

The logo for IEEE P2P'10 features a red stylized 'P' icon on the left, followed by the text 'IEEE P2P'10' in a bold, black, sans-serif font. Below this, the text 'IEEE INTERNATIONAL CONFERENCE ON PEER-TO-PEER COMPUTING' is written in a smaller, black, sans-serif font, and 'AUGUST 25-27, 2010 • DELFT, NETHERLANDS' is written in the same font at the bottom.

IEEE P2P'10
IEEE INTERNATIONAL CONFERENCE ON PEER-TO-PEER COMPUTING
AUGUST 25-27, 2010 • DELFT, NETHERLANDS




IEEE P2P'10 Conference


Amir H. Payberah
amir@sics.se

Papers and Topics Distribution

- Full papers: 27
- Industrial papers: 3
- Short papers: 7



Countries Distribution



BitTorrent

The Impact of Caching on BitTorrent-like Peer-to-Peer Systems

Frank Lehrieder, György Dán, Tobias Hoßfeld, Simon Oechsner, and Vlad Singeorzan

Best
Paper

- They study the **efficiency** of caches in term of decreasing the **inter-ISP traffic**?
- They develop a **fluid model** of the system dynamics of **BitTorrent**-like file-sharing systems that incorporates the effects of P2P caches.
- They show that the impact of caches **cannot be accurately assessed** without considering the effects of the caches on the system **dynamics**.
 - They may have **negative effects**: increase the amount of outgoing transit traffic of an ISP.



Can P2P-Users Benefit from Locality-Awareness?

Frank Lehrieder, Simon Oechsner, Tobias Hoßfeld, Zoran Despotovic, Wolfgang Kellerer, and Maximilian Michel

- They study the impact of **locality-awareness** mechanisms in P2P CDN from the **user's point of view**.
- A set of studies shows that the **locality** can lead to a **win-win** or **win-no lose** situations: **the ISPs benefit and the users do not suffer**.
- They show that in **real-life**, even a **win-no lose situation is difficult to achieve**.
 - They provide **mainly a gain for the ISPs**.
 - Some P2P users may benefit, some may lose.

Unravelling BitTorrent's File Unavailability: Measurements and Analysis

Sebastian Kaune, Rubén Cuevas Rumín, Gareth Tyson, Andreas Mauthe, Carmen Guerrero, and Ralf Steinmetz

- They study BitTorrent's **content unavailability** problem.
 - The available lifespan of most torrents is between **30-300 hours**.
- They perform measurement studies including **46K torrents** and **29M users**.
- In **86%** of cases, leechers are **unable** to reconstruct files in the **absence of seeders**, however, in 14% of cases, they can.
- In **64%** of torrents, unavailability is not **immutable**.
- **23.5%** of users affected by a lack of seeders actually being able to complete their downloads.



On Tracker Selection for Peer-to-Peer Traffic Locality

Haiyang Wang, Jiangchuan Liu, Bo Chen, Ke Xu, and Zhen Ma

- How ISPs can control the **neighbour selection** of most peers to reduce the inter-ISP traffic?
- **Modifying the Internet trackers**, but
 - **how many** trackers should be modified for a given ISP?
 - **which** trackers should be modified?
 - there is **no grantee** that the peers will always choose the modified trackers as we expected.
- They find that some trackers have very **similar peer distribution** that should be clustered and modified at the same time.
 - A **machine learning** based model to quantify the tracker similarity.



Can Realistic BitTorrent Experiments Be Performed on Clusters?

Ashwin Rao, Arnaud Legout, and Walid Dabbous

- In contrast to experiment testbeds such as Planetlab, the experiments performed on **dedicated clusters** are **reproducible**: **Are they realistic?**
- They study the impact of **network latency** and **packet loss** on the **download time** of a file using BitTorrent.
- They show that they have a **marginal impact** on the download time: **dedicated clusters can be safely used.**

Understanding Peer Exchange in BitTorrent Systems

Di Wu, Prithula Dhungel, Xiaojun Hei, Chao Zhang, and Keith W. Ross

- **Peer Exchange (PEX)**: peers directly **exchange** with each other lists of active peers in the torrent.
 - Used in modern bittorrent clients, e.g., **Vuze** and **uTorrent**.
- They study the **impact and properties** of BitTorrent PEX.
- They also show that PEX can increase the **download speed** of **40%** torrents. The average reduction of **download time** is about **7%**.



Measurements, Analysis and Modeling of Private Trackers

Xiaowei Chen, Yixin Jiang, and Xiaowen Chu

- Private Trackers (PT): much faster download speed.
- Ratio Enforcement (SRE): an **incentive mechanism** to overcome the free-riding issue.
 - PT banned a peer if its **upload-to-download** ratio is less than a threshold.
 - A peer with higher share ratio will be awarded.
- They compared system behaviours among 13 private trackers and 2 public trackers, i.e., **thePirateBay** and **TorrentPortal**, for over 6 months:
 - user viscosity, torrents evolution, user behaviours, and content distribution.



香港浸會大學
ONG KONG BAPTIST UNIVERSITY

Improving Accuracy and Coverage in an Internet-Deployed Reputation Mechanism

Rahim Delaviz, Nazareno Andrade, and Johan A. Pouwelse

- **BarterCast** is a **reputation mechanism** used by **Tribler** file-sharing client: to select good partners and to prevent free-riding.
- They propose **three modifications** to BarterCast and evaluate the reputation properties:
 - **Accuracy**: how well a peer can approximate objective reputation values when calculating the reputation of other peers.
 - **Coverage**: the fraction of peers for which an interested peer is able to compute reputation values.

Measurement and Analysis of BitTorrent Traffic in Mobile WiMAX Networks

Seungbae Kim, Xiaofei Wang, Hyunchul Kim, Taekyoung Kwon, and Yanghee Choi

Short paper

- They measure the traffic of BitTorrent service in **WiMAX** networks.
- They show links in mobile WiMAX are quite **unstable** that degrades the download performance.
- They also show that **control packets** are unnecessarily increased and waste bandwidth.



Streaming

QoE in Pull Based P2P-TV Systems: Overlay Topology Design Tradeoffs

Rossella Fortuna, Emilio Leonardi, Marco Mellia, Michela Meo, and Stefano Traverso



- They analyse the performance of **pull-based P2P live** systems in scenarios:
 - heterogeneous peer bandwidth
 - network latencies
 - properties of encoded video streams are accounted
- They propose a **latency/bandwidth-aware overlay** topology design strategy that partially **localizes** the traffic, while improving the user **QoE**.
- They show that by **prioritizing chunks** that contain more **valuable** information at the scheduler level, system performance can be slightly improved in overloaded conditions.



Spotify--Large Scale, Low Latency, P2P Music-on-Demand Streaming

Gunnar Kreitz and Fredrik Niemelä

- **Mesh** structure, nodes have fixed maximum degree.
- Play a random track
 - Request **first piece** from Spotify **servers**.
 - Meanwhile, search P2P for remainder.
 - Switch back and forth between Spotify servers and peers as needed.
 - Towards end of a track, start pre-fetching next one.
- Finding peers:
 - Partial central index (**Napster-style**).
 - Broadcast query in small neighborhood (**Gnutella-style**).
 - Limited broadcast for local (LAN) peer discovery.
- **Cache** plays an important role in Spotify.



Multi-Source Scheduling in Streaming Erasure-Coded Video over P2P Networks

Man Lok Ma and Jack Y. B. Lee

- They propose a **push-based multi-source** streaming model that does not need the network resource estimation, e.g., upload bandwidth.
- They use of **erasure codes** to encode the media.
- The peers transmit data at the maximum rate allowed, and the client wait until sufficient amount of encoded media data are received, and then decode the data to obtain the original media data for playback.
- They decrease the number of **duplicate blocks** by introducing **Disjoint-Prefix randomized scheduler**.



NAP: An Agent-based Scheme on Reducing Churn-Induced Delays for P2P Live Streaming

Fei Huang, Binoy Ravindran, and Maleq Khan

- **Churn-induced delay:** delay from channel switch (**channel-churn**) and streaming recovery (**peer-churn**).
- They present **NAP** as a solution to reduce the churn-induced delay by **facilitating the bootstrapping process**.
- **Channel-churn:** peers proactively starts bootstrapping in **other channels**, retrieving list of peers in the channel while viewing current channel.
- **Peer-churn:** peers proactively bootstraps to the overlay of **current channel**.

Boosting Gossip for Live Streaming

Davide Frey, Rachid Guerraoui, Anne-Marie Kermarrec, and Maxime Monod

- They show that **gossip alone is unable** to offer satisfactory performance in the context of video streaming applications.
- They present **gossip++** as a reliable gossip-based streaming systems for real world environments.
- Gossip++:
 - **Codec**: adds redundant coded chunks to the stream (**erasure coding**).
 - **Claim2**: allows nodes to **re-request missing content** by re-contacting the nodes from which they received advertisements for the corresponding chunks. Instead of stubbornly requesting the same sender, the requesting node re-requests nodes in the set of proposing nodes in a **round-robin manner**.

Do BitTorrent-like VoD Systems Scale under Flash-Crowds?

Lucia D'Acunto, Tamas Vinko, and Johan Pouwelse

Short paper

- They study the **scalability** of BitTorrent-like VoD systems under flash-crowds.
- They show that their scalability is constrained by:
 - The **initial service capacity**.
 - The efficiency of **piece exchange** of the underlying P2P protocol.

Security

Cryptographically Enforced Permissions for Fully Decentralized File Systems

Bernhard Amann and Thomas Fuhrmann

- They present a **P2P file system** that provides **unix-like file access** to its users.
- They extend their system, **IgorFS**, that is unaware of users and groups.
 - IgorFs is based on **Igor** DHT.
- All objects in the IgorFS are cut into variably sized chunks, encrypted and inserted into the DHT.
- They propose:
 - An **integrity verification** algorithm checks the validity of the current file system state.
 - A **cryptographic data protection** scheme preserves the privacy of the file system content.

Distributed Access Enforcement in P2P Networks: When Privacy Comes Into Play

Marc Sánchez-Artigas

- They present protocols to **delegate access control to intermediaries** in such a way that **intermediates do not learn the privileges** and the **requesters do not learn the access policy**.
- Example:
 - **Alice** is a requester and has private input **a**.
 - **Bob** is another user who has a policy **P** distributed among a set of **w** delegates **S1..Sw**.
 - To gain access to a resource owned by Bob, Alice needs the knowledge of a secret value **s** specified by Bob.

$$F_{S_\ell}(\mathcal{P}) = \perp, \quad F_{Alice}(a) = \begin{cases} s & \text{if } a \text{ satisfies } \mathcal{P}; \\ \perp & \text{otherwise.} \end{cases}$$

On the Privacy of Peer-Assisted Distribution of Security Patches

Di Wu, Cong Tang, Prithula Dhungel, Nitesh Saxena, and Keith W. Ross

- **Peer-assisted patch distribution:** when a peer A requests a patch from another peer B, it announces to B its vulnerability, which B can **exploit instead of providing the patch**.
- They propose two solution for this problem:
 - **Honeypots:** discover malicious peers.
 - **Anonymizing network:** hiding the id of vulnerable hosts from malicious hosts.
 - It is based upon **onion routing**:
 - Routing path consists of a series of **relay nodes**.
 - The **first relay node** learns the source of the message but not the message itself.
 - The **last relay node** learns the message and its destination but not the source.
 - Each **intermediary relay** node only learns the id of the relay node from which it received and which the packet is to be forwarded to.



Spads: Publisher Anonymization for DHT Storage

Pascal Felber, Martin Rajman, Etienne Rivière, Valerio Schiavoni, and José Valerio

- **Publisher anonymization:** it allows users to send their sensitive data, but at the same time hides her identity from the others in the system.
 - e.g., crash report in web browsers
- **Rate limitation:** limits the data inserted by users into the system.
- They present **SPADS**, a publisher anonymizing and rate-limiting publication interface for untrustworthy clients willing to send privacy-sensitive data to a DHT.
- They used combination of **cryptography techniques** and **random path selection** (as used in **onion routing**).

Towards Plugging Privacy Leaks in the Domain Name System

Yanbin Lu and Gene Tsudik

- **DNS privacy leak:** Each DNS query generated reveals the origin and the target of that query.
 - Users' communication patterns might become exposed.
- They propose Privacy-Preserving DNS (**PPDNS**) based on a **DHT**. It uses the DHT index structure to provide name resolution query privacy.
- PPDNS also uses computational private information retrieval (**cPIR**) technique that helps clients **reduce** their communication **overhead**.



Verifiable Encryption for P2P Block Exchange

Gertjan Halkes and Johan Pouwelse

Short paper

- **Large View Exploit (LVE)**: a peer connects to as many other peers as possible to increase the chance to get free data.
 - Tit-for-tat does not work here.
- A solution to prevent LVE is to use **encryption** to send blocks, and provide the key for decryption after proof of upload by the remote peer.
 - Problem: lack of verifying the received blocks.
- They present an encryption method that does allow **(partial) verification** of the encrypted data.

Storage

Online Data Backup: a Peer-Assisted Approach

László Toka, Matteo Dell'Amico, and Pietro Michiardi

- They study the benefits of a **peer-assisted** approach to online backup applications, e.g., **dropbox** and **ubuntu one**.
 - Spare bandwidth and storage space of end-hosts
- They showed that, by using **adequate bandwidth allocation** policies in which storage space at a cloud provider is only used **temporarily**, a peer-assisted backup application can achieve performance comparable to client-server architectures at a fraction of the costs.
- They also explore the impact of **data placement** policies on system performance and fairness (**random vs tit-for-tat**):
 - The impact of system-wide fairness on performance is **negligible**.



PeerDedupe: Insights into the Peer-assisted Sampling Deduplication

Yuanjian Xing, Zhenhua Li, and Yafei Dai

- Server-side **deduplication**: splits a backup data stream into chunks and removes duplicate chunks (to improve storage utilization).
- **PeerDedupe**: a peer-assisted sampling for **inter-peer deduplication**.
- A peer sends the IDs of his **locally new chunks** to only a small number of other peers, called **MVH**, for further inter-peer deduplication.
 - Uses **probabilistic estimation** algorithm based on the **MinHash** sampling method to select MVHs.
 - The **MinHash** is a compressed representation of a set from which one can approximate the **resemblance of two sets**.
- PeerDedupe can remove over **98%** inter-peer duplication with each peer coordinating with no more than **5** MVHs.



Optimizing Near-Duplicate Detection for P2P Networks

Odysseas Papapetrou, Sukriti Ramesh, Stefan Siersdorfer, and Wolfgang Nejdl

- **Near Duplicate Detection (NDD)**: searching for **very similar files** over large file repositories.
 - Different recordings of the same movie.
- Most of current solutions employ a family of algorithms, called **Locality Sensitive Hashing (LSH)**, to map resources to bit strings, and to build an index suitable for answering **K-Nearest Neighbour (KNN)** queries.
- They present **POND**, an algorithm for NDD that **tunes the LSH parameter** to minimize the network usage.

Availability and Redundancy in Harmony: Measuring Retrieval Times P2P Storage Systems

Lluís Pamies-Juarez, Pedro García-López, and Marc Sánchez-Artigas

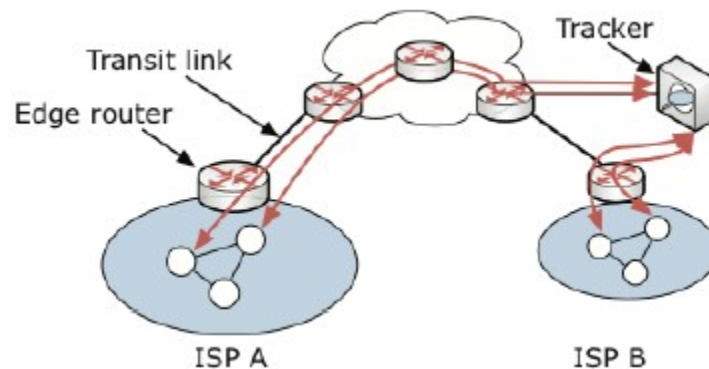
- Data availability → Redundancy → Storage costs and communication.
- They present an **analytical framework** to measure object **retrieval times** under different **redundancy** and **churn** circumstances.

Peer Discovery

IMP: ISP-Managed P2P

Shakir James and Patrick Crowley

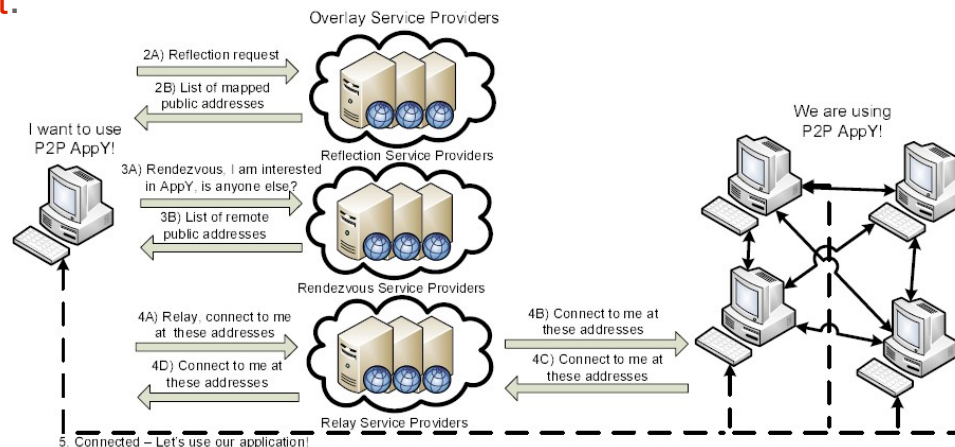
- They present **IMP**, a **transparent peer discovery** service that returns peers favourable to ISPs: reducing cross-ISP traffic.
- IMP does not require any modification in the current client application.
- For example, ISP A installs IMP alongside its edge router to reduce cross-ISP. IMP identifies and redirects tracker requests to its ISP Oracle component. The ISP Oracle responds with **a set of network-local peers** and **one non-local peer**, instead of a random subset from the tracker.



Addressing the P2P Bootstrap Problem for Small Overlay Networks

David Isaac Wolinsky, Pierre St. Juste, P. Oscar Boykin, and Renato Figueiredo

- They study the bootstrap problem for **small-scale** and **private** P2P networks, where there is **not dedicated bootstrap server**.
- They show that these peers can use the **existing public overlays**, if they support:
 - **Reflection**: provides a peer with a globally-addressable identifier,.
 - **Rendezvous**: A method for identifying peers interested.
 - **Relaying**: NAT traversal services.
- **XMPP** and **Brunet**.




Misc.

Hivory: Range Queries on Hierarchical Voronoi Overlays

Matteo Mordacchini, Laura Ricci, Luca Ferrucci, Michele Albano, and Ranieri Baraglia

- The current DHT solutions are not suitable for handling **multi-attribute range queries**, or if they are, they have a high maintenance, high replication costs or need a high number of messages to solve range queries.
- They present **Hivory**, a P2P support for multidimensional range query based on a hierarchy of **Voronoi** structure.



ASAP Top-k Query Processing in Unstructured P2P Systems

William Kokou Dédzoé, Philippe Lamarre, Reza Akbarinia, and Patrick Valduriez

- **Top-k query processing:** techniques are useful in unstructured P2P systems, to avoid overwhelming users with too many results.
 - **Problem: long waiting times**, because results are returned only when all queried peers have finished processing the query.
- They propose **ASAP** to deal with this problem.
- It uses a threshold-based scheme that considers the **score and the rank of intermediate results** to return quickly the high quality results to users.

Local Access to Sparse and Large Global Information in P2P Networks: a Case for Compressive Sensing

Rossano Gaeta, Marco Grangetto, and Matteo Sereno

- **Compressive Sensing (CS)**: is a technique for finding sparse solutions to underdetermined linear systems.
- They present a solution to give peers **local** access to **global** large and **sparse information** at a given rate.
- They use **CS** and **Random Walk (RW)**.
 - **CS**: it allows one to collect and compress the information in a distributed fashion.
 - **RW**: it distributes this compressed information with a controlled communication overhead.

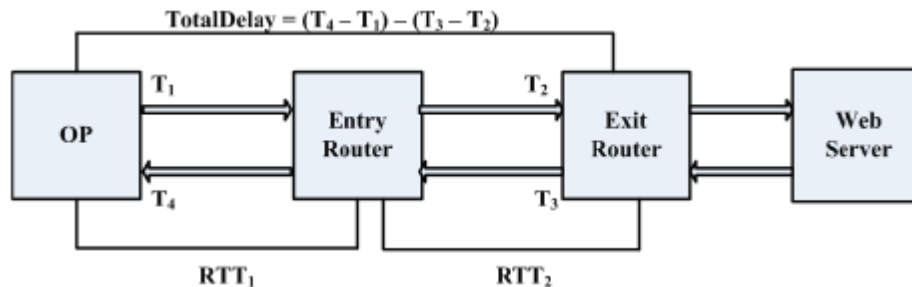


Waiting for Anonymity: Understanding Delays in the Tor Overlay

Prithula Dhungel, Moritz Steiner, Ivica Rimać, Volker Hilt, and Keith W. Ross

Short paper

- Tor: the overlay for providing **anonymity** services.
 - Problem: **high delay**.
- They study the **delays** in Tor.
- They show that the Tor's **routers delay** and the **overlay latency** both play role in delays in Tor.




Alcatel-Lucent
NYU·poly
POLYTECHNIC INSTITUTE OF NYU

Autonomous NAT Traversal

Andreas Müller, Nathan Evans, Christian Grothoff, and Samy Kamkar

Short paper

- Autonomous NAT traversal: NAT traversal **without a third party**.



Analyzing the DC File Sharing Network

Pavel Gurvich, Noam Koenigstein, and Yuval Shavit

Short paper

- **Direct Connect (DC):** P2P file sharing that clients connect to a central **hub** and can download files directly from one another.
 - Npster-like system
- They study the DC network characteristics:
 - distribution of users in hubs
 - hubs geography
 - queries distribution
 - trends in shared folder size



Spring for Vivaldi--Orchestrating Hierarchical Network Coordinates

Benedikt Elser, Andreas Förschler, and Thomas Fuhrmann

Short paper

- **Vivaldi**: a solution that uses nodes' coordinate to help any two nodes to estimate the network latency between them.
- They propose **hierarchical Vivaldi** that optimizes Vivaldi's peer selection process.



Universität
Karlsruhe (TH)





???