GnuNet Brahms: random node sampling

Alexander Galtsev and Ilya Migal
Technical University of Munich

Course: Peer-2-Peer Systems & Security

June 24, 2014

Overview

- Problem
- The project
- Other peer selection mechanisms
- Conclusion

RPS Problem

- In the gossip protocols there is a need for randomness in peer selection.
- Bad people might attempt to influence your choice, reducing the randomness.
- Byzantine fault tolerant random peer sampling is needed to prevent that.
- Brahms is byzantine fault tolerant, but hasn't been implemented yet.

System

- Each peer has a view V with a set of IDs it knows.
- V is asymptotically smaller than the system size.
- Adversary controls fraction f of the nodes.
- System uses both PUSH and PULL gossip, to prevent star topologies.

Brahms PUSH Defense

- Limit rate at which nodes can PUSH.
- If more PUSHes are received than expected in a given time interval, ignore all of them.

Brahms PULL Defense

- Control contribution of PUSH IDs ($\alpha|V|$)
- Control contribution of PULL IDs ($\beta|V|$)
- Use history samples ($\gamma|V|$)
- $\alpha + \beta + \gamma = 1$.
- If history contains non-faulty nodes, the attacker failed.

Other peer selection mechanisms

- ALTO protocol
- P4P
- BitTorrent
- Biased Neighbors Selection
- many others

ALTO Protocol

- The Application Layer Traffic Optimization.
- The ALTO architecture:
  - ALTO Server that responds to queries from ALTO Clients
  - ALTO Service Discovery entity used to discover the location of the server.

ALTO Architecture [10]

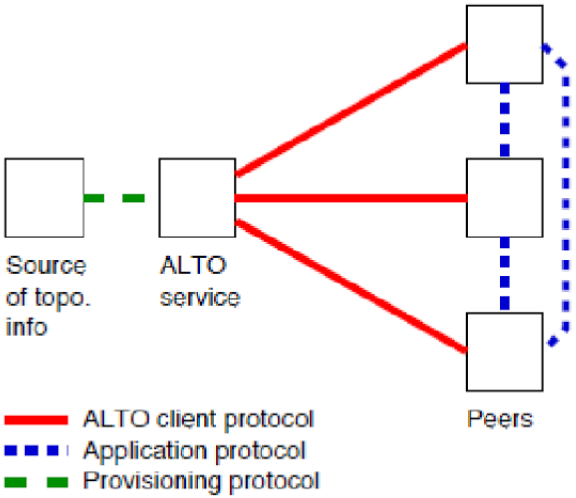

(a) Application without tracker

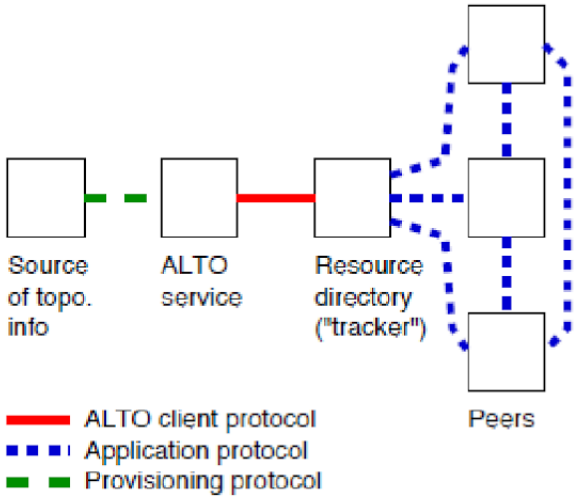

(b) Application with tracker

ALTO Server

- The ALTO Server, is operated by the Service Provider and provides network related information to it's clients.
- The server has *network map* and *cost map.*
- Network map groups hosts together, and assigns identifiers.
- Cost map assigns path costs to connections between networks in the network map.
- A peer-to-peer client can use the information provided by the service of the ALTO server to determine which of the other known peers are good candidates to choose as neighbors.

ALTO Architecture for P2P [10]

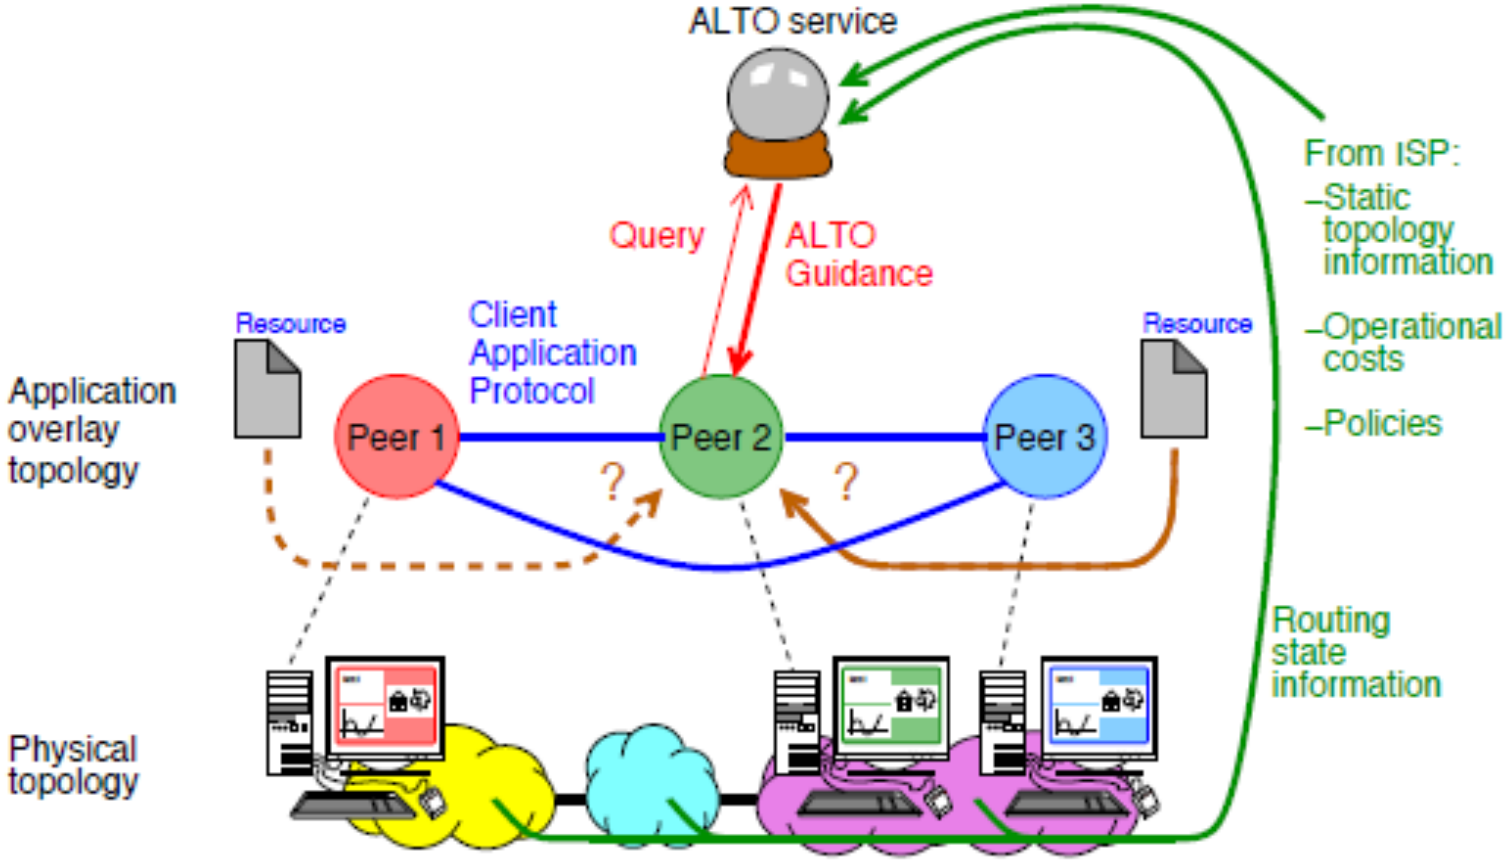

P4P framework

- *Proactive network Provider Participation for P2P* (P4P) is an architecture that allows ISPs to make information about the network available to users and application providers.

- The central component of the framework is the *iTracker*, a service that is provided by the ISP.

P4P framework

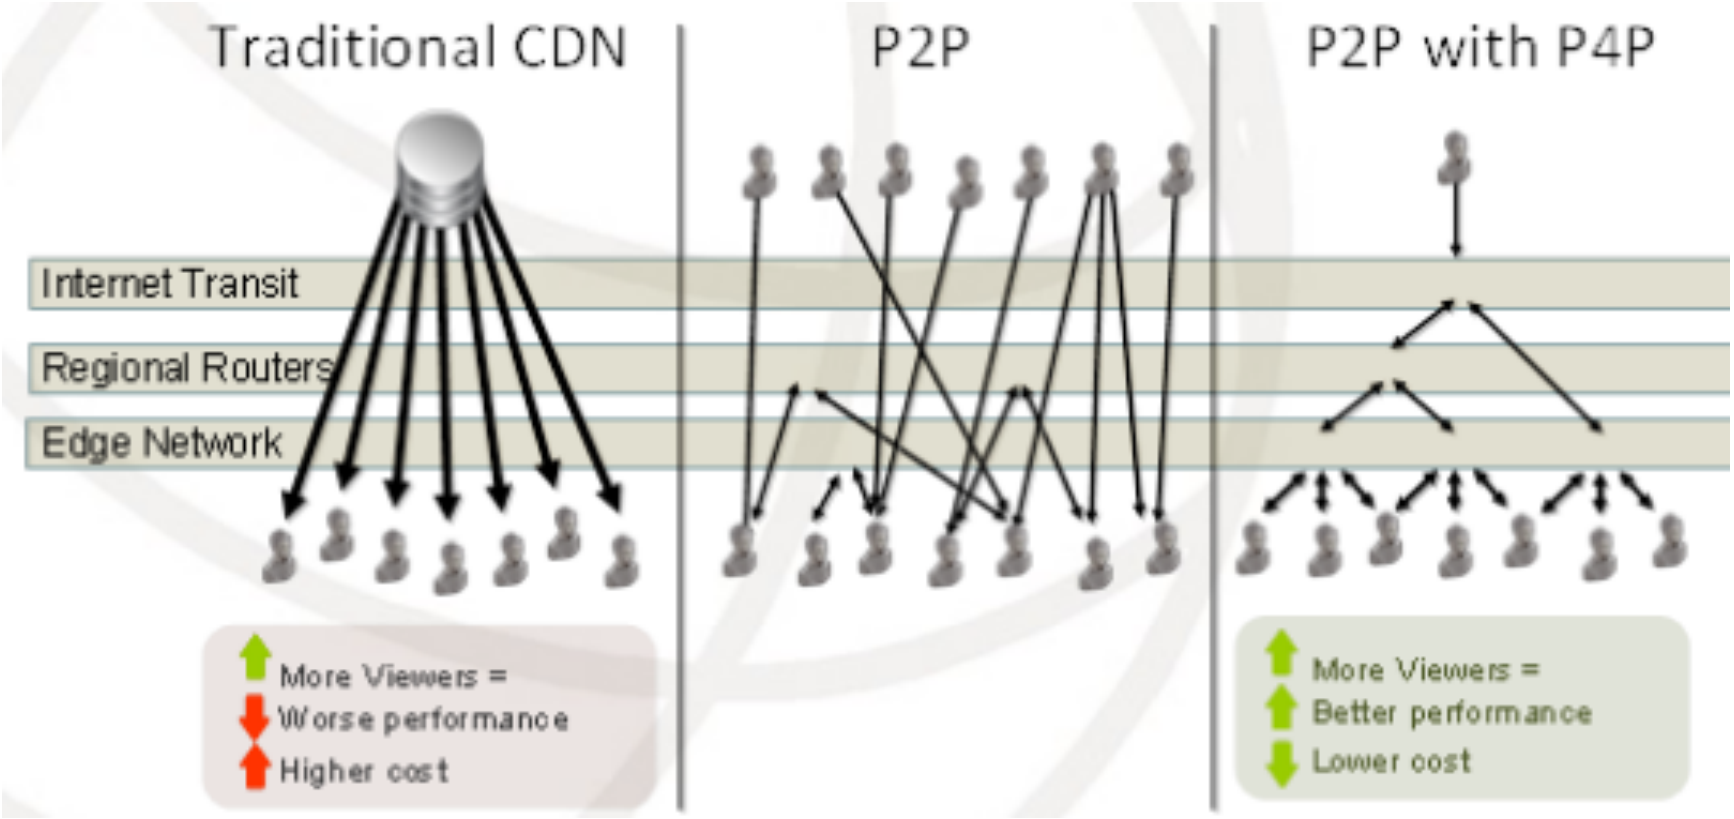

iTracker's interfaces

- *policy*: this interface gives hints about the preferences of the ISP with respect to the usage of links to other networks. The policy can be dynamic, e.g. it can depend on the congestion in certain parts of the network.

iTracker's interfaces

- *p4p-distance*: the users or applications can query this interface to compare the costs and distance to several peers as seen from the iTracker. This information can be used by peers or trackers in the peer selection process.

iTracker's interfaces

- *capability*: this interface can be used by ISPs to publish information about services they provide to support the P2P systems. These services can be used by users or application providers to accelerate the functions of the P2P system.

Entities and interactions between P4P components

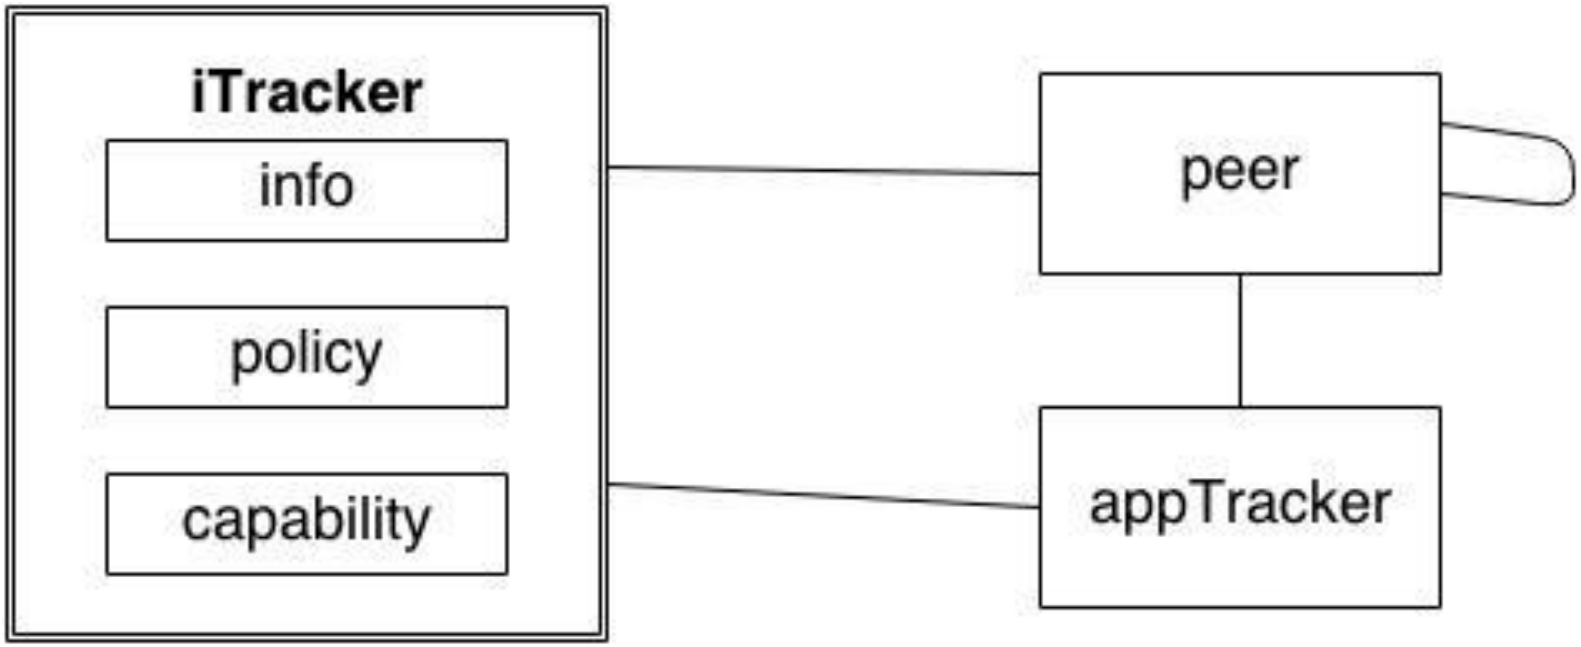

BitTorrent [6]

- The BitTorrent = many peers and one or more central trackers.
- Peers (the swarm) can be subdivided into seeders and leechers.
- Seeders are only uploading
- Leechers are both uploading and downloading.
- Peer selection - through tracker, DHT, PEX

BitTorrent

- New peer appears - looks for neighbors.
- Request the parts they do not have yet.
- Choking algorithm

BitTorrent peer selection [6]

- Peer selection situations in BitTorrent:
  - The tracker selects randomly a list of peers to send to a new peer;
  - The peer selects randomly a neighbor set from all known peers;
  - The peer selects peers to be regular unchoked (based on their performance);
  - The peer selects randomly an other peer to be optimistic unchoked.

Biased Neighbors Selection [6]

- In BitTorrent the tracker provides peers with a set of neighbors
- This neighbor set is a randomly chosen subset of all connected peers interested in this torrent.
- Biased neighbor selection (BNS)
- Internal peers and the external peers

Random Sampling difference [6]

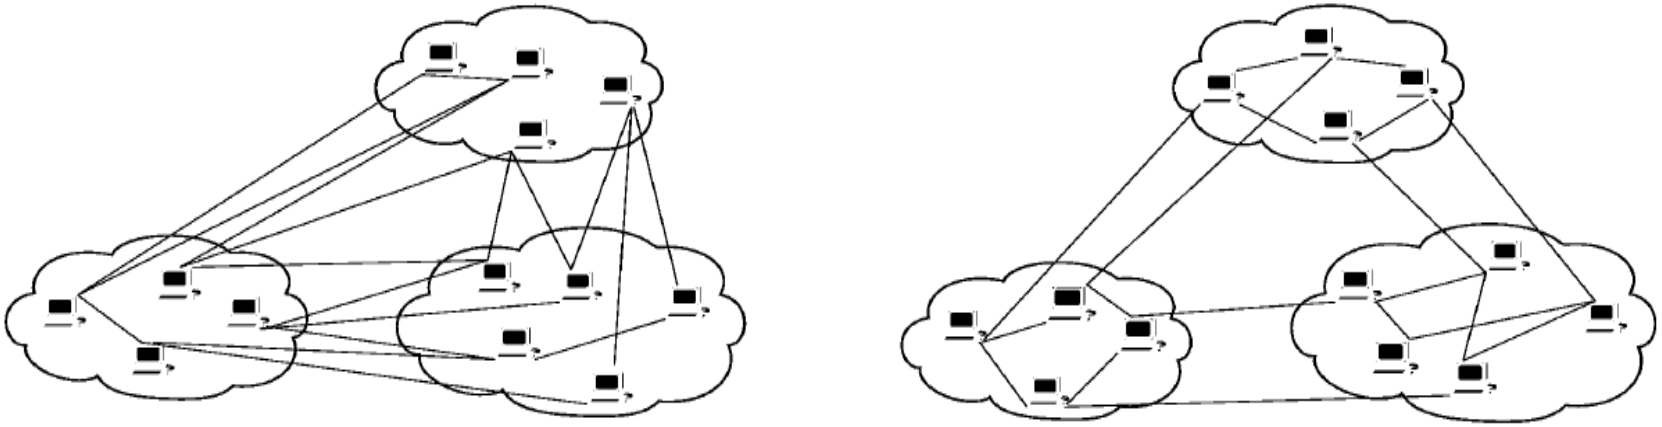

Uniform random neighbor selection (left) vs. biased neighbor selection (right).

Problem

- Each of those protocols' random peer sampling fails at one of two things: either they don't expect churn, or they don't expect byzantine attacks.
- Brahms assumes both of those problems.

Architecture

- Brahms will be a java module started with the gnunet-arm –s command.
- Ideally should have a DB for storing samples.
- Connects to nodes and exchanges the samples via GnuNet Transport subsystems.
- Check the info using Peer Info
- Pass on to Brams module where magic happens

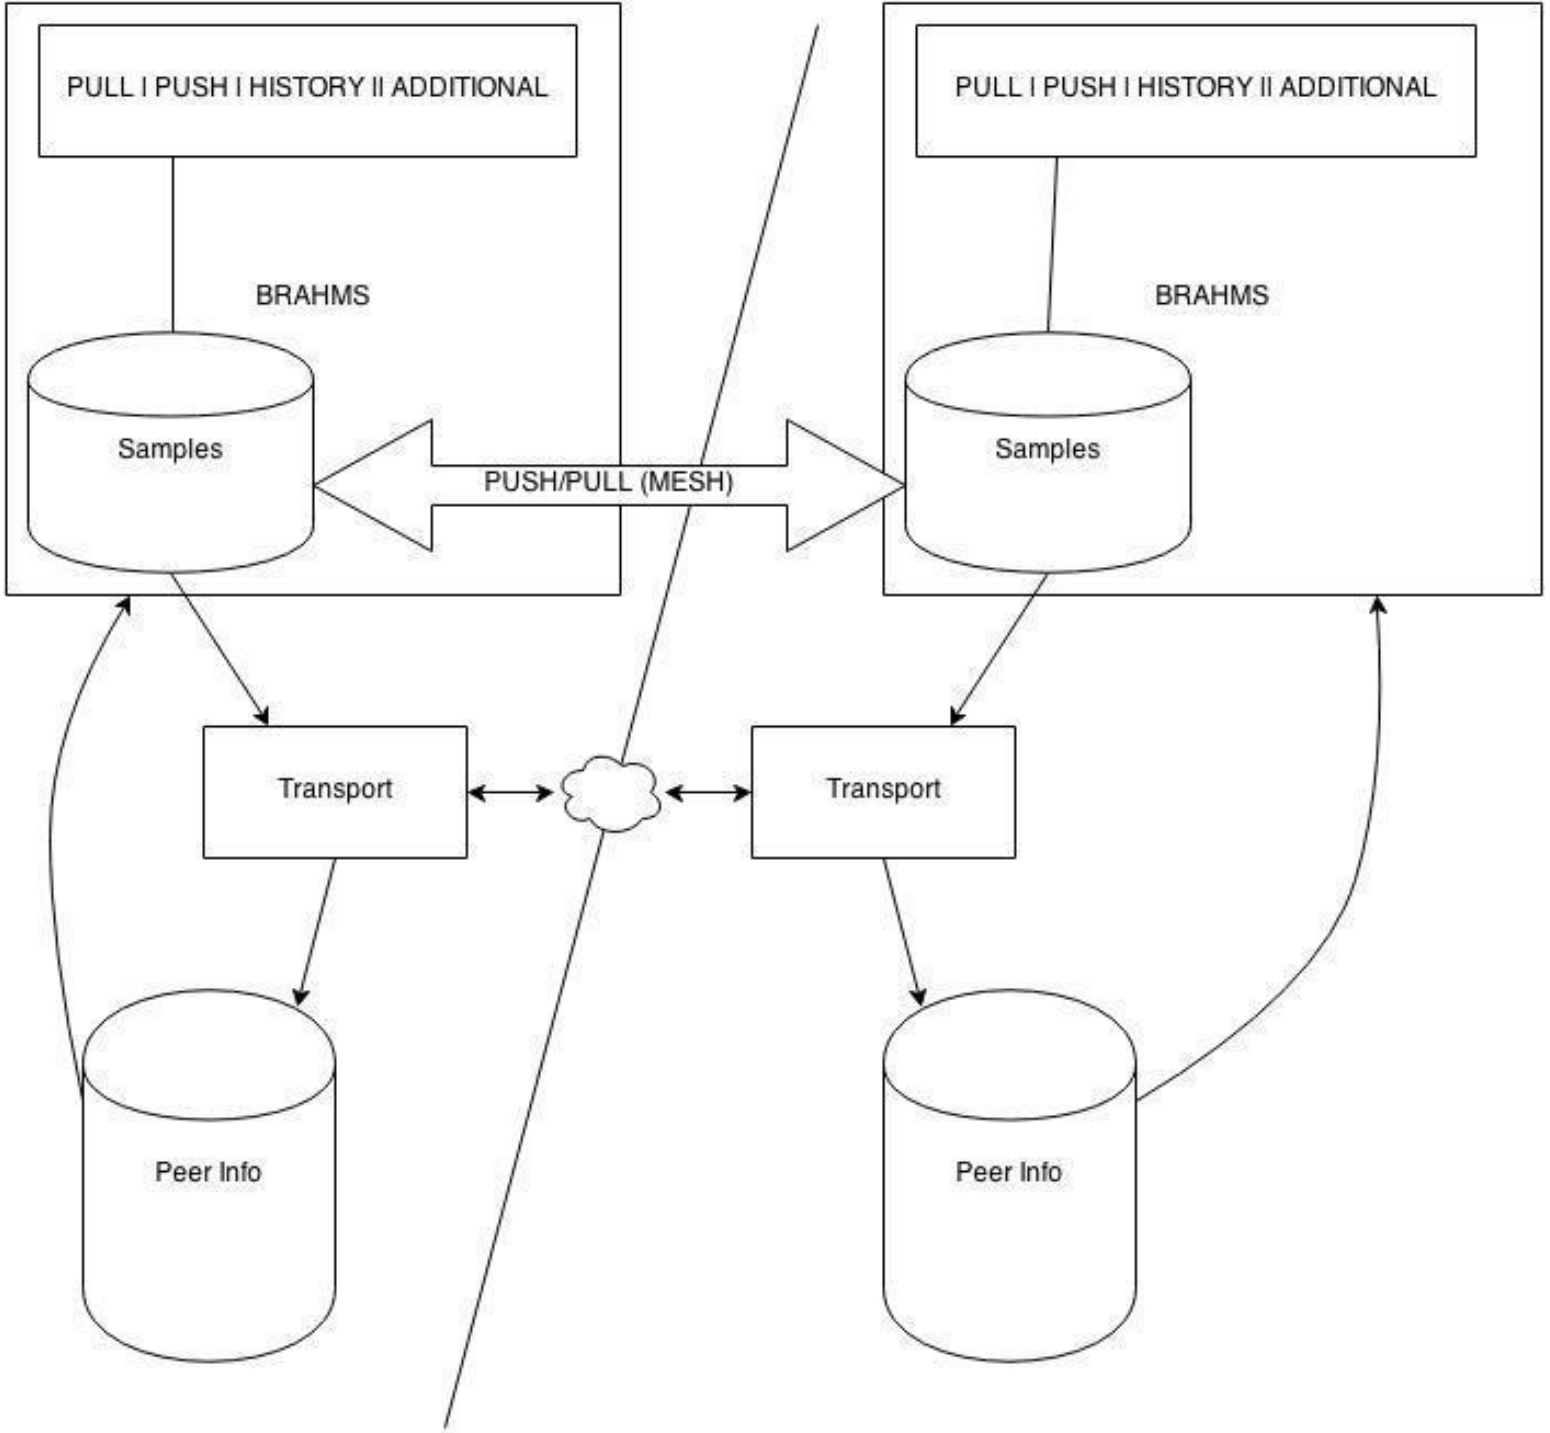
PULL I PUSH I HISTORY II ADDITIONAL
BRAHMS
Samples
PUSH/PULL (MESH)
PULL I PUSH I HISTORY II ADDITIONAL
BRAHMS
Samples
Transport
Transport
Peer Info
Peer Info

Milestones and Development Plan

- Create the packet template for sampling
- Create and Test message exchange between peers
- Create a sample storage mechanism (encrypted txt, sql, etc.)
- Create a documentation describing the implementation and the usage

Thank you!

- Comments?
- Questions?

References

[1] 2008, Bortnikov, Edward and Gurevich, Maxim and Keidar, Idit and Kliot, Gabriel and Shraer, Alexander (view online)

Brahms: Byzantine Resilient Random Membership Sampling

[2] 2009, Gian Paolo Jesi and Alberto Montresor

Secure Peer Sampling Service: The Mosquito Attack

18th IEEE International Workshops on Enabling Technologies: Infrastructures for Collaborative Enterprises, WETICE 2009, Groningen, The Netherlands, 29 June - 1 July 2009, Proceedings

[3] 2012, van de Bovenkamp, Ruud and Kuipers, Fernando and Van Mieghem, Piet (view online)

Gossip-Based Counting in Dynamic Networks

References

[4] 2007, Jelasity, Mark and Voulgaris, Spyros and Guerraoui, Rachid and Kermarrec, Anne-Marie and van Steen, Maarten (view online)

Gossip-based Peer Sampling

[5] 2010, Kermarrec, Anne-Marie and Leroy, Vincent and Thraves, Christopher (view online)

Ensuring Uniformity in Random Peer Sampling Services

[6] 2011, Menasche, Daniel Sadoc and de A. Rocha, Antonio A. and de Souza e Silva, Edmundo A. and Towsley, Don and Meri Le\"ao, Rosa M. (view online)

Implications of Peer Selection Strategies by Publishers on the Performance of P2P Swarming Systems

References

[6] 2004, Márk Jelasity and Rachid Guerraoui and Anne-marie Kermarrec and Maarten Van Steen

The Peer Sampling Service: Experimental Evaluation of Unstructured Gossip-Based Implementations

[7] 2004, Haase, Peter and Siebes, Ronny and van Harmelen, Frank (view online)

Peer Selection in Peer-to-Peer Networks with Semantic Topologies.

[8] 2010, Baldoni, Roberto and Platania, Marco and Querzoni, Leonardo and Scipioni, Sirio (view online)

Practical Uniform Peer Sampling under Churn.

References

[9] 2007, Vivek Vishnumurthy and Paul Francis

A comparison of structured and unstructured P2P approaches to heterogeneous random peer selection

[9] 2012, Oguz, Barlas and Anantharam, Venkat and Norros, Ilkka (view online)

Stable, distributed P2P protocols based on random peer sampling.

[10] 2010, J. Seedorf, S. Niccolini, M. Stiemerling, E. Ferranti, and R. Winter.

Quantifying operational cost-savings through ALTO-guidance for P2P live streaming.