

The Peer-to-Peer Trace Archive: Design and Comparative Trace Analysis

Boxun Zhang
TU Delft, the Netherlands
B.Zhang@tudelft.nl

Alexandru Iosup
TU Delft, the Netherlands
A.Iosup@tudelft.nl

Johan Pouwelse
TU Delft, the Netherlands
J.A.Pouwelse@tudelft.nl

Dick Epema
TU Delft, the Netherlands
D.H.J.Epema@tudelft.nl

Categories and Subject Descriptors

C.2.0 [Computer-Communication Networks]: General; C.2.3 [Network Operations]: Network Monitoring; H.2.8 [Databases]: Database Application; H.3.1 [Online Information Services]: Data sharing.

General Terms

Measurement, Performance, Reliability.

Keywords

peer-to-peer, data archiving, measurement.

1. INTRODUCTION

Peer-to-Peer (P2P) systems have gained a phenomenal popularity in the past few years; BitTorrent serves daily tens of millions of people and generates an important fraction of the Internet traffic. Measurement data collected from real P2P systems are fundamental for gaining solid knowledge of the usage patterns and the characteristics of these systems, and can improve the modeling, the design, and the evaluation of P2P systems. Although many P2P measurements have been carried out in the last decade [1, 2, 4, 6, 7], few measurement data sets are publicly available, and for these few the data are presented in different formats. This situation hampers the exchange, study, and reuse of existing traces. Furthermore, due to the lack of available datasets, many P2P studies have been based on unrealistic assumptions about the characteristics and usage patterns of P2P systems; as a consequence, many P2P algorithms and methods lack a realistic evaluation. To remedy this situation, we have set to create the P2P

Trace Archive (P2PTA), a virtual meeting place that facilitates the collection and exchange of P2P traces.

2. THE P2P TRACE ARCHIVE

P2PTA consists of a unified data format, many datasets, and associated tools; we describe each, in turn.

2.1 A Unified Trace Format

The P2PTA provides a unified data format for P2P traces. Our data format was designed [8] to reflect the structure of diverse P2P systems and to be extensible for future traces. The format can store data traced at four operational levels: interaction between the P2P application and the resources it uses (network packets, other I/O, etc.), individual peers, swarms (groups of peers that interact with each other to achieve a specific goal, for example sharing a movie), and communities (groups of weakly coordinated swarms, for example PirateBay). The format allows the specification of meta-information such as author, measurement period, etc. The format already allows storing both generic P2P system information, such as peer arrival and session length, and application-specific information, such as names of shared files for P2P file-sharing systems. The format further defines two data aggregation levels, trace and community dataset. A trace is the result of a single measurement collected from a P2P community. A community dataset is a set of traces collected from the same P2P community.

Researchers can add information fields to the format, but currently there is no community-validated entity to manage the sharing of added fields. We are looking for cooperation and help from the community, towards developing a standard format and its managing entity.

2.2 Archive Contents

Currently, the P2PTA hosts 23 traces corresponding to a variety of P2P systems and applications. Table 1 gives an overview of ten of these traces; for a full description of each trace present in the P2PTA we refer the reader to our technical report [9]. For P2P file-

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

ACM CoNEXT Student Workshop, November 30, Philadelphia, USA.

Copyright 2010 ACM 978-1-4503-0468-9/10/11 ...\$10.00.

ID	Trace description (type)	Period	Sampling	Files	Sessions
T1'03	SuprNova [6] (BitTorrent)	06 Dec 2003 to 17 Jan 2004	2.5 min	12	28,423,470
T2'05	ThePirateBay [3] (BitTorrent)	05-11 May 2005	2.5 min	4,800	35,881,338
T3'05	FileList.org (BitTorrent)	14 Dec 2005 until 4 Apr 2006	6 min	3,000	2,172,738
...					
T12'04	Gnutella [4]	Mar 19 2004 to Mar 28 2004	n/a	2,896,885	n/a
T13'03	eDonkey [2]	Oct 14 2003 to Oct 16 2003	n/a	1,282,420	n/a
T14'07	PPLive [7] (Streaming & VoD)	Jan 2007 (1 day)	10 min	n/a	67,051
T15'05	Skype [1] (VoIP)	Sep 12 2005 to Oct 12 2005	30 min	n/a	29,218

Table 1: Summary of the traces in the Peer-to-Peer Trace Archive.

sharing systems besides BitTorrent, T12 is a Gnutella trace [4], and T13'03 and T13'04 were collected [2, 5] from eDonkey. For other P2P applications, T14 was collected [7] from PP Live (video streaming), and T15 was collected [1] from Skype (phone over Internet).

The Archive includes a rich collection of community datasets collected from BitTorrent, a popular file-sharing system. From the community perspective, the BitTorrent traces focus on communities with either general or very specific types of content, and communities that are accessible either to everyone or that are closed (open to registered users). From the community size perspective, the traces have been collected from the largest communities in the world at the time of the data collection down to small communities, both in terms of the number of users and number of shared files. The BitTorrent traces were collected through various measurement techniques. In addition, we anonymize the sensitive information in these traces (see our technical report [9]).

2.3 Comparative Trace Analysis

The P2PTA includes various tools for automatic trace analysis. These tools can analyze the common features of P2P systems: peer arrival and departure, application upload and download bandwidth, etc. For P2P file-sharing, the existing analysis tools focus on content size and popularity, and on download completion ratios (also per session); for BitTorrent, the automated analysis includes the peer seeding and seeding-after-leaching times. Our tools can use data from entire community datasets and from single traces, and aggregate results for swarms and communities. The results of the analysis include basic statistics and fitting of well-known probability distributions to the data (see our report [9]).

The P2PTA analysis tools enable a wide variety of trace-based investigations; we have already started with the following: a study [9] of characteristics of different systems (by comparing different community datasets, or different traces belonging to the same dataset), a study [9] of the evolution of P2P systems (by comparing traces from one community dataset but collected in different years), a study of measurement techniques (by comparing traces within one community dataset but collected with different measurement techniques), etc. In near future we plan to provide tools for converting

traces into our unified data format.

3. RESEARCH ENABLER

The Archive paves the way for comparative studies of P2P systems, which allows researchers to consider various (types of) P2P systems and to capture their characteristics simultaneously.

The P2PTA can help complement the current model-based system design with a trace-based approach. In this way, the hidden patterns that exist in real traces of existing P2P systems will be implicitly used to improve the design, testing, and tuning of P2P systems.

4. AVAILABILITY

The Peer-to-Peer Trace Archive is available online at <http://p2pta.ewi.tudelft.nl>. We invite contributions of P2P traces and knowledge to the Archive.

5. REFERENCES

- [1] S. Guha et al. An experimental study of the Skype Peer-to-Peer VoIP system. In *IPTPS*, 2006.
- [2] S. Handurukande et al. Peer sharing behaviour in the eDonkey network, and implications for the design of server-less file sharing systems. In *EuroSys*, 2006.
- [3] A. Iosup et al. Correlating topology and path characteristics of overlay networks and the internet. In *GP2PC*, 2005.
- [4] A. Klemm et al. Characterizing the query behavior in peer-to-peer file sharing systems. In *IMC*, 2004.
- [5] F. Le Fessant et al. Clustering in peer-to-peer file sharing workloads. In *IPTPS*, 2004.
- [6] J. Pouwelse et al. The BitTorrent P2P file-sharing system: Measurements and analysis. In *IPTPS*, 2005.
- [7] L. Vu et al. Measurement and modeling a large-scale overlay for multimedia streaming. In *QShine*, 2007.
- [8] B. Zhang et al. A unified format for traces of peer-to-peer systems. In *LSAP*, pages 27–34, 2009.
- [9] B. Zhang et al. The peer-to-peer trace archive: Design and comparative trace analysis. Technical Report PDS-2010-003, TU Delft, April 2010. [Online] Available: pds.twi.tudelft.nl/reports/2010/PDS-2010-003.pdf.