Markovich, Natalia M.; Krieger, Udo R.

Analysis of Packet Transmission Processes in Peer-to-Peer Networks by Statistical Inference Methods

Date of secondary publication: 22.04.2026

Accepted Manuscript (Postprint), Bookpart

Persistent identifier: urn:nbn:de:bvb:473-irb-114801x

Primary publication

Markovich, Natalia M.; Krieger, Udo R. (2013): Analysis of Packet Transmission Processes in Peer-to-Peer Networks by Statistical Inference Methods, in: Ernst Biersack, Christian Callegari, und Maja Matijasevic (Ed.), Data traffic monitoring and analysis : from measurement, classfication, and anomaly detection to qualtiy of experience, Berlin ; Heidelberg: Springer, pp. 104–119, doi: 10.1007/978-3-642-36784-7_5.

Analysis of Packet Transmission Processes in Peer-to-Peer Networks by Statistical Inference Methods

Natalia M. Markovich[1] and Udo R. Krieger[2,⋆]

[1] Institute of Control Sciences, Russian Academy of Sciences, Moscow, Russia
markovic@ipu.rssi.ru
[2] WIAI, Otto-Friedrich-University, Bamberg, Germany
udo.krieger@ieee.org
http://www.uni-bamberg.de/ktr

Abstract. Applying advanced statistical techniques, we characterize the peculiarities of a locally observed peer population in a popular P2P overlay network. The latter is derived from a mesh-pull architecture. Using flow data collected at a single peer, we show how Pareto and Generalized Pareto models can be applied to classify the local behavior of the population feeding a peer. Our approach is illustrated both by file sharing data of a P2P session generated by a mobile BitTorrent client in a WiMAX testbed and by video data streamed to a stationary client in a SopCast session. These techniques can help us to cope with an efficient adaptation of P2P dissemination protocols to mobile environments.

Keywords: heavy hitter model, Generalized Pareto distribution, peer-to-peer network, change-point detection.

1 Introduction

In recent years modern dissemination platforms employing peer-to-peer (P2P) overlay protocols have gained increasing interest. They have been derived from BitTorrent and its ramifications, for instance, GoalBit, Zattoo, PPLive, SopCast, Voddler or Skype (c.f. [4], [22], [24]). The latter have become mature middleware systems to distribute real-time media streams among interested clients. The estimation of the traffic matrices arising from such P2P based streaming or file sharing services and the integration of these services into mobile environments demand a solid analysis of the generated packet streams and the required capacity on an IP underlay network (cf. [15], [27]).

Components of P2P teletraffic engineering comprise basic elements such as monitoring and analysis of overlay and underlay structures, the design of topology aware routing and QoS driven peer- as well as piece-selection algorithms and the modeling of the resulting workloads (cf. [3], [7], [13], [17], [25], [26]). The

⋆ The authors acknowledge the partial financial support by the ESF-project COST IC0703.

latter tasks involve the statistical characterization of P2P packet flows and the determination of an effective bandwidth regarding aggregated P2P traffic on IP network links. For this purpose the required delay-loss profiles of a media service such as the tolerable packet and frame losses and the playback delay bounds of the media players should be taken into account.

Compared to a parametric teletraffic approach, purely measurement based concepts may provide an alternative. The latter can be integrated easily into real-time control components of the P2P middleware and handle the evolution of the P2P protocols more rapidly. To respond to this fast deployment of P2P overlay structures, we have developed a comprehensive P2P traffic measurement, modeling and teletraffic analysis concept (cf. [20]). It integrates four orthogonal dimensions to cope with the analysis of P2P structures and P2P traffic characterization: (1) traffic measurements at the packet layer combining passive and active monitoring techniques by the open source tool Atheris [1], [8], (2) data extraction, analysis and inspection of P2P overlays based on a hierarchical multi-layer modeling concept [20], (3) a characterization of the overlay structure by techniques and metrics of complex networks, and, last but not least, (4) a non-parametric teletraffic modeling approach based on the statistical characterization of P2P traffic [19]. The latter relies on the analysis and estimation of the bivariate distribution $\mathbb{P}\{X_i \leq x, Y_i \leq y\}$ of packet inter-arrival times X_i and packet lengths Y_i extracted from corresponding aggregated flows of P2P conversations and their collected i.i.d samples $\{(X_1, Y_1), ..., (X_n, Y_n)\}$.

Considering the control plane of a P2P system at the side of mobile clients, these building blocks constitute basic elements of a middleware concept derived from observation-driven, sound control-theoretical design principles. The latter should automatically adapt itself to the states of a very dynamic mobile environment, in particular to handover management, using the fundamental feedback principles of a Luenberger state observer and modeling results of teletraffic theory (cf. [3], [7], [9], [23]). To limit complexity, the information of the controls should only stem from local observations at the packet and session layers of the clients. Our recently developed system RapidStream [2] which offers a P2P video service for Android smartphones illustrates this necessity (cf. [10]). By these means we want to enhance the heuristic control approaches developed in the area of P2P protocols. Improvements concerning adaptive peer-selection and chunk scheduling can be implemented into an enhanced version of the prototype Rapidstream and its performance can be evaluated by the distributed monitoring capabilities of a currently developed enhanced variant of Atheris [1].

It is the objective of this paper to elaborate on new building blocks of our statistical approach. It is motivated by a statistical characterization of the local peculiarities of a peer population feeding a single client. The latter concept can be combined with the bandwidth measurement functionality of the Atheris approach [8] and a control-theoretical concept to identify dynamically the most important peers and to adapt peer-selection policies in mobile networks. These peers should provide a sufficiently large packet input rate to feed the media player without performance degradation (cf. [4], [9]). This limitation to useful peers is of particular importance if a mobile environment is considered (cf. [10], [11]).

We have partially validated the developed classification concept by experimental testbeds collecting a rich set of traces (see, e.g., [9], [15], [20]). Due to the renewal process character of chunk request-reply patterns during the lifetime of a P2P session we have seen that one does not gain more information if a huge set of traces is collected. The basic locally observed dissemination features of a P2P protocol that we want to investigate are already manifested in small sets. Here we use two traces as an example to illustrate some of these results. On one hand, a trace of P2P streaming traffic observed at a monitored stationary SopCast client has been used as representative of a high-speed wireline network (cf. [20]). On the other hand, a P2P file sharing trace collected in a WiMAX testbed in Seoul, Korea, is evaluated as example of a wireless environment (cf. [9], [15]). There a mobile BitTorrent client has been traveling in a bus. Regarding the provisioning of mobile high-speed services in an urban environment, this scenario provides an interesting example. This Asian setting is different from the normal European one and WiMAX an interesting competitor of UMTS or LTE networks. Based on the corresponding measurement studies we provide further characteristics of the latter peer-to-peer file sharing and streaming services. The data analysis illustrates the features of our analysis approach, partly validates former findings (cf. [9], [26]) and extends our concepts on P2P traffic characterization and measurement-driven classification of a locally observed peer population by adequate teletraffic models.

The rest of the paper is organized as follows. In section 2 the statistical analysis and modeling of packet flows of a locally observed peer population in terms of a Pareto model is presented. It is illustrated by superimposed packet flows feeding a monitored SopCast client in an observed P2P session. In section 3 the characterization of the packet transmission processes which are locally observed at a mobile client of a typical BitTorrent session in a WiMAX environment by means of a generalized Pareto model is stated. Principal component analysis is applied to prove that both volume and count based attributes of the monitored flows will provide an equivalent classification of the locally activated peer population. Finally, some conclusions on peer-to-peer protocols and their application to mobile networks are drawn.

2 Characterizing the Size-Count Disparity of a P2P Session by a Pareto Model

In the following we consider overlay networks derived from the mesh-pull architecture of a modern P2P dissemination service and use streaming data of a SopCast session as illustrative example (cf. [20]). Such an overlay network normally embeds a monitored peer, here called home peer p_0, requesting a certain object like a media file or a video stream into a dense mesh of n feeding peers p_i from a finite peer population $\mathcal{U}$ which share common interest in a specific object.

In a streaming context this strategy will guarantee a reliable and timely supply of chunk sequences to the streaming engine of a host. For instance, our

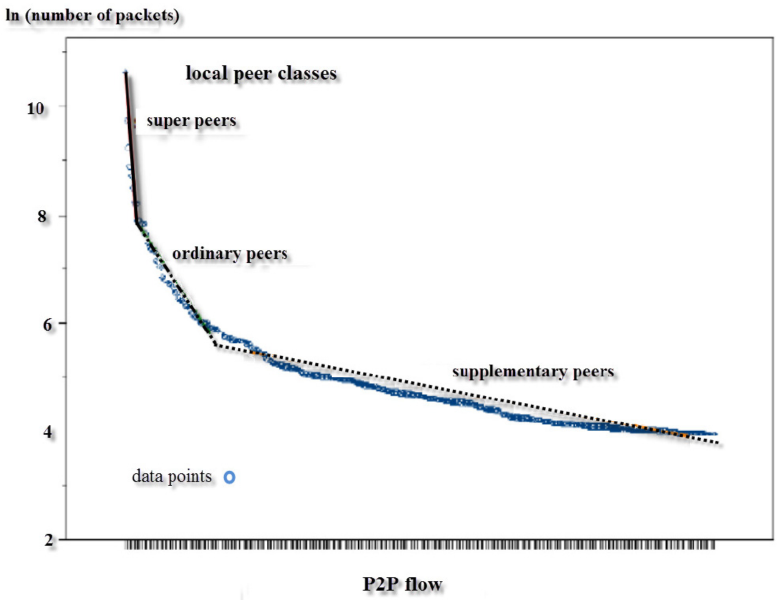

Fig. 1. Classification of peer relations based on the packet flows of all active peers feeding a stationary home peer during a SopCast session (cf. also [20, Fig. 8])

analysis [20] that will be used as illustrative example and subsequently be further extended has revealed that SopCast clients are typically connected to up to thousand different peers during the lifetime of a session and the exchanged flow data exhibit a hierarchical pattern.

This feature and the instantiated packet flow relations have been observed for other P2P applications as well and are caused by the applied peer and piece selection strategies of the underlying P2P protocol (cf. [7], [26]). Therefore, it is necessary to investigate the preference relationship among the packet flows of the active peers in order to understand the inherent local hierarchical structure of the mesh-pull topology. As basic metric one can use the number of transferred packets of all active flows feeding the home peer p_0 and the intensity or the volume of these flows which are depicted on a logarithmic scale.

Our previous investigations have revealed that the first criterion, for example, is simple to monitor and allows us to distinguish three local levels of peers $p_i \in \mathcal{U}$ associated with a home peer $p_0 \in \mathcal{U}$ during a session, namely his individual *super peers*, *ordinary peers* and *supplementary peers* (see fig. 1, also [20, Fig. 8]). Such hierarchical flow patterns are observed for other protocols as well, and thus bear a certain generality.

The investigation of the transferred accumulated volumes among these n active peers of the realized flow graph G_V, in particular the number of packets and the byte volumes flowing inbound and outbound to a monitored home peer p_0, obviously reveals the local hierarchical structure of the overlay and expresses the implemented selection strategies (see fig. 1, cf. [20], [26]).

We intend to describe these observed structures by a universal teletraffic model that can be used in further design studies or become a load model of other performance investigations, e.g. in a simulation of mobile clients. If we arrange the n flows according to the number of exchanged packets on a logarithmic scale, we can realize the hierarchy of the locally instantiated peer classes. Interpreting the relative number of transferred packets as frequency f_i to select a feeder p_i, we can model the ranked selection process by a versatile heavy-tailed distribution of a random variable (rv) Y on the integers $\mathbb{N}$. It should obey a distribution function (df) of a generalized Zipf type. We may choose, for instance, a special case of the zero-truncated Lerch distribution with probability mass function (pmf)

$$p_k = \mathbb{P}\{Y = k\} = C \cdot \frac{p^k}{(a+k)^\alpha} , \qquad k \in \mathbb{N}, \tag{1}$$

the parameters $a > -1$, $p \in (0, 1]$, and the tail coefficient $\alpha \in \mathbb{R}$. The normalization constant $C = [p\Phi(p, a + 1, \alpha)]^{-1}$ is defined in terms of the Lerch' transcendent $\Phi(p, a, \alpha) = \sum_{k=0}^{\infty} p^k \cdot (a + k)^{-\alpha}$, (cf. [5], [28]). Selecting the parametrization $p = 1, a \geq 0, \alpha > 1$, we get the Zipf-Mandelbrot pmf $\mathbb{P}\{Y = k\} = C \cdot (a + k)^{-\alpha}$, $k \in \mathbb{N}, C^{-1} = \Phi(1, a+1, \alpha)$ and with the restriction $a = 0$ the well-known Zipf law

$$\mathbb{P}\{Y = k\} = C \cdot \frac{1}{k^\alpha} , \quad k \in \mathbb{N}, \quad C^{-1} = \Phi(1, 1, \alpha). \tag{2}$$

If we plot the rank-frequency relationship of the relative number f_i of packets of all inbound flows $\phi(p_i, p_0)$ transferred from a peer p_i to the home peer p_0 on a log-log scale, we can identify in this case a linear relation between $\ln f_i$ and $\ln i$ of the related ranks i of the feeders p_i (see also [20]). Thus, a pmf of Zipf type (2) will adequately describe the local hierarchical peer structure seen by the home peer.

In our example of a SopCast session depicted in fig. 1 this feature has been validated (cf. [20])). Subsequently, we will see that the basic BitTorrent protocol itself, which has inspired the development of SopCast, behaves slightly different and must be modelled in a modified manner.

If we interpret the transferred number of packets of a flow $\phi(p_i, p_0)$ as realization x_i of an equivalent income $X_i \in \mathbb{R}$ of the feeding peer $p_i, i \in \{1, \ldots, n\}$, we can represent the local P2P packet model of a session by a corresponding heavy-tailed physical model of Pareto type with a rv X and its sample $\{X_1, \ldots, X_n\}$ (cf. [20], [21]). We denote the distribution function of this Pareto model associated with the Lerch ranking law by

$$F(x) = \mathbb{P}\{X \leq x\} = \int_{x_0}^{x} f(t)dt = 1 - Cx^{-\beta} = 1 - \left(\frac{x}{x_0}\right)^{-\beta} \tag{3}$$

with $x \geq x_0 > 0$ and tail index $\beta = 1/\alpha$.

The corresponding $q = 1 - p$ quantile function of this Pareto model $F_q(x_p) = 1 - F(x_p) = p \in (0,1)$, i.e. $x_p = F_q^{-1}(x_0, p, \beta) = x_0 \cdot p^{-1/\beta}$, can be used to define the local classes of the feeding peers. Taking, for instance, the 2.5% or 5% as well as 10% and 20% levels for p, their quantiles $x_{0.025}, x_{0.5}, x_{0.1}, x_{0.2}$ may specify the break points of the local classes of super peers, ordinary and supplementary peers associated with a home peer. Alternatively, these break points may be determined automatically by change-point detection algorithms (see [12]).

Using the transferred numbers of packets $x_1 \geq x_2 \ldots \geq x_n$ of the flows or approximating the pmf $p_i = \mathbb{P}\{Y = i\}$ by the empirical values $\{f_i, i = 1, \ldots, n\}$ of the n flows, we can estimate the tail index β by Hill's estimate $\widehat{\beta}^{-1} = \frac{1}{n-1}\left(\sum_{k=1}^{n-1} \ln(\frac{x_i}{x_n})\right) = \frac{1}{n-1}\sum_{k=1}^{n-1} \ln(x_i) - \ln(x_n)$ or in terms of Newman's estimate $\widehat{\alpha} = \frac{1}{n}\sum_{k=1}^{n} \ln(\frac{f_i}{f_{\min}})$ where $f_{\min} = f_n$ represents the minimal measured value (cf. [18], [21]). It corresponds to the value of the gradient of the linear segment in the rank-frequency plot (cf. [20]).

To investigate the size-count disparity issues in more detail, we can apply the excess wealth function $W_X(\cdot)$ (cf. [16]). Regarding the Pareto model (3) or any heavy-tailed, nonnegative income variable X with finite first moment $\mu = \mathbb{E}(X)$ it is determined by the corresponding continuous distribution function $F(x)$ in terms of the excess wealth transform

$$W_X(q) := \int_{x_q}^{\infty} (1 - F(x))\, dx = W_X(F(x_q)) \tag{4}$$

of X. Here $x_q = F^{-1}(q), q \in (0,1)$ is the qth quantile of $F(x)$. The random variable $X_{ew} := W_X(F(X)) = \int_X^{\infty} (1 - F(t))\, dt$ defined for $F(X) = q \in (0,1)$, which is a uniformly distributed rv U on $(0,1)$, is called observed excess wealth. By the underlying integral transform of the survival function $1 - F(x) = \mathbb{P}\{X > x\}$ of X,

$$\pi_X(t) := \int_t^{\infty} (1 - F(x))\, dx = W_X(F(t)), \quad t \geq 0, \tag{5}$$

called stop-loss transform, it determines the observed excess wealth $X_{ew} = \pi_X(X)$. Hence, the survival function of X_{ew} can be simulated easily since it is the inverse of the excess wealth transform, i.e.

$$\mathbb{P}\{X_{ew} > y\} = \mathbb{P}\{W_X(U) > y\} = \mathbb{P}\{U \leq W_X^{-1}(y)\} = W_X^{-1}(y), \quad y \in (0, \mu).$$

Then we can conclude that for the Pareto model (3), or any other heavy-tailed model, the normalization by the mean is determined by

$$W_N(x_p) = \frac{W_X(1-p)}{\mu} = \frac{\pi_X(x_p)}{\pi_X(x_0)} = \frac{\int_{x_p}^{\infty} x f(x) dx}{\int_{x_0}^{\infty} x f(x) dx} = (\frac{x_p}{x_0})^{-(\beta-1)}. \tag{6}$$

Here x_p denotes the $q = 1 - p$ quantile, i.e. $\mathbb{P}\{X > x_p\} = p \in (0,1)$. It yields a size weighting of the wealth excess beyond x_p. It can be easily used to study size-count disparity issues in more detail by simply fitting the tail index β in (3) to the available data (cf. [20]).

In our context it means that the fraction of those packets sent by the most active part of the peers p_i is specified in terms of (6) by the ratio of the packet load of the fraction of flows exceeding level x_p compared to the total packet load. Hence, given the Pareto model (3) the most active $p \cdot 100$ % of the flows determine the fraction $W_N(x_p)$ of the sent packets by means of the $1-p$th quantile x_p in terms of $W_N(x_p) = p^{(\beta-1)/\beta}$. The latter term is also related to the Lorenz curve (cf. [21]).

Motivated by this analysis, we subsequently investigate as an alternative a parametric analysis and classification method based on a generalized Pareto model. We shall illustrate its relevance by means of the fundamental BitTorrent protocol in a mobile network environment.

3 Modeling Local Peculiarities of a Peer Population by a Generalized Pareto Distribution

In the previous section we have shown that the traffic flows which are realized in a P2P overlay network by a population $\mathcal{U}_n := \{p_0, p_1, \ldots, p_n\} \subseteq \mathcal{U}$ of active clients generate a characteristic load pattern at the single measurement point of an involved, monitored client p_0 during a P2P session. The classification of the observed population of feeding peers from $\mathcal{U}_{p_0} := \mathcal{U}_n \setminus \{p_0\}$ can be based on this information and determine an understanding of the underlying flow graph G_V of the P2P network (see [20], [26]). Since we can record different statistical attributes of the flows exchanged with the monitored client, such as the number of incoming or outgoing packets, the overall transferred packets, or the exchanged byte volumes, the question arises whether there is a most informative conversation attribute which is arising from the collected flow statistics.

3.1 Principal Component Analysis of the Exchanged Flow Data

Let R_i denote the number of packets that a monitored home peer p_0 has received from a feeding peer $p_i \in \mathcal{U}_{p_0}$ in the overlay network during a typical P2P session. Let $V_i^{(i)}$ be the volume of the inbound traffic received from peer p_i, $V_i^{(o)}$ be the volume of the outbound traffic sent from p_0 to peer p_i and $V_i^{(e)}$ be the volume of the overall traffic exchanged with peer p_i. We assume that the random variables of each attribute are governed by a common distribution and denote the underlying generic random variables by $R, V^{(i)}, V^{(o)}, V^{(e)}$, respectively. The classification of the observed population of feeding peers $p_i \in \mathcal{U}_{p_0}$ can be based on this information.

To investigate the issue whether there is a most informative entity in the flow statistics, one should study the potential correlation among the attribute variables of the packet flows feeding a home peer p_0 during a P2P session. If the volume oriented random variables $Y \in \{V_i^{(i)}, V_i^{(o)}, V_i^{(e)}\}$ and the simple counting statistics R_i are linearly correlated, then there should exist a corresponding linear function $y = f_k(x) = a_k x + b_k$, $k \in \{i, o, e\}$ such that the regression

$$V^{(k)} = f_k(R) = a_k R + b_k, \quad k \in \{i, o, e\} \tag{7}$$

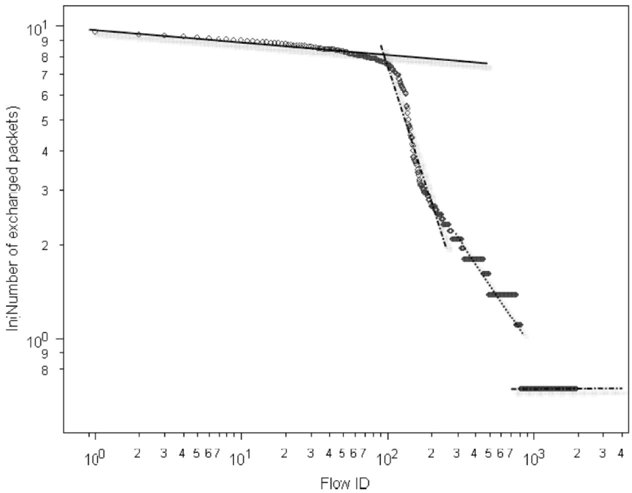

Fig. 2. Exchanged flow data of the WiMAX bus trace on a double logarithmic scale

holds. In this case a visual representation of the relationships $(R, V^{(k)}), k \in \{i, o, e\}$ should arrange the data of the peer flows along straight lines.

Principal component analysis (PCA) or other methods arising from factor analysis can be applied to study this issue in a rigorous manner.

3.2 Application to BitTorrent Data of a WiMAX Network

In the following we use the flow data captured by some BitTorrent bus trace in the WiMAX testbed in Seoul on March 17, 2010, as illustrative example (cf. [9], [15]). Studying the packet flows ϕ_i exchanged among the home peer p_0 and a feeder p_i in fig. 2, we see that there are at least four different tail regimes if we depict the ordered flow data on a double logarithmic scale. The shown structure illustrates that only the top 500 flows of the first two classes of super peers and dominant peers provide relevant information on the classification of the most dominant feeders within the peer population. Therefore, we will focus on the latter subset of the peer population in the subsequent investigations.

Applying the regression concept (7) to all flows ϕ_i of packets exchanged with the home peer p_0 that are arising from a feeder p_i and captured by the BitTorrent bus trace, for instance, we have realized (see [12, Fig. 3]) that there exists a perfect linear matching among all these variables as expected.

A factor analytic study of this data set using principal component analysis shows us that 99.49% of the variance in the flow data is explained by one major principal component depending on the four attributes $R, V^{(k)}, k \in \{i, o, e\}$. Each attribute of the flow data contributes between 24.7 to 25.1% to this dimension of the dominant principal component 1. The other components simply reflect the influence of the volume variables $V^{(k)}, k \in \{i, o, e\}$. Since all attributes nearly equivalently explain the variance of the dominant principal component, we see that there is no preferred variable carrying more information than all others. It confirms the result of the dependence plot in [12, Fig. 3].

In conclusion, we see that the simplest attribute 'R_i=number of packets incoming from (or exchanged with) a certain peer $p_i \in \mathcal{U}_{p_0}$' is sufficient to perform a sound classification analysis. Hence, a measurement campaign aiming to achieve the classification objective can be based on this simple packet counting metric R. There is no loss of information to apply it as foundation of a local peer classification approach and for further control actions in a mobile setting.

3.3 Generalized Pareto Modeling of Data Exchange Patterns

In the following we develop a more general parametric method to classify the exchange patterns within a local peer population $\mathcal{U}$ of an overlay network feeding a monitored home peer $p_0 \in \mathcal{U}$. We are looking for a mathematically sound procedure which can generate in a very efficient manner a reasonable partitioning of the set of all active peers $\mathcal{U}_n \subseteq \mathcal{U}$ feeding the observed home peer p_0 by means of their packet flows.

From a statistical point of view, we consider a sample $\mathcal{X}^{(n)} = \{X_1, X_2, \ldots, X_n\}$ of iid random variables $X_i, i \in \{1, \ldots, n\}$, which are governed by a common distribution function $F(x)$ of an underlying generic rv X. If we interpret X_i as number of transferred packets to or from peer $p_i, i \in \{1, \ldots, n\}$, during a monitored session, we have seen in the previous section that there is evidence that the related df obeys a heavy-tailed law of Pareto-type, i.e. $1 - F(x) = \mathbb{P}\{X > x\} \sim l(x)x^{-1/\gamma}$, for large $x \in \mathbb{R}^+$. Here $l(x)$ is a slowly varying function satisfying $\frac{l(tx)}{l(x)} \to 1$ as $x \to \infty, \forall t > 0$, and $f(x) \sim g(x)$ denotes the asymptotically identical growth rate, $\lim_{x\to\infty} f(x)/g(x) = 1$, of f and g (cf. [18, Def. 10, p. 4]). The positive tail index $\alpha = 1/\gamma > 0$ characterizes the slow decay of the tail of the distribution at infinity as basic model parameter. Different estimation techniques can be applied to determine the corresponding extreme-value index (EVI) $\gamma > 0$ by means of the sample (cf. [18, §1.2, p. 6f]).

Following classical extreme-value theory, we consider the absolute excess $X - u$ of the heavy-tailed rv X over a high threshold $u > 0$ subject to the condition of the exceedence $X > u$, i.e. the conditional rv $Z := X - u \mid X > u$. By the peaks-over-threshold method (POT) of extreme-value theory [18, p. 14], in particular Pickands' theorem, we can conclude that for increasing thresholds $u \to \infty$ the asymptotic distribution of rv Z is determined by a Generalized Pareto df (GPD(γ, σ))

$$\Psi(z) = \mathbb{P}\{Z \le z\} = 1 - (1 + \gamma z/\sigma)^{(-1/\gamma)}, \quad z \ge 0. \tag{8}$$

Here the scaling parameter $\sigma > 0$ is depending on the threshold u, i.e. $\sigma(u)$, and the EVI $\gamma > 0$ can be determined by the sample (cf. [6], [18, Sec. 1.2.3, p. 13]). Then we know that the transformation

$$Y = \frac{1}{\gamma} \ln\left(1 + \frac{\gamma}{\sigma} Z\right)$$

yields an exponentially distributed rv with unit scale parameter $\lambda = 1$, i.e. $\mathbb{P}\{Y \le y\} = 1 - \exp(-y), y \ge 0$, (cf. [14, Chap. 19.5.5, p. 240]).

Let us assume that Z obeys a GPD(γ, σ) law. Then we see that

$$\mathbb{P}\{Z - w \le z \mid Z > w\} = \frac{\mathbb{P}\{w < Z \le z + w\}}{\mathbb{P}\{Z > w\}} = \frac{\Psi(z+w) - \Psi(w)}{1 - \Psi(w)} \tag{9}$$

$$= \frac{(1+\gamma w/\sigma)^{(-1/\gamma)} - (1+\gamma(z+w)/\sigma)^{(-1/\gamma)}}{(1+\gamma w/\sigma)^{(-1/\gamma)}} = 1 - (1 + \gamma z/(\sigma + \gamma w))^{(-1/\gamma)}$$

holds and, hence, for all $w > u > 0$ the conditional rv $Z(w) := Z - w \mid Z > w$ satisfies also a GPD law with the modified scale parameter

$$\sigma(w) = \sigma + \gamma w > 0 \tag{10}$$

depending in a linear manner on the EVI γ. It means that we can consistently study GPD-like tail models of our original sample for all higher thresholds $w \ge u$ after identifying an appropriate initial threshold value $u > 0$, e.g. as appropriate high quantile of the empirical distribution of the flow data, such that the GPD hypothesis approximately holds. For this purpose we have to apply a linear scaling (10) of the initial scale parameter σ (see [6], [18]).

Thus we can consider a sample $\mathcal{X}^{(n)} = \{X_1, X_2, \ldots, X_n\}$ of n iid heavy-tailed rvs such as the number of packets R_i that a monitored home peer p_0 has received and that were sent by a feeding peer $p_i, i \in \{1, \ldots, n\}$. Then we select for a predetermined threshold $0 < u \in \mathbb{R}$ the corresponding sequence of N_u exceedences $\{X_{i_1}, \ldots, X_{i_{N_u}}\}$, where

$$\begin{aligned} i_1 &:= \min\{i \in \{1, \ldots, n\} \mid X_i > u\} = \min\{i \in \{1, \ldots, n\} \mid \mathbf{1}_{(u,\infty)}(X_i) = 1\} \\ i_{j+1} &:= \min\{i \in \{1, \ldots, n\} \mid i > i_j, \mathbf{1}_{(u,\infty)}(X_i) = 1\}, \quad j = 1, \ldots, N_u - 1 \\ N_u &:= \sum_{i=1}^{n} \mathbf{1}_{(u,\infty)}(X_i) \end{aligned}$$

determines the latter indices of the sample.

Based on this information we can estimate the unconditional tail distribution $\bar{F}(x) = \mathbb{P}\{X > x\}$ by the sample $\mathcal{X}^{(n)}$ in terms of $\widehat{H}(x) = \frac{N_u}{n}\left(1 + \frac{\widehat{\gamma}}{\widehat{\sigma}}(x-u)\right)$ and the logarithmically transformed excess process $\widehat{Y} = \frac{1}{\widehat{\gamma}} \ln(\frac{n}{N_u}\widehat{H}(x)) = \frac{1}{\widehat{\gamma}} \ln\left(1 + \frac{\widehat{\gamma}}{\widehat{\sigma}}(x-u)\right)$ (see [12], [18, (1.18), p. 14]). Here estimates $\widehat{\gamma}$ of the EVI γ and $\widehat{\sigma}$ of the scale parameter σ of the GPD are used (see (8), (10)).

3.4 Application to BitTorrent Data of a WiMAX Bus Trace

We will illustrate the classification concept derived from this generalized Pareto model by flow data exchanged with a monitored home peer in the overlay network of the swarm-like P2P protocol BitTorrent. We are again regarding a client moving by bus through Seoul on March 17, 2010, and consider the data set collected in the WiMAX testbed as representative example (cf. [9], [15]).

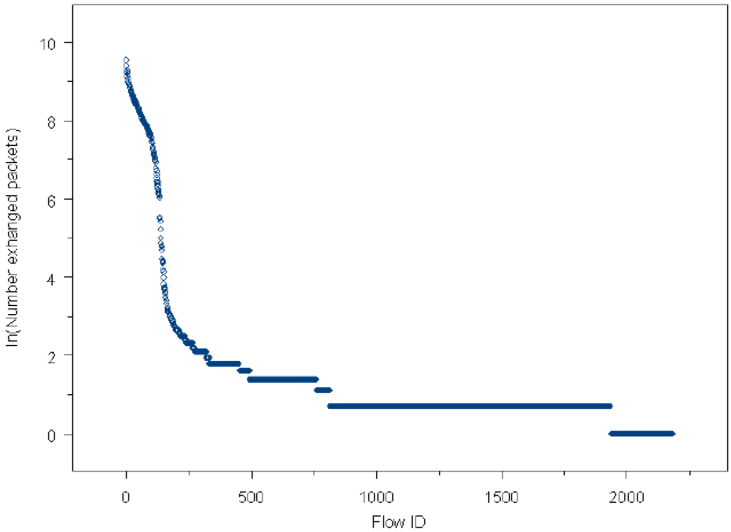

Fig. 3. Ln-transformed number of packets R_i exchanged with a peer p_i by a flow ϕ_i

A plot of the ordered number of packets R_i exchanged with active peers $p_i \in \{1, \ldots, n\}$ during this session characterized by a flow identifier i and its logarithmic transformation $\ln R_i$ in figure 3 illustrates that we can expect a heavy-tailed behavior of the most dominant peers due to the piecewise linear shape. Moreover, we see that there occurs a substantial change in the flow behavior within a part of the dominant flows (on the lhs in fig. 3). It is due to the separation of the feeding peer population into the groups of super peers, ordinary peers and the supplementary ones (see [9]).

An analysis of the dominant flows by a peaks-over-threshold (POT) methodology with a relatively high threshold u should reveal this behavior. Here we have chosen $u = 4000$ exchanged packets as separator of the flows. A fit of the available packet flow data $R_i, i = 1, \ldots, n = 2184$ of the WiMAX trace by the maximum likelihood method implemented by the routine 'fitgpd' of the R-package POT yields the parameter estimates $\widehat{\sigma} = 2344, \widehat{\gamma} = 0.02107$ of a Generalized Pareto distribution regarding the generic conditional excess variable $Z = R - u \mid R > u$ of exchanged packets. A related visual comparison of the counted flow data and the GPD model by a probability plot and a QQ plot with associated confidence intervals in fig. 4 illustrates that the data of all dominant flows are relatively well covered by the GPD distribution (see also [12, Sec. 3.2]).

If the threshold is chosen even higher, e.g. $u = 6000$, the data can also be covered by a GPD model, as shown in fig. 5 by the plots of the conditional excess and the tail distribution $\mathbb{P}\{Z \geq x\}$, arising from the data and the fitted GPD model. However, now the parameters of the fitted model have slightly changed as indicated by (9).

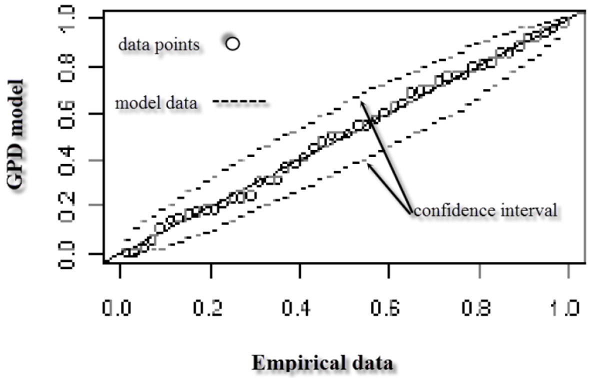

(a) Comparison by a probability plot

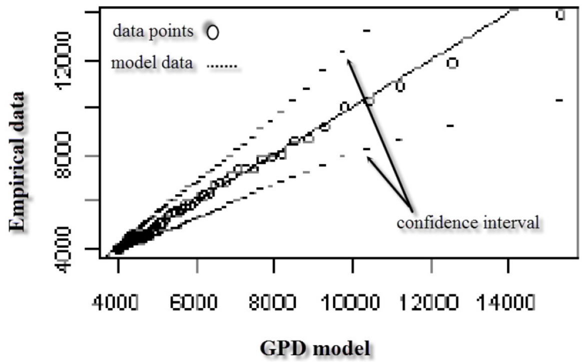

(b) Comparison by a QQ plot

Fig. 4. Fitted GPD tail model for threshold $u = 4000$

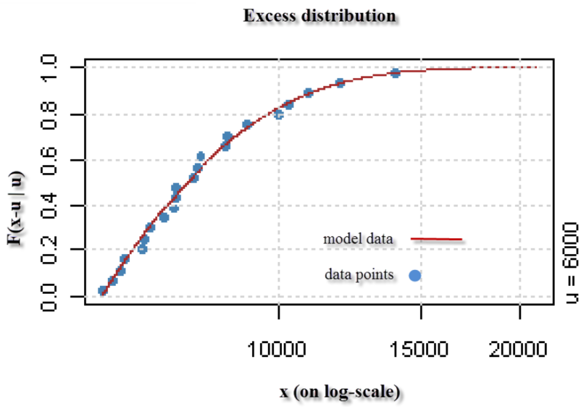

(a) Excess distribution on a logarithmic scale

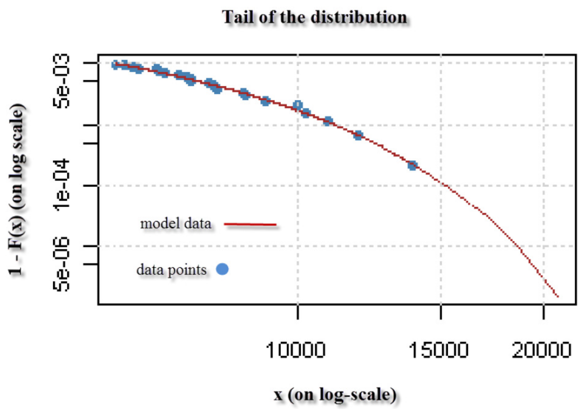

(b) Tail of the distribution on a double logarithmic scale

Fig. 5. Checking the accuracy of a fitted GPD tail model for threshold $u = 6000$

4 Conclusions

Looking at the rapid deployment of advanced multimedia applications by peer-to-peer overlay networks derived from a mesh-pull architecture and their adaptation to mobile networks, improved design approaches must be developed. To address the related challenges in teletraffic engineering, we have provided the new public source monitoring tool Atheris [1], [8] and recently presented a comprehensive analysis concept [20]. It integrates modeling, measurement, and statistical analysis of peer-to-peer traffic flows at the packet and session levels focusing on the insight gained by a single monitoring point in a P2P overlay network.

In this paper we have extended this approach and investigated improved statistical techniques to characterize the peculiarities of a locally observed peer population in popular P2P overlay networks. Using collected flow data at a single peer, we have shown how Pareto and Generalized Pareto models can be applied to isolate the most dominant part of the feeding population of a peer. Our scheme can be used to support the dynamic adaptation of a peer selection strategy to the real needs of streaming media feeding moving clients (cf. [9], [10]).

We have illustrated our classification approach by real flow data of P2P sessions generated by a mobile BitTorrent client in a WiMAX testbed and a stationary client in a SopCast streaming environment. Regarding the SopCast flows and their exchanged volumes, we have shown that there are strong indications that the dominant and most useful part of the peer population obeys a power law. Its distribution can be described by a Pareto model using the wealth excess transform. Considering BitTorrent flows, a generalized Pareto distribution can be applied successfully as appropriate model.

In this respect our modeling approach supplements the practical findings of Tang et al [26] about the overlay topology of SopCast and theoretical findings of Couto da Silva et al [7]. The latter have shown that a fluid-flow modeling approach can be used to describe the behavior of mesh-based P2P streaming systems like SopCast and PPLive and that bandwidth-aware peer scheduling compares favorably to location-aware and random peer selection schemes. Our statistical analysis discussed here and the previous measurement findings on SopCast [20] confirm da Silva's conclusion [7, p. 452] that the observed hierarchical peer structures are strongly influenced by the access bandwidth. According to our insights super peers have mainly high access bandwidth and an institutional embedding and can thus provide a useful upload service to ordinary peers behind standard access links. The active bandwidth measurement technique implemented by our open source tool Atheris [1] can be used by other investigators to study these relations for new P2P streaming protocols.

In [9] we have further realized that a BitTorrent-like protocol, especially BitTorrent's choking algorithm, is not well adapted to the fluctuating conditions in a mobile network. In this regard our sketched dynamic classification techniques have the potential to cope with an efficiency awareness of a P2P dissemination protocol in such an environment. Considering the control plane of a P2P system at a mobile client side, its realization by the tool RapidStream [2] and the support by new device-to-device protocols of mobile clients, these findings will guide the implementation of an improved middleware concept. It is derived from observation-driven, sound control-theoretical design principles. In future work we plan to support our teletraffic results by a distributed measurement campaign in UMTS and LTE networks (cf. [11]). The latter will guide the required efficient adaptation of peer-to-peer signaling and transport protocols in the overlay network spanning a mobile environment with gigabit links.

In conclusion, we are convinced that our integrated monitoring and analysis approach can provide a deeper insight into the dynamics of P2P multimedia applications and support the rapid development of appropriate teletraffic models and control algorithms at the packet and session levels regarding the edge link in front of a peer and the structural level of the overlay.

Acknowledgement. The authors express their sincere thanks to P. Eittenberger at Otto-Friedrich-University for his support.

References

1. Atheris, `http://sourceforge.net/projects/atheris/`
2. RapidStream, `http://www.ktr.uni-bamberg.de/project/rapidStream.shtml`
3. Aalto, S., et al.: Segmented P2P Video-on-Demand: Modeling and Performance. In: Proc. 22nd International Teletraffic Congress, Amsterdam, Netherlands (2010)
4. Ali, S., Mathur, A., Zhang, H.: Measurement of commercial peer-to-peer live video streaming. In: Proc. of ICST Workshop on Recent Advances in Peer-to-Peer Streaming, Waterloo, Canada (2006)
5. Aksenov, S.V., Savageau, M.A.: Some properties of the Lerch family of discrete distributions. University of Michigan (February 2008)
6. Castillo, E., Hadi, A.S.: Fitting the Generalized Pareto Distribution to Data. Journal of the American Statistical Association 92(440), 1609–1620 (1997)
7. Couto da Silva, A.P., et al.: Chunk distribution in mesh-based large-scale P2P streaming systems: A fluid approach. IEEE Trans. on Parallel and Distributed Systems 22(3), 451–463 (2011)
8. Eittenberger, P.M., Krieger, U.R.: Atheris: A First Step Towards a Unified Peer-to-Peer Traffic Measurement Framework. In: 19th Euromicro International Conference on Parallel, Distributed and Network-Based Processing, PDP 2011, Ayia Napa, Cyprus, February 9-11, pp. 574–581 (2011)
9. Eittenberger, P.M., Kim, S., Krieger, U.R.: Damming the Torrent: Adjusting BitTorrent-like Peer-to-Peer Networks to Mobile and Wireless Environments. Advances in Electronics and Telecommunications 2(3), 14–22 (2011)
10. Eittenberger, P.M., Herbst, M., Krieger, U.R.: RapidStream: P2P Streaming on Android. In: Proc. 19th International Packet Video Workshop, PV 2012, Munich, Germany, May 10-11, pp. 125–130 (2012)
11. Eittenberger, P.M., Schneider, K., Krieger, U.R.: Location-aware traffic analysis of a peer-to-peer streaming application in a HSPA network. In: 21st Euromicro International Conference on Parallel, Distributed, and Network-Based Processing, PDP 2013, Belfast, Northern Ireland, February 27-March 1 (to appear, 2013)
12. Eittenberger, P.M., Krieger, U.R., Markovich, N.M.: Teletraffic modeling of peer-to-peer traffic. In: Laroque, C., et al. (eds.) Proc. Winter Simulation Conference, WSC 2012, Berlin, Germany, December 9-12 (2012)
13. Hei, X., Liang, C., Liang, J., Liu, Y., Ross, K.W.: A measurement study of a large-scale P2P IPTV system. IEEE Tran. on Multimedia 9(8), 1672–1687 (2007)
14. Johnson, N.L., Kotz, S.: Distributions in Statistics - Continuous Univariate Distributions I. Wiley, New York (1970)
15. Kim, S., et al.: Measurement and analysis of BitTorrent traffic in mobile WiMAX networks. In: 10th IEEE International Conference on Peer-to-Peer Computing, IEEE P2P 2010, pp. 1–4. IEEE (2010)
16. Li, X., Shaked, M.: The observed total time on test and the observed excess wealth. Statistics & Probability Letters 68, 247–258 (2004)

17. Liu, F., Li, Z.: A Measurement and Modeling Study of P2P IPTV Applications. In: Proc. of the 2008 International Conference on Computational Intelligence and Security, vol. 1, pp. 114–119 (2008)
18. Markovich, N.M.: Nonparametric Analysis of Univariate heavy-tailed Data. Wiley, Chichester (2007)
19. Markovich, N.M., Krieger, U.R.: Statistical Analysis and Modeling of Skype VoIP Flows. Special Issue Heterogeneous Networks: Traffic Engineering and Performance Evaluation. Computer Communications 33, 11–21 (2010)
20. Markovich, N.M., Biernacki, A., Eittenberger, P., Krieger, U.R.: Integrated Measurement and Analysis of Peer-to-Peer Traffic. In: Osipov, E., Kassler, A., Bohnert, T.M., Masip-Bruin, X. (eds.) WWIC 2010. LNCS, vol. 6074, pp. 302–314. Springer, Heidelberg (2010)
21. Newman, M.E.J.: Power laws, Pareto distributions and Zipf's law. Contemporary Physics 46, 323–351 (2005)
22. Peltotalo, J., et al.: Peer-to-peer streaming technology survey. In: Proc. of the Seventh International Conference on Networking, ICN 2008, pp. 342–350. IEEE Computer Society, Washington (2008)
23. Qiu, D., Srikant, R.: Modeling and performance analysis of BitTorrent like peer-to-peer networks. In: Proc. ACM SIGCOMM 2004, pp. 367–378 (2004)
24. Sentinelli, A., Marfia, G., Gerla, M., Kleinrock, L., Tewari, L.: Will IPTV ride the peer-to-peer stream? IEEE Communications Magazine 45(6), 86–92 (2007)
25. Silverston, T., et al.: Traffic analysis of peer-to-peer IPTV communities. Computer Networks 53(4), 470–484 (2009)
26. Tang, S., Lu, Y., Hernández, J.M., Kuipers, F., Van Mieghem, P.: Topology Dynamics in a P2PTV Network. In: Fratta, L., Schulzrinne, H., Takahashi, Y., Spaniol, O. (eds.) NETWORKING 2009. LNCS, vol. 5550, pp. 326–337. Springer, Heidelberg (2009)
27. Xu, K., et al.: A Model Approach to Estimate Peer-to-Peer Traffic Matrices. In: Proc. IEEE INFOCOM 2011, Shanghai, April 10-15, pp. 676–684 (2011)
28. Zörnig, P., Altmann, G.: Unified representation of Zipf distributions. Computational Statistics & Data Analysis 19, 461–473 (1995)