a distribution \hat{F} to the data,
 Transform the data with PIT: $\hat{U} = \hat{F}(X)$

$$E\% \leftarrow 100 \frac{\sum_{t=1}^n |U_t - \hat{U}_t|}{n E_{max}}$$

return $E\%$

If the distribution \hat{F} is a perfect fit, then the PIT transforms the random sample to a Uniform distribution. That is, given an arbitrary stochastic variable R with CDF $F_X(x)$, then

$$F_X(R) = P[X \leq R] = Y \quad (6)$$

where $Y \sim U[0, 1]$. However, since perfect fittings rarely occur in reality, the transformed distribution, \hat{U} , only approximates the Uniform distribution. The discrepancies between U and \hat{U} are computed and their average is normalized to the highest theoretical error, $E_{max} = \int_0^1 \sup \{U(x), 1 - U(x)\} dx = \frac{3}{4}$.

4. Measurements

The P2P research group at BTH has implemented a dedicated measurement infrastructure to measure application layer messages at the link layer [11].

The measurement network infrastructure consists of two measurement nodes; one for Gnutella and one for BitTorrent (Fig. 1). Both nodes are connected to a 10/100 Mbit switch, which in turn is connected to the Internet via the BTH access router.

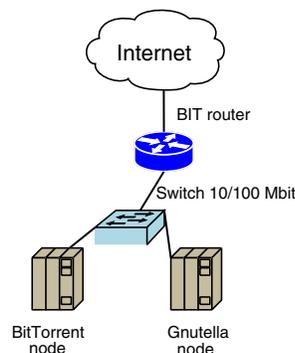


Figure 1. Measurement infrastructure

The infrastructure used for the measurements reported here is based on the one reported in [11] with the addition of several enhancements, most notably ability to decode HSEP messages [18] and improved performance in application flow reassembly.

The measurements for the models reported in the paper were performed during July 1–12, 2005 at the BTH P2P research lab. The observed traffic was flowing across one of the peer nodes at BTH running as a Gnutella ultrapeer node. Measurements were collected using `tcpdump` [12], and no other network applications were running during the measurement.

The trace was collected using the approach described in [8]. The total amount of PCAP data is approximately 33 GB. The PCAP data generated approximately 45 GB log files. The recorded traffic contains 234 million IP datagrams. The log files show 604 thousand Gnutella sessions that were used to exchange 267 million Gnutella messages. A total of 423 thousand sessions (70%) were unable to perform a successful Gnutella handshake. The main reasons for the unsuccessful handshakes were connection queue full¹ and refusal to accept uncompressed connections². The remaining sessions consisted of 181 thousand sessions where both peers used compression, 22 where one of the peers used compression and 10 uncompressed sessions.

¹Code 409: “Vendor would exceed 60% of our slots”.

²Gnutella compresses the application data sent on a connection using the *deflate* algorithm [9, 15].

5. Message Type Transfer Rates

In this section we present the transfer byte rate modeling results for four Gnutella messages types: PING, PONG, QUERY and QUERY_HIT. Models for both incoming and outgoing message rates are reported. All Complementary Cumulative Distribution Function (CCDF) plots use $\log_{10} x$ - and y -axis.

5.1. PING-PONG Messages

Table 2 provides summary statistics of the of PING and PONG byte rates.

Table 2. PING-PONG Byte rate statistics

Type	Dir	Max	Min	Mean	Median	Stddev
PING	IN	1665	0	92	92	50
PING	OUT	503	0	19	0	45
PONG	IN	17043	0	1213	1173	541
PONG	OUT	26050	0	2235	2162	1179

In Table 3, the modeling results for PING and PONG byte rates are presented. While both PONG directions and the incoming PING rates are modeled by a single Log-normal distribution, the outgoing PING message distribution is more complicated. Our investigations show that it cannot be described by neither a single distribution nor a finite mixture. Most of the probability mass, 62.9%, is accounted for by zero byte samples. This indicates that there are periods of one second or longer when no PING messages are sent. An additional 24.3% of the probability mass is centered on 23 bytes, which is the size of a single PING message without extensions. Samples with values ranging from 40 to 90 bytes use another 7.6% of the probability mass. We have therefore decided to model only the upper 5% of the tail (the 95% quantile) with a Normal distribution, resulting in a 4.6% error.

Table 3. PING-PONG modeling results

Message	Model	$E_{\%}$
In PING	$LN_X(x; 4.5, 0.43)$	5.8%
Out PING	$N_X(x; 192, 71.5)$	4.6%
In PONG	$LN_X(x; 7.07, 0.37)$	3.9%
Out PONG	$LN_X(x; 7.64, 0.5)$	3.8%

The errors reported in Table 3 may be considered to be fairly large according to the criteria given in [7]. However, recall that for the purposes of this paper, the distribution tail is of higher importance than the body. Inspection of the CCDFs in Fig. 2 clearly indicates that the majority of the errors are in the distribution body, since the errors in the tail are of little significance.

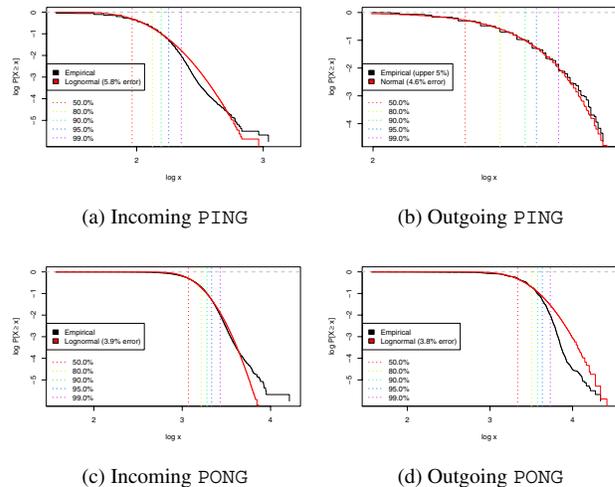


Figure 2. PING/PONG message byte rates

5.2. QUERY-QUERY_HIT Messages

The most salient feature of the summary statistics reported in Table 4 is the similarity of the statistics for incoming and outgoing QUERY messages.

Table 4. QUERY-QUERY_HIT byte rate statistics

Type	Dir	Max	Min	Mean	Median	Stddev
QUERY	IN	24101	0	4441	4317	1426
QUERY	OUT	46424	0	5088	4702	2511
QUERY_HIT	IN	1736791	0	4868	1912	23917
QUERY_HIT	OUT	360235	0	3355	1837	5229

The appearance of a Pareto distribution in the incoming QUERY_HIT rates shown in Table 5 is likely due to some misbehaving peers sending larger responses than is allowed according to the specification. The maximum allowed Gnutella message size is 4 kB [13], while our measurement peer has received several messages larger than 30 kB. The measurement ultrapeer drops all such large messages from the connected nodes, which explains the less heavy tail for the outgoing rates.

Table 5. QUERY-QUERY_HIT modeling results

Message	Model	$E_{\%}$
In QUERY	$LN_X(x; 8.37, 0.29)$	0.4%
Out QUERY	$LN_X(x; 8.46, 0.46)$	1.7%
In QUERY_HIT	$0.95LN_X(x; 7.46, 1.17) + 0.05PA_X(x; 0.91, 0, 7377)$	2.0%
Out QUERY_HIT	$LN_X(x; 7.5, 1.1)$	3.0%

The slope of the upper percentiles in Fig. 3(c) is a clear

visual indicator of the Paretian influence, as compared to the more sub-exponential (Log-normal) behavior observed in the rest of Fig. 3.

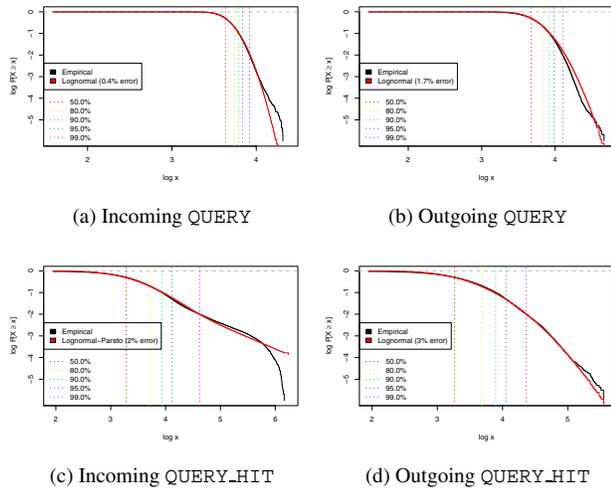


Figure 3. QUERY/QUERY_HIT message byte rates

6. Gnutella Byte Rate Distribution

This section presents models for the aggregate Gnutella byte rates, i.e., the total byte rates for all incoming and outgoing messages. While the average and median outgoing rates are slightly higher than the corresponding incoming rates, the variance and maximum outgoing rate indicate a less bursty traffic pattern.

Table 6. Gnutella byte rate statistics

Dir	Max	Min	Mean	Median	Stddev
IN	1745341	0	12883	10113	24287
OUT	370825	0	13338	12062	7624

The reported byte rate models in Table 7 show more heavy-tail behavior. This is indicated by the appearance of a Pareto distribution in the models for both incoming and outgoing rates as well as by the tails of the CCDFs shown in Fig. 4.

Table 7. Gnutella byte rate modeling results

Message	Model	$E\%$
In	$0.76 \text{LN}_X(x; 9.26, 0.37) + 0.23 \text{PA}_X(x; 1.06, 0, 4003)$	5.2%
Out	$0.81 \text{LN}_X(x; 9.43, 0.39) + 0.19 \text{PA}_X(x; 0.63, 0, 3704)$	5.3%

The aggregate Gnutella models in Table 7 are less well-behaved, in the sense that they are more difficult to model

than the separate message models. This is further emphasized by the CCDFs shown in Fig. 4.

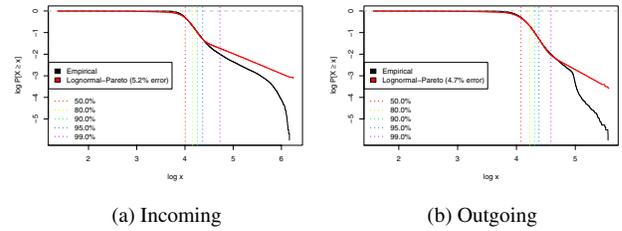


Figure 4. Gnutella byte rates

7. Network Layer Transfer Rates

Table 8 provides the summary statistics for the IP byte rates. It is interesting to note that the mean and median IP byte rates are very similar to the corresponding Gnutella byte rates.

Table 8. IP Byte rate statistics

Dir	Max	Min	Mean	Median	Stddev
IN	249522	0	11536	10961	4075
OUT	176986	0	12668	12037	5722

The heavy-tail behavior present in the Gnutella rates are still present in the IP byte rates reported in Table 9, albeit less pronounced. We conjecture that this is due to the compression of the Gnutella application messages (either single or in bulk) resulting in less data being sent to the IP layer. Accordingly, the maximum byte rates and variances are smaller. This is particularly evident in the incoming IP and Gnutella byte rates, where the maximum and variance differ by an order of magnitude.

Table 9. IP byte rate modeling results

Message	Model	$E\%$
In	$0.89 \text{LN}_X(x; 9.33, 0.26) + 0.11 \text{PA}_X(x; 0.32, 0, 2774)$	3.5%
In (upper 30%)	$\text{PA}_X(x; 0.27, 12812, 2180)$	2.2%
Out	$0.86 \text{LN}_X(x; 9.45, 0.27) + 0.14 \text{PA}_X(x; 0.87, 0, 1662)$	2.3%

Upon inspection of the CCDF for the incoming IP byte rate, the model appeared not to be very accurate in the upper 10%. However, it was found that a single Pareto distribution can accurately model 30% of the tail with no more than 2.2% error. The fitness quality is clearly visible in Fig. 5(a)

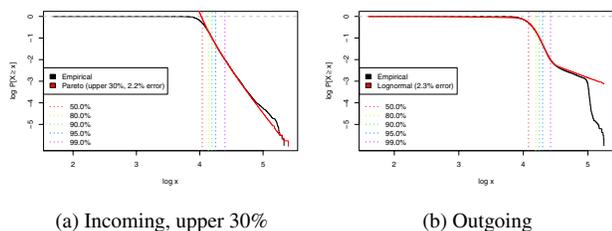


Figure 5. IP byte rates

8. Conclusions

The paper has reported detailed models for several Gnutella protocol message rates, including PING, PONG, QUERY and QUERY_HIT, in both up- and downstream directions, for traffic collected at BTH. Additionally, models for the aggregate byte rates at the IP layer have been presented. Results show that the Log-normal distributions make up a large part of the distributions bodies, and that the Pareto distribution appears in several distribution tails. Most notably, the downstream IP byte rates show a very distinct power-law behavior for the upper 50% of the probability mass.

References

- [1] C. A. Parker. P2P in 2005. <http://www.cachelogic.com/research/slide9.php>, May 2005.
- [2] N. B. Azzouna and F. Guillemin. Experimental analysis of the impact of peer-to-peer applications on traffic in commercial IP networks. *European Transactions on Telecommunications: Special Issue on P2P Networking and P2P Services*, 2004.
- [3] J. Banks, J. S. Carson II, B. L. Nelson, and D. M. Nicol. *Discrete-Event System Simulation*. Prentice-Hall, Inc., Upper Saddle River, New Jersey 07548, U.S.A, 3rd edition, 2001. ISBN: 0-13-088702-1.
- [4] J. Beran. *Statistics for Long-Memory Processes*. Chapman & Hall, 1994. ISBN: 0-412-04901-5.
- [5] Clip2. *The Annotated Gnutella Protocol Specification v0.4*. The Gnutella Developer Forum (GDF), 1.8th edition, July 2003. http://groups.yahoo.com/group/the_gdf/files/Development/.
- [6] R. B. D'Agostino and M. A. Stephens. *Goodness-of-Fit Techniques*. Marcel Dekker, Inc., 1986. ISBN: 0-8247-7487-6.
- [7] D. Erman, D. Ilie, and A. Popescu. BitTorrent session characteristics and models. In D. Kouvatso, editor, *Technical Proceedings. HET-NETs '05 - 3rd International Working Conference on Performance Modelling and Evaluation of Heterogeneous Networks*, 2005.
- [8] D. Erman, D. Ilie, and A. Popescu. Peer-to-peer traffic measurements. Technical report, Blekinge Institute of Technology, Karlskrona, Sweden, 2005.
- [9] J.-I. Gailly and M. Adler. zlib. <http://www.gzip.org/zlib>, August 2005.
- [10] K. P. Gummadi, R. J. Dunn, S. Saroiu, L. H. M. Gribble, Steven D., and J. Zahorjan. Measurement, modeling and analysis of a peer-to-peer file-sharing workload. In *19th ACM Symposium on Operating Systems Principles*, pages 314–329, Bolton Landing, NY, U.S.A, Oct. 2003.
- [11] D. Ilie, D. Erman, and A. Popescu. Traffic measurements of P2P systems. *Swedish National on Computer Networking Workshop (SNCNW04)*, November 2004.
- [12] V. Jacobsen, C. Leres, and S. McCanne. Tcpdump. <http://www.tcpdump.org>, August 2005.
- [13] T. Klingberg and R. Manfredi. *Gnutella 0.6*. The Gnutella Developer Forum (GDF), 200206-draft edition, June 2002. http://groups.yahoo.com/group/the_gdf/files/Development/.
- [14] A. M. Law and W. D. Kelton. *Simulation Modeling and Analysis*. McGraw-Hill, 3rd edition, 2000. ISBN: 0-07-059292-6.
- [15] R. Manfredi. *Gnutella Traffic Compression*. The Gnutella Developer Forum (GDF), Jan. 2003. http://groups.yahoo.com/group/the_gdf/files/Development/.
- [16] C. Rohrs. *Query Routing for the Gnutella Network*. Lime Wire LLC, 1.0 edition, May 2002. http://groups.yahoo.com/group/the_gdf/files/Development/.
- [17] S. Saroiu, P. K. Gummadi, and S. D. Gribble. Measuring and analyzing the characteristics of napster and gnutella hosts. *Multimedia Systems Journal*, 8(5), 2002.
- [18] T. Schürger. *Horizon size estimation on the Gnutella network v0.2*, Mar. 2004. <http://www.menden.org/gnutella/hsep.html>.
- [19] A. Singla and C. Rohrs. *Ultrapeers: Another Step Towards Gnutella Scalability*. Lime Wire LLC, 1.0 edition, Nov. 2002. http://groups.yahoo.com/group/the_gdf/files/Development/.
- [20] D. M. Titterton, A. F. M. Smith, and U. E. Makov. *Statistical Analysis of Finite Mixture Distributions*. John Wiley & Sons, 1985. ISBN: 0-471-90763-4.