Markovich, N. M.; Biernacki, A.; Eittenberger, P.; Krieger, U. R.

Integrated Measurement and Analysis of Peer-to-Peer Traffic

Date of secondary publication: 24.04.2026

Accepted Manuscript (Postprint), Conferenceobject

Persistent identifier: urn:nbn:de:bvb:473-irb-114818x

Primary publication

Markovich, N. M.; Biernacki, A.; Eittenberger, P.; Krieger, U. R. (2010): Integrated Measurement and Analysis of Peer-to-Peer Traffic, in: Evgeny Osipov, Andreas Kassler, Thomas Michael Bohnert, Xavier Masip-Bruin (Ed.), Wired / Wireless Internet Communications : 8th international conference, WWIC 2010, Luleå, Sweden, June 1 - 3, 2010 ; proceedings, Berlin ; Heidelberg: Springer, pp. 302–314, doi: 10.1007/978-3-642-13315-2_25.

Integrated Measurement and Analysis of Peer-to-Peer Traffic

N.M. Markovich[1], A. Biernacki[2], P. Eittenberger[2], and U.R. Krieger[2,⋆]

[1] Institute of Control Sciences, Russian Academy of Sciences, Moscow
[2] WIAI, Otto-Friedrich Universität, D-96045 Bamberg, Germany
`udo.krieger@ieee.org`

Abstract. We present a comprehensive traffic measurement and analysis concept to cope with the rapid deployment of peer-to-peer multimedia applications and their overlay structures. It integrates four orthogonal dimensions: traffic measurements at the packet layer, data extraction, flow analysis and inspection of peer-to-peer overlays based on a hierarchical multi-layer modeling concept, a characterization of overlay structures by techniques and metrics of complex networks, and a statistical characterization of peer-to-peer traffic features.

Keywords: Traffic measurements, peer-to-peer overlay, peer-to-peer traffic characterization, IPTV, SopCast.

1 Introduction

Currently, the rapid development of new multimedia services including VoIP, video streaming and video on demand as well as IPTV and the integration of triple play into the service portfolios of network operators indicate an evolutionary path towards next generation networking. The dissemination of multimedia content can be performed by diverse communication channels according to different paradigms including broadcast, multicast, unicast or peer-to-peer communication patterns (see Fig. 1). While network operators and content providers have advocated in favour of multicast architectures adopting the technology of content distribution network providers, peer-to-peer technology has gained an increasing attention in recent years.

Peer-to-peer (P2P) overlay networks are used with growing intensity to implement the content dissemination and control planes of many new portals and their multimedia service components like GoalBit, Zattoo, PPLive, SopCast, Voddler or Skype (see, e.g., [1], [8], [12], [13]). The availability of DSL-based high speed access to the Internet in residential networks and powerful stationary and mobile terminals have fertilized the rollout of these service platforms and will intensify the demand in the next years.

To respond to this rapid deployment of P2P overlay structures, we have developed a comprehensive P2P traffic measurement, modeling and teletraffic analysis

---

[⋆] The authors acknowledge the support by the projects BMBF MDA08/015 and COST IC0703.

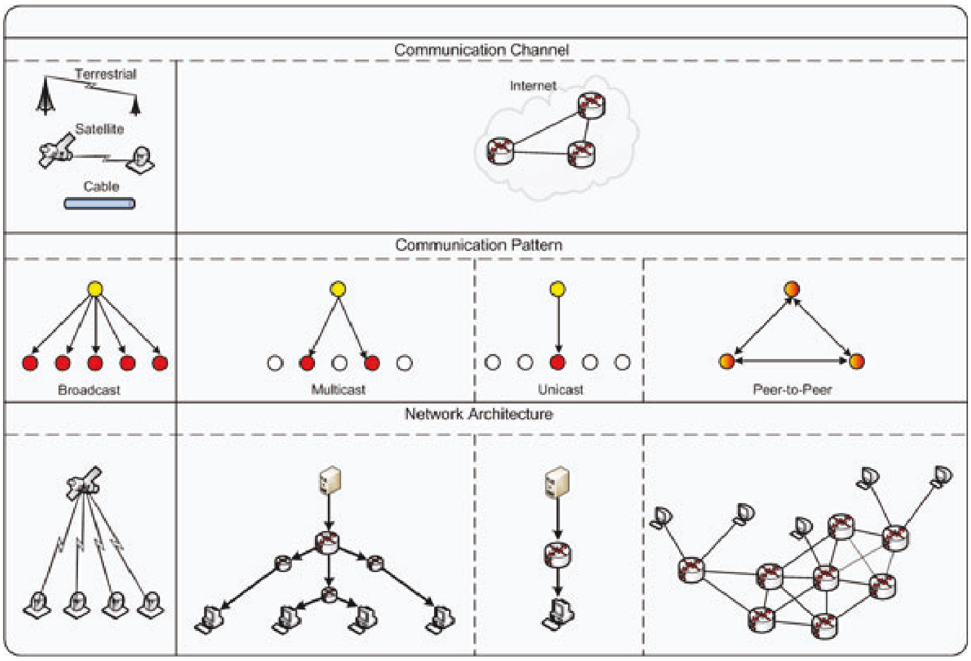

Fig. 1. Network architectures supporting multimedia content distribution

concept. 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, (2) data extraction, analysis and inspection of P2P overlays based on a hierarchical multi-layer modeling concept, (3) a characterization of the overlay structure by techniques and metrics of complex networks, and (4) teletraffic modeling based on a statistical characterization of P2P traffic features. It incorporates the analysis and estimation of the bivariate distribution $\mathbb{P}\{X_i \leq x, Y_i \leq y\}$ of packet interarrival times X_i and packet lengths Y_i extracted from corresponding flow data $\{(X_1, Y_1), ..., (X_n, Y_n)\}$ of collected i.i.d samples.

In recent years numerous measurement studies of different VoIP and P2PTV applications have been performed, e.g. by Ciullo et al. [3], Hei et al. [5], Liu et al. [9], Silverston et al. [14], and Tang et al. [15] among others. We have also recorded a moderate number of traces of prototypical streaming applications in 2007 and 2009 and performed a statistical analysis of the peer-to-peer transport services (cf. [4], [8], [10]). In this paper we present our P2P modeling approach and an integrated measurement and teletraffic analysis concept. Based on a measurement study in a home environment we further discuss some characteristic traffic features arising from a SopCast session. The data analysis illustrates the features of our approach, partly validates former findings and provides new teletraffic models of the transport service applied by SopCast.

The rest of the paper is organized as follows. In Section 2 we provide an overview of the modeling and measurement concept and its integration into the tool Atheris. In Section 3 the data analysis and modeling of typical P2PTV packet traffic arising from a SopCast session is discussed. Finally, some conclusions are drawn.

2 A P2P Modeling, Measurement and Analysis Concept

2.1 Modeling P2P Overlay Networks

Peer-to-peer multimedia applications establish an overlay on top of the packet-switched TCP/IP transport network based on a tree or mesh topology. Then they normally apply a pull mechanism to distribute the video content among all peers looking to a certain live TV channel or recorded video stream. The latter is described by an object space $\mathcal{O} = \{O_1, \ldots, O_m\}$.

From a logical perspective, this approach generates a multi-tier peer-to-peer architecture of the streaming infrastructure where the overlay network is mapped by a dissemination topology of the replicated streamed video data onto the used transport infrastructure derived from the TCP/IP protocol suite (see Fig. 2).

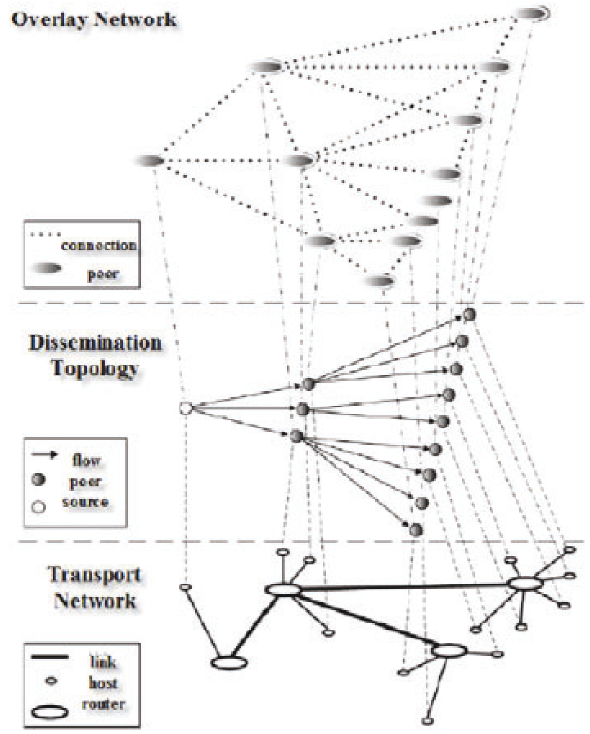

Fig. 2. Multi-tier peer-to-peer architecture of a content dissemination infrastructure

Inspired by the approach of Tang et al. [15], we can describe this hierarchical multi-layer infrastructure to analyze the P2P overlay structure, to describe the P2P traffic relations, and to perform a traffic characterization as well as the interactive investigation of corresponding traffic flows during P2P sessions. It covers the observed parts of the overlay structure by means of an undirected *neighborhood graph* $G_N = (V_N, E_N)$ among an object specific, time dependent peer population $\mathcal{P}^{(O_i)}(t) = \{p_1(t), \ldots, p_{n(t)}(t)\} \subset \mathcal{U}$ within a finite universe $\mathcal{U}$, see Fig. 2. The engaged peers form an overlay structure of common interest called *neighborhood community* (cf. [15]). It dynamically evolves in time since peers may join or leave this community of common interest related to an object O_j, e.g. due to channel hopping. The exchange of content data is described by a directed *dissemination flow graph* $G_V = (V_V, E_V, f, c)$, $V_V \subseteq V_N$, among pairs (p_i, p_j) of peers derived from a dissemination topology (V_V, E_V) (cf. [15]). $e = (p_i, p_j) \in E_V$ means that a flow $\phi(e)$ of requested video chunk sequences called *micro-flows* is transferred from peer p_i to p_j on request of p_j, see Fig. 3. In our single-site monitoring approach p_j represents a home peer. Further, we can study the feeding relations of this home peer as source (see Fig. 2). G_V includes a capacity function $c : V_V \times V_V \longrightarrow \mathbb{R}_0^+$ and a flow function $f : V_V \times V_V \times T \longrightarrow \mathbb{R}_0^+$ which assigns to each flow $\phi(p_i, p_j)$ from p_i to p_j its intensity as flow rate $f(p_i, p_j, t) \geq 0$ and an attribute $t \in T = \{t_1, \ldots, t_h\}$

determining the flow type. Thus we can distinguish between different kinds of control and content flows.

Finally, the underlying TCP/UDP-IP transport network connecting the peers p_i, p_j by flow paths $p = w(p_i, p_j) = (e_1, \ldots, e_{n_p})$ along capacity constrained links (e_k, c_k) is taken into account(see Fig. 2).

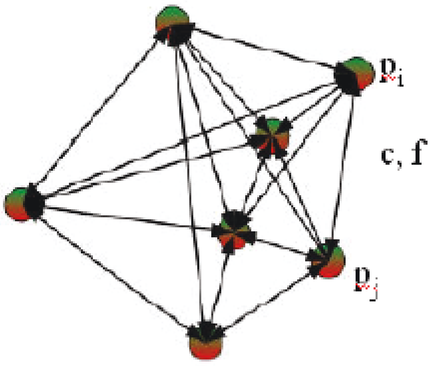

Fig. 3. Flow dissemination graph of a P2P model

The capacity function c of a link in the overlay network is determined by the bottleneck capacity of the underlying path $w(p_i, p_j)$ in the router network and normally unknown. Thus, one has to measure it by appropriate bandwidth estimation techniques and to determine the length (i.e. hop count) $L[w(p_i, p_j)]$ and round trip delay $R[w(p_i, p_j)]$ of the route in the transport network by path-pinging the corresponding hosts of the peers. To get c for both directions, it is necessary to implement a measurement instrumentation in the peers feeding the home peer demanding an expensive instrumentation of the complete infrastructure, for instance, by PlanetLab experiments which is out of scope.

In 2007 and during the second quarter of 2009 we have performed measurement studies to collect representative traces of Joost, Zattoo, PPLive and SopCast sessions covering both a wired and a wireless access to the Internet. The used test bed of the latter campaign has been developed to study the basic operation of the P2PTV streaming system SopCast in the typical home environment of a German customer (cf. [4]). Traces with volumes of 136 to 138 MB arising from live-streamed soccer matches during SopCast sessions of approximately 30 minutes have been gathered at the portable host of the SopCast client. In our further illustrations we will use the corresponding packet flows extracted from this traces.

In summary, the modeling approach of an overlay structure yields a hierarchical aggregated flow model of superimposed micro-flows in upward and downward direction to a home peer. Using the monitored data, it allows us to describe the exchange of chunk sequences among the peers by appropriate teletraffic models.

2.2 Determination of the P2P Flow Graph

In our measurement concept we have combined a passive single point monitoring of a P2P overlay network at a selected *home peer* p_1 with the active probing of the connections $\{p_1, p_j\} \in E_N$ to this observed peer p_1 in a time interval $[0, t]$. To identify the flow graph from the perspective of this single monitoring site, we apply trace routing at the IP layer of the transport paths $p = (e_1, \ldots, e_{n_p})$ starting from (or leading to) p_1. Further, we are able to estimate dynamically the available bandwidth $c = \min\{c_i \mid e_i, i = 1, \ldots, n_p\}$ of the path p along its constituting links e_i with the inferred (uplink) capacities c_i.

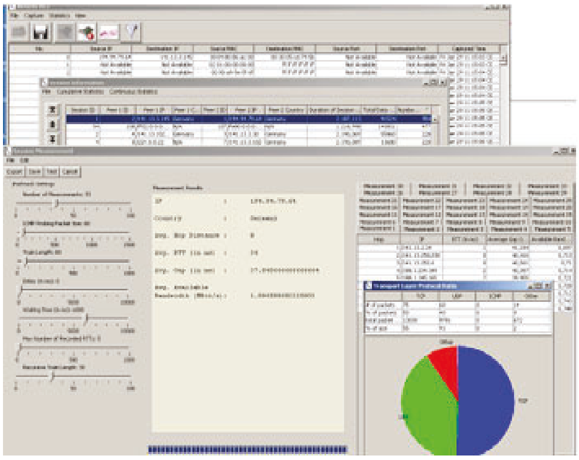

Fig. 4. Path probing

Fig. 5. Visualization of a bandwidth estimate

Since an end-to-end monitoring approach with two active probing and evaluating points is not available for bandwidth estimation, we have integrated the sender based approach implemented by the tool *Pathneck* (cf. [6]).

The individual flows of multimedia content feeding the home peer or fed by the latter have to be identified in the collected traces by post-processing steps. For this purpose, the flows can be specified as sequence of related IP packets determined by the conventional 5-tuples (A_s, P_s, A_d, P_d, pr) of the source and destination addresses A_s, A_d, the corresponding ports P_s, P_d and the related protocol identifier pr. They characterize the endpoints of a TCP or UDP transport connection along a path p, segment and assemble the packet sequences into flows based on various timeout methods, e.g. a fixed elapsed time determining flow expiration. Within a set of traces we have to derive the latter structures. We can, for instance, apply the corresponding C and Perl functions like *crl_flow* of the flows stack software implemented by the universal CAIDA suite *CoralReef* for passive Internet traffic monitors (cf. [7]). It offers the capability to determine the flow length, FlowType and CounterType of the packet streams in time intervals of predefined fixed length and to store it into Tables as well as to get access to the data and process them efficiently (cf. [7]). By appropriate filtering a flow analysis generates lists of detected flows within a P2P session of a home peer including the packet and volume statistics of each flow. Further, we can calculate the instantaneous traffic rates $\lambda_{j,i}(t) = \lim_{\Delta t \to 0} \frac{1}{\Delta t} \mathbb{P}\{N_{j,i}(t + \Delta t) - N_{j,i}(t) = 1 \mid N_{j,i}(t), \sigma(t)\}$ that are related to the packet counting process $N_{j,i}(t)$ of a flow $\phi(p_j, p_i)$ by the ratio of the observed number of packets in a finite interval of length $\Delta t > 0$ (based on the past history $\sigma(t)$ of the packet arrivals in $[0, t]$). Additionally, a flow selection arising from a set of most active peers by Table operations, flow aggregation into superimposed streams to a home peer, the construction of an IP-matrix table or the mapping of addresses to countries are possible, too (cf. [7]). This functionality has been provided by a virtual machine encapsulating the corresponding functions of CoralReef.

We have further integrated into our concept the capability to investigate flows within peer-to-peer sessions by adequate inspection techniques using empirical statistical characteristics like histograms (see Figs. 4, 5). We are also able to analyze the geographical distribution of the peer population feeding a monitored home peer.

2.3 Integration of Measurement and Analysis by the Tool Atheris

The concept has been used to develop the new versatile, extensible JAVA measurement and analysis tool Atheris based on a modular object-oriented software design. It has been implemented both for a Linux and Windows OS (see Fig. 6). The Linux version is built on top of libpcap and its Java wrapper Jpcap and integrates an adjusted C version of the bandwidth estimation tool Pathneck. The Windows version relies on WinPcap, Jpcap, and a Java based implementation of the bandwidth estimation system.

The major advantage of Atheris arises form the possibility to simultaneously monitor and analyze in an on-line manner P2P applications and to probe the dissemination paths p to a home peer. One key analysis feature concerns the capability to estimate the available bandwidth c of specific paths p and to inspect interactively selected traffic flows $\phi(p_j, p_i, t)$ of a certain type t.

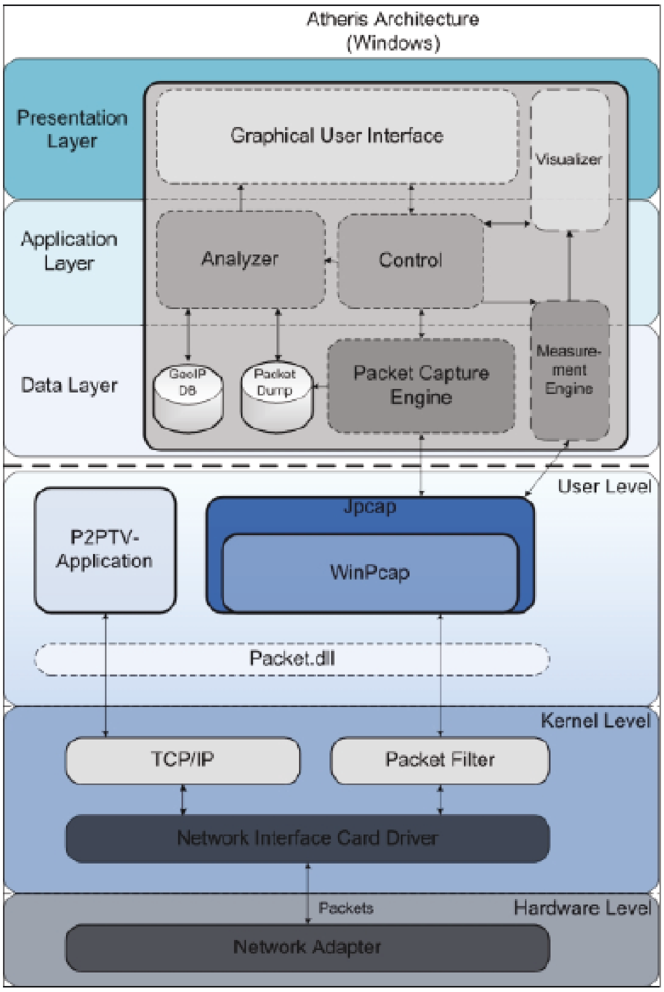

Fig. 6. Design of Atheris

3 Analysis and Teletraffic Modeling of P2P Flows

We show in the following section that single site observations in peer-to-peer networks can already reveal many details about the operational characteristics of a dissemination system despite all limitations of this approach (cf. [4], [15]).

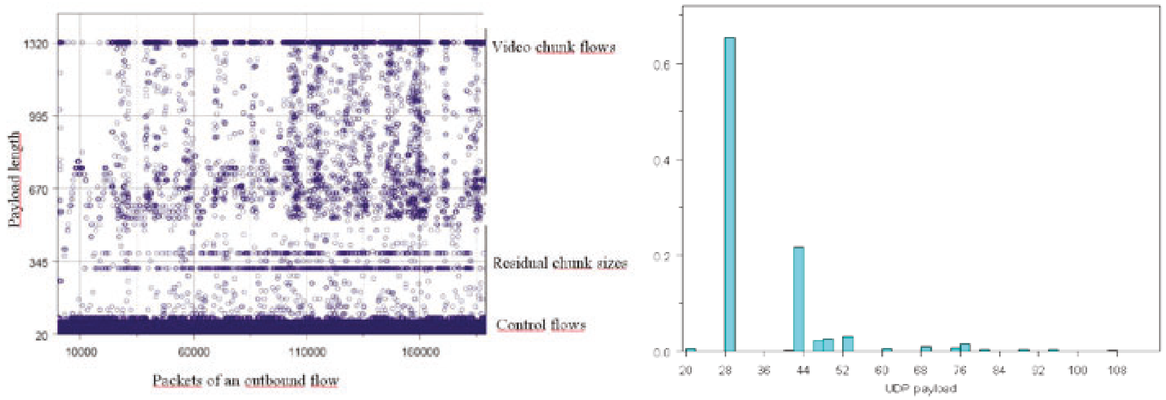

Fig. 7. Scatter plot of the UDP payload (in bytes) of frames of the superimposed outbound flows from a home peer (left) and a histogram of the UDP payload of all its signaling packets (right) during a typical SopCast session

3.1 Identification of Mesh-Pull Functions

A modern mesh-pull dissemination system like SopCast establishes a dense peer-to-peer mesh network to operate the delivery of the multimedia content in terms of smaller chunks (cf. [4], [15]). It applies a pull mechanism to distribute the content among all peers looking to a certain live TV channel or recorded video stream. SopCast employs mainly UDP as transport protocol to perform the topology discovery and maintenance tasks of the peer-to-peer mechanism as well as the control and delivery tasks of the video content dissemination among the involved peers. Empirical statistical analysis based on scatter plots and histograms and the visual inspection of the frame lengths of the packet flows instantiated during a typical P2P session between the observed client and the communicating peers can reveal valuable insights on the functioning of a mesh-pull system (see Fig. 7). For example, the typical SopCast session data depicted in Fig. 7 show the characteristic footprints of the frame length at 62, 70, 84, 88, 90, 94 bytes for control and 1362 bytes for content flows with a UDP payload of 20, 28, 42, 46, 48, 52 and 1320 bytes in agreement with [15]. A flow-based analysis is required to study the intrinsic structure of typical flows (cf. [4]).

3.2 Hierarchical Structure of a P2P Session

Overlay networks normally embed a home peer requesting a video stream into a dense mesh of feeding peers to guarantee a reliable and timely supply of chunk sequences to the streaming engine of a host. For instance, our analysis has revealed that SopCast clients are typically connected to up to thousand different peers during the lifetime of a session Therefore, it is necessary to investigate the preference relationship among the peer flows to understand the inherent hierarchical structure of the mesh-pull topology. One can use as metric the number of transferred packets of all active flows or those feeding the home peer, and the intensity or the volume of these flows depicted on a logarithmic scale.

The first criterion, for example, is simple to monitor and allows us to distinguish three local levels of peers associated with a home peer $p_1 \in V_V$ during

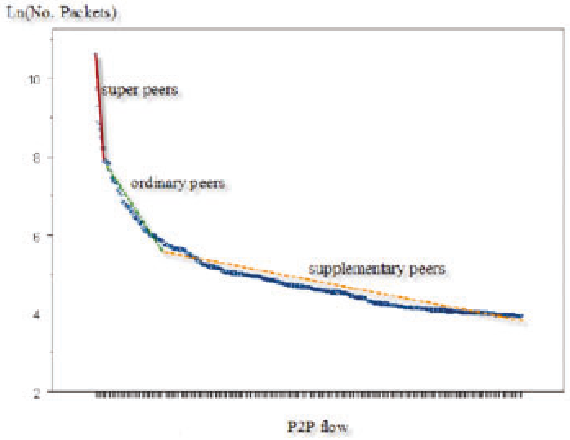

Fig. 8. Classification of peer relations based on flows feeding a home peer

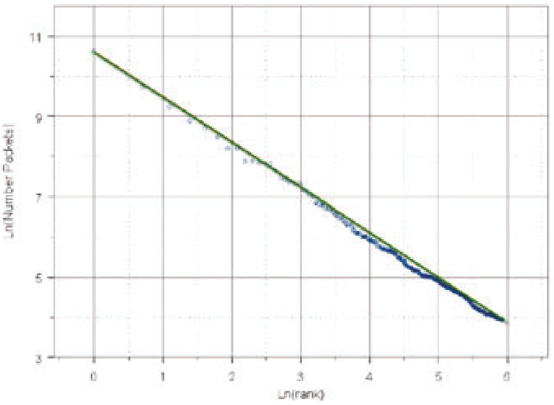

Fig. 9. Rank-frequency plot of flows feeding a home peer

a session, namely his individual *super peers*, *ordinary peers* and *supplementary peers* (see Fig. 8). The investigation of the transferred accumulated volumes among the peers of the flow graph G_V, in particular the number of packets and the byte volumes flowing inbound and outbound to a home peer reveal the hierarchical structure of the overlay. In our SopCast session example it is illustrated by a home peer p_1 at 10.59.1.106 whose inbound traffic is generated by $n = 375$ superimposed flows from distinct feeders p_i.

If we arrange the n flows according to the number of exchanged packets on a logarithmic scale, we can realize the hierarchy of the 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 (r.v.) 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 $a > -1$, $p \in (-1, 1]$, 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. [2], [16]). Selecting the parametrization $p = 1, a \geq 0, \alpha > 1$, we get the Zipf-Mandelbrot df $\mathbb{P}\{Y = k\} = C \cdot (a+k)^{-\alpha},\ k \in \mathbb{N}, C^{-1} = \Phi(1, a+1, \alpha)$ 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 $\phi(p_i, p_1)$ transferred from a peer p_i to the home peer p_1 on a log-log scale, we can identify in this case (2) a linear relation between $\ln f_i$ and $\ln i$ of the related ranks i of the feeders p_i, see Fig. 9. Thus, a pmf of Zipf type

(2) adequately describes the local hierarchical peer structure seen by the home peer in our SopCast session.

If we interpret the transferred number of packets of a flow $\phi(p_i, p_1)$ 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 P2P model of a session by a corresponding heavy-tailed physical model of Pareto type with a r.v. X and its sample $\{X_1, \ldots, X_n\}$ (cf. [11]). We denote the distribution function of this Pareto model associated with the Zipf ranking law by

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

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

The corresponding $1-p$ quantile function of this Pareto model $x_p = F^{-1}(x_0, p, \beta) = \left(\frac{p}{C}\right)^{-1/\beta} = x_0 \cdot p^{-1/\beta}$ can be used to define the local classes of the feeding peers. Taking, e.g., the 2.5% or 5% and 10% or 20% levels for p, their quantiles $x_{0.025}, x_{0.5}, x_{0.1}, x_{0.2}$ specify the break points of the local classes of super peers, ordinary and supplementary peers associated with a home peer.

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

$$\widehat{\beta} = \left[\frac{1}{n-1}\left(\sum_{k=1}^{n-1} \ln(\frac{x_i}{x_n})\right)\right]^{-1} = \left[\frac{1}{n-1}\sum_{k=1}^{n-1} \ln(x_i) - \ln(x_n)\right]^{-1}$$

or in terms of Newman's estimate [11] $\widehat{\alpha} = \frac{1}{n}\sum_{k=1}^{n} \ln(\frac{f_i}{f_{\min}})$ where $f_{\min} = f_n$ represents the minimal measured value (cf. [11]).

In our Sopcast example it follows $x_0 = 51$, $\beta = 1.2 > 1$ and, thus, the mean $\mathbb{E}(X) = \beta x_0/(\beta - 1) = 306$ exists. Based on this Pareto model (3) of the peer preference relation the median is determined by $x_{1/2} = 51 \cdot 2^{1/\beta} = 90.87$.

Applying the size-weighted df

$$W(x_p) = \frac{\int_{x_p}^{\infty} xf(x)dx}{\int_{x_0}^{\infty} xf(x)dx} = (\frac{x_p}{x_0})^{-(\beta-1)}, \tag{4}$$

size-count disparity issues can be studied. It means that the fraction of those packets sent by the most active part of the peers p_i is specified in terms of (4) by the ratio of the packet load of the fraction of flows exceeding level x_p to the total packet load. Hence, the most active $p \cdot 100$ % of the flows determine the fraction $W(x_p)$ of the sent packets by means of the $1-p$th quantile x_p, i.e. $1 - F(x_p) = p \in (0, 1)$, in terms of $W(x_p) = p^{(\beta-1)/\beta}$, related to the Lorenz curve (cf. [11]).

In our SopCast example we get $W(x_{0.5}) = 2^{-(\beta-1)/\beta} = 0.89$. It means that 89 % of the mean packet load are sent by the 188 most active flows. For the top 10 flows we get $p = 10/375 = 0.0266$ and $W(x_{0.0266}) = 0.0266^{0.2/1.2} = 0.546$, i.e more than half of the packets are sent by the top 10 flows. The top 20% of

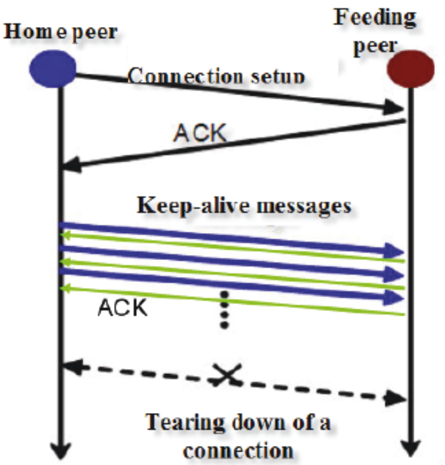

Fig. 10. SopCast connection setup

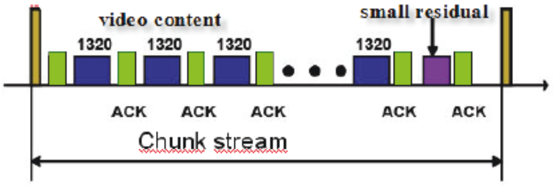

Fig. 11. Chunk stream of video content (see also [15])

the flows, i.e. 75 out of 375, achieve an activity ratio of $W_{x_{0.2}} = 0.76$ of the issued packets which confirms the 80/20 rule of heavy-tailed modeling (cf. [11]). In summary, we see that the first group with around 2.6 % of the flows provides a bit more than 54 % of the number of packets exchanged with the home peer and consists of less than 11 peers with high upload capacity, the second group of the top 20% of the flows excluding top 10 offers additionally 22 % of the packet load and the third group the residual less than 24 %. (see Fig. 8).

3.3 Teletraffic Analysis of a Dominant Packet Flow

One can study the UDP flows during a P2PTV session from the perspective of the initiating home peer p_1. The investigations confirm that first several control flows

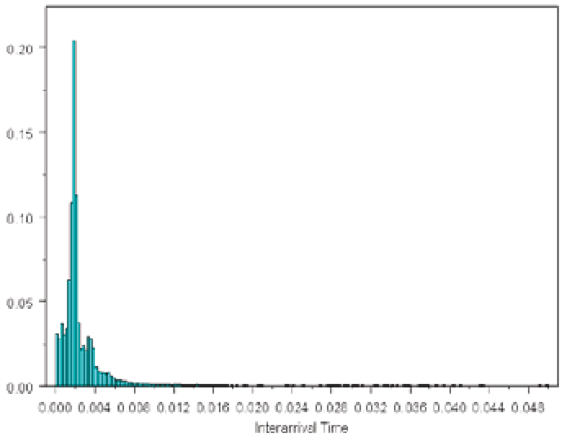

Fig. 12. Packet inter-arrival times arising from the downward flow of a dominant super peer

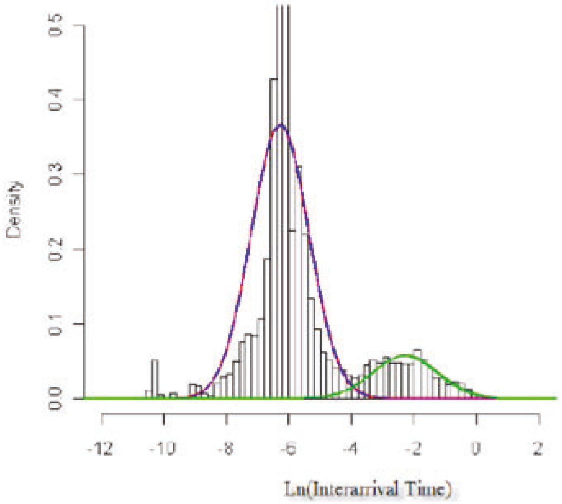

Fig. 13. Approximation of the packet inter-arrival times on a logarithmic scale

are instantiated for the initial opening of several connections to the ordinary and super peers indicated during registration in the downloaded peer list (see Fig. 10, cf. [15]). After the establishment of the neighborhood graph the communication on the flow graph is started implying the simultaneous existence of several flows $\phi(p_j, p_1)$. The latter are used to exchange video chunks between the monitored home peer p_1 and its neighbors p_j. Each flow consists of multiple consecutive packet sequences, i.e. micro-flows, and comprises a pair of a video chunk request on the uplink and the requested chunk packets on the downlink (cf. [4]). These pairs can be identified by a request packet with a frame length of 88 bytes which decomposes the incoming flow of a certain peer p_j into several chunk sequences (cf. [15]). They are answered by video contents with a maximal frame packet length of 1362 byte, i.e. 1320 bytes UDP payload, and acknowledged by p_1 with packets of 28 byte UDP payload, see Fig. 11. We suspect that the latter message applies a TCP semantics calling the next requested video block and acknowledges all correctly received previous ones. Exchanged between the feeding peers p_j and p_1, these flows create the superimposed inbound flow to the home peer p_1.

If we look into the downward flows and sort them according to their contribution we can recognize a dominant super peer at the IP address 150.237.180.164. The analysis of its packet inter-arrival times at the normal scale does not provide much structure whereas a logarithmic transformation can reveal its compositional character (see. Figs. 12, 13). If one disregards the small correlation up to lags 100 among the packets (see. Fig. 14), one can approximate the inter-arrival time of such a dominant downward flow by a recurrent stream whose packet inter-arrival time follows a mixture of two log-normal distributions (see. Fig. 13). It enables the development of appropriate simple teletraffic models of P2P traffic at the IP packet level by renewal streams with independently marked packet lengths that are derived from the histogram estimates of the IP datagram length with an 80% atom at 1348 bytes and several atoms of the control packets determined by a histogram estimate (see Fig. 7).

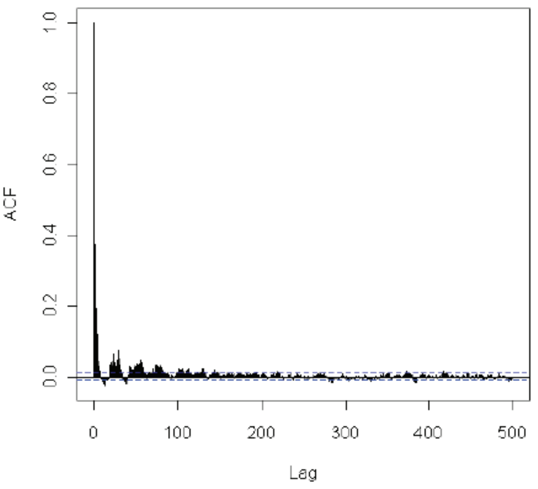

Fig. 14. ACF of packet inter-arrival times arising from the downward flow of a dominant super peer

4 Conclusions

Responding to the rapid deployment of multimedia applications provided by peer-to-peer overlay structures, we have presented a comprehensive concept integrating modeling, measurement, and teletraffic analysis of peer-to-peer traffic flows at the packet and session levels. A first version of this concept has been implemented in the JAVA tool Atheris on Windows and LINUX operating systems. It incorporates passive and active measurements at a single point in the overlay network. The data analysis focuses on the identification, data extraction and inspection of flows upward and downward to an observed home peer applying corresponding functions of CoralReef. We have illustrated the latter concept by traffic characteristics of individual streams issued by feeding peers in a SopCast overlay network.

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 at the packet and session levels as well as at the structural level of the overlay.

References

1. 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)
2. Aksenov, S.V., Savageau, M.A.: Some properties of the Lerch family of discrete distributions. University of Michigan, Ann Arbor (February 2008)
3. Ciullo, D., et al.: Understanding P2P-TV systems through real measurements. In: Proc. GLOBECOM 2008, pp. 2297–2302. IEEE Computer Society, Los Alamitos (2008)
4. Eittenberger, P., Krieger, U.R., Markovich, N.M.: Measurement and Analysis of Live-Streamed P2PTV Traffic. In: Czachórski, T. (ed.) Performance Modelling and Evaluation of Heterogeneous Networks, Proc. HET-NETs 2010, Zakopane, Poland, January 14-16 (2010)
5. Hei, X., Liang, C., Liang, J., Liu, Y., Ross, K.W.: A measurement study of a largescale P2P IPTV system. IEEE Tran. on Multimedia 9(8), 1672–1687 (2007)
6. Hu, N., Steenkiste, P.: Evaluation and Characterization of Available Bandwidth Probing Techniques. IEEE Journal on Selected Areas in Communications 21, 879–894 (2003)
7. Keys, K., et al.: The Architecture of CoralReef: An Internet Traffic Monitoring Software Suite. Coral Reef Documentation, CAIDA (2001), `http://www.caida.org/tools/measurement/coralreef/doc/doc/index.html`
8. Krieger, U.R., Schweßinger, R.: Analysis and Quality Assessment of Peer-to-Peer IPTV Systems. In: Proc. 12th Annual IEEE International Symposium on Consumer Electronics (ISCE2008), Algarve, Portugal, April 14-16 (2008)
9. 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)
10. 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 (submitted)
11. Newman, M.E.J.: Power laws, Pareto distributions and Zipf's law. Contemporary Physics 46, 323–351 (2005)
12. Peltotalo, J., et al.: Peer-to-peer streaming technology survey. In: Proc. of the Seventh International Conference on Networking ICN '08, pp. 342–350. IEEE Computer Society, Washington (2008)

13. 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)
14. Silverston, T., et al.: Traffic analysis of peer-to-peer IPTV communities. Computer Networks 53(4), 470–484 (2009)
15. Tang, S., Lu, Y., Hernández, J.M., Kuipers, F.A., Van Mieghem, P.: Topology dynamics in a P2PTV network. In: Fratta, L., Schulzrinne, H., Takahashi, Y., Spaniol, O. (eds.) IFIP-TC 6. LNCS, vol. 5550, pp. 326–337. Springer, Heidelberg (2009)
16. Zörnig, P., Altmann, G.: Unified representation of Zipf distributions. Computational Statistics & Data Analysis 19, 461–473 (1995)