

P4P – Peer for Peer

Introduction

- P2P applications establish connections with other peers without the knowledge of the underlying network resources
- P2P applications have not the ability to retrieve useful information about the network from providers
- This leads to inefficient use of resources
- P4P is a set of techniques and architectures aiming at exploiting the knowledge about the networks by clients

Introduction

- P2P traffic can cause traffic to unnecessarily follow multiple links within a provider's network
- Moreover, a large (unnecessary) amount of inter-domain traffic is possible

P4P

- P2P systems can try considering locality of peers to optimize resource utilization
- Current estimations of peer locality are inaccurate
- P4P aims at cooperation between P2P networks and underlying networks
- This cooperation requires explicit communication between P2P networks and providers
- Benefits for *providers*:
 - Fair and efficient usage of network resources
 - Cost reduction
- Benefits for *clients*:
 - Achieve the best possible application performance

P4P Principles

- Proactive network Provider Participation for P2P (P4P) can be achieved by intervening on:
 - The control level
 - Cooperation between P2P and providers for traffic control
 - The data level
 - Routers give fine-grained feedback to P2P
 - Routers can mark packets (e.g. TCP headers or a field in a P2P header)
 - P2P packets must be recognizable, implicit cooperation is required
 - Clients then adjust their flow rates accordingly

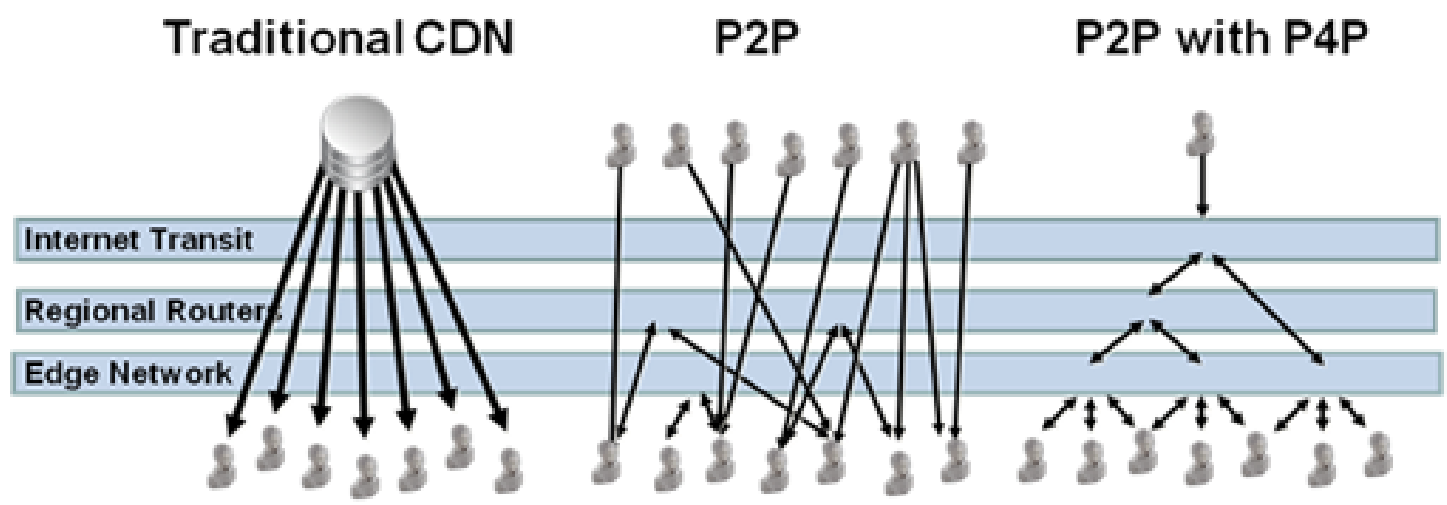
Other requirements

- P4P solutions should also offer:
 - Scalability when a large number of peers is involved
 - Extensibility: different P2P protocols should be supported
- Provider participation:
 - a network provider could promote the use of P4P techniques by also participating in the file exchange, offering high bandwidth servers for file “caching”

P4P Principles

- Modern P2P systems allow retrieving parts of the same file from multiple sources
- Source selection is typically performed with a random approach
- The optimized selection of sources can lead to both a successful download and an efficient use of network resources
- P2P clients can use network information for identifying the data routes the network provider prefers and connections to avoid in a given time period
- P2P clients can then cooperate by connecting to closer peers (or cheaper for the network provider)

P4P Principles



P4P Principles – Architectures

- Local peers can be identified in different ways:
- Client-Based. Each peer retrieves network information from a P4P server and autonomously selects peers for file exchange
 - P2P clients can independently exploit the information, depending on their internal policies
 - Clients are free to use provider information
 - P4P servers have no full control on the use of the provided information

P4P Principles – Architectures

- P2P System-Based:
- P2P servers retrieve network information from a P4P server, and use this information for forging appropriate responses to querying peers
 - Peers' selection policies are implemented by P2P servers / P2P system
 - P2P clients are forced to use the provided sub-set of peers
 - When the index is distributed, every single peer cooperates in the enforcement of the peers' selection policies

P4P Principles – Architectures

- Mixed Client- P2P System-Based:
- Each peer retrieves network information from a P4P server and send this information to a P2P server, also providing the information about the file to be downloaded
 - Peers' selection policies are implemented by P2P servers / P2P system
 - P2P clients are forced to use the provided sub-set of peers
 - P4P servers have no control on the reliability of the received information

P4P Principles – Architectures

- P4P Server-Based:
- P2P clients send a request to a P4P servers providing the information about the file to be downloaded, and the P4P servers provides the client the set of peers to be contacted for file sharing
 - P2P clients are forced to use the provided sub-set of peers
 - Peers' selection policies are implemented by P4P servers
 - P4P servers can control and limit the sharing of some files
 - Potential bottleneck on the P4P server
 - Providers do not disclose explicit network information

P4P Optimization

- P4P effectiveness depends on what information is communicated through the P4P server and what algorithms are used for controlling traffic
- The optimization algorithm jointly executed by the P2P / P4P infrastructure must take into account:
 - Swarm statistics
 - Peers' bandwidth
 - Peers' position
 - Current levels of background traffic on different links
 - Network status

P4P Optimization

- The objective function can be:
 - Minimization of inter-domain traffic
 - Minimization of other cost-related metrics
 - Minimization of the maximum link utilization
 - Maximization of overall throughput
- Random peers selection should be partially preserved to improve robustness

Caching

- Some ISPs have attempted to reduce P2P traffic by placing caches at the ISP's gateway to the Internet or by using network appliances for forging TCP RST messages, which trick clients into closing connections to remote peers
- Requires cooperation
 - Protocols must not be obfuscated nor encrypted
 - Some solutions exploits HTTP caching, but requires protocol adaptations

Caching

- Alternatively, it is possible to implement the creation of a node running the same “client” software of standard peers
- The provider actively participate in the distribution of files
- However, the identification of most popular files to be downloaded and provided to customers is not easy, and would require deep packet inspection (neutralized by obfuscation or encryption) or some sort of cooperation with P2P networks (e.g. a P4P approach)

Existing Projects

- Ono Project
- Not a strictly P4P approach
- Modification of BitTorrent protocol
- Determines vicinity of peers querying CDNs, such as Akamai and Limelight
- CDNs use dynamic DNS redirection or URL rewriting to send clients to low-latency replica servers located in thousands of ISPs worldwide
- If two client computers are sent to the same CDN server, they are likely to be close to each other (*CDN as Oracle*)
- It relies on preexisting infrastructures (CDNs) that perform extensive Internet measurements

Application-Layer Traffic Optimization (ALTO)

- IETF Working Group
- *Application-Layer Traffic Optimization (ALTO)*
- The Working Group will design and specify an Application-Layer Traffic Optimization (ALTO) service that will provide applications with information to perform better-than-random initial peer selection

- Different optimizations factors:
 - Maximum bandwidth
 - Minimum cross-domain traffic
 - Lowest cost to the user
 - etc.
- Supported systems:
 - BitTorrent
 - Tracker-less P2P
 - Content delivery networks (CDN)
 - Mirror selection

Application-Layer Traffic Optimization (ALTO)

- ALTO will focus on:
- Request/response protocol for querying the ALTO service to obtain information useful for peer selection, and a format for requests and responses
- A specification of core request and response formats and semantics to communicate network preferences to applications
- Initially the WG will consider:
 - IP ranges to prefer and to avoid
 - Ranked lists of the peers requested by the client
 - Information about topological proximity and approximate geographic locations

Application-Layer Traffic Optimization (ALTO)

- The Map Service **Cost Map** query is a batch operation in which the ALTO Server returns the **Path Cost for each pair of source/destination PID** defined by the ALTO Server

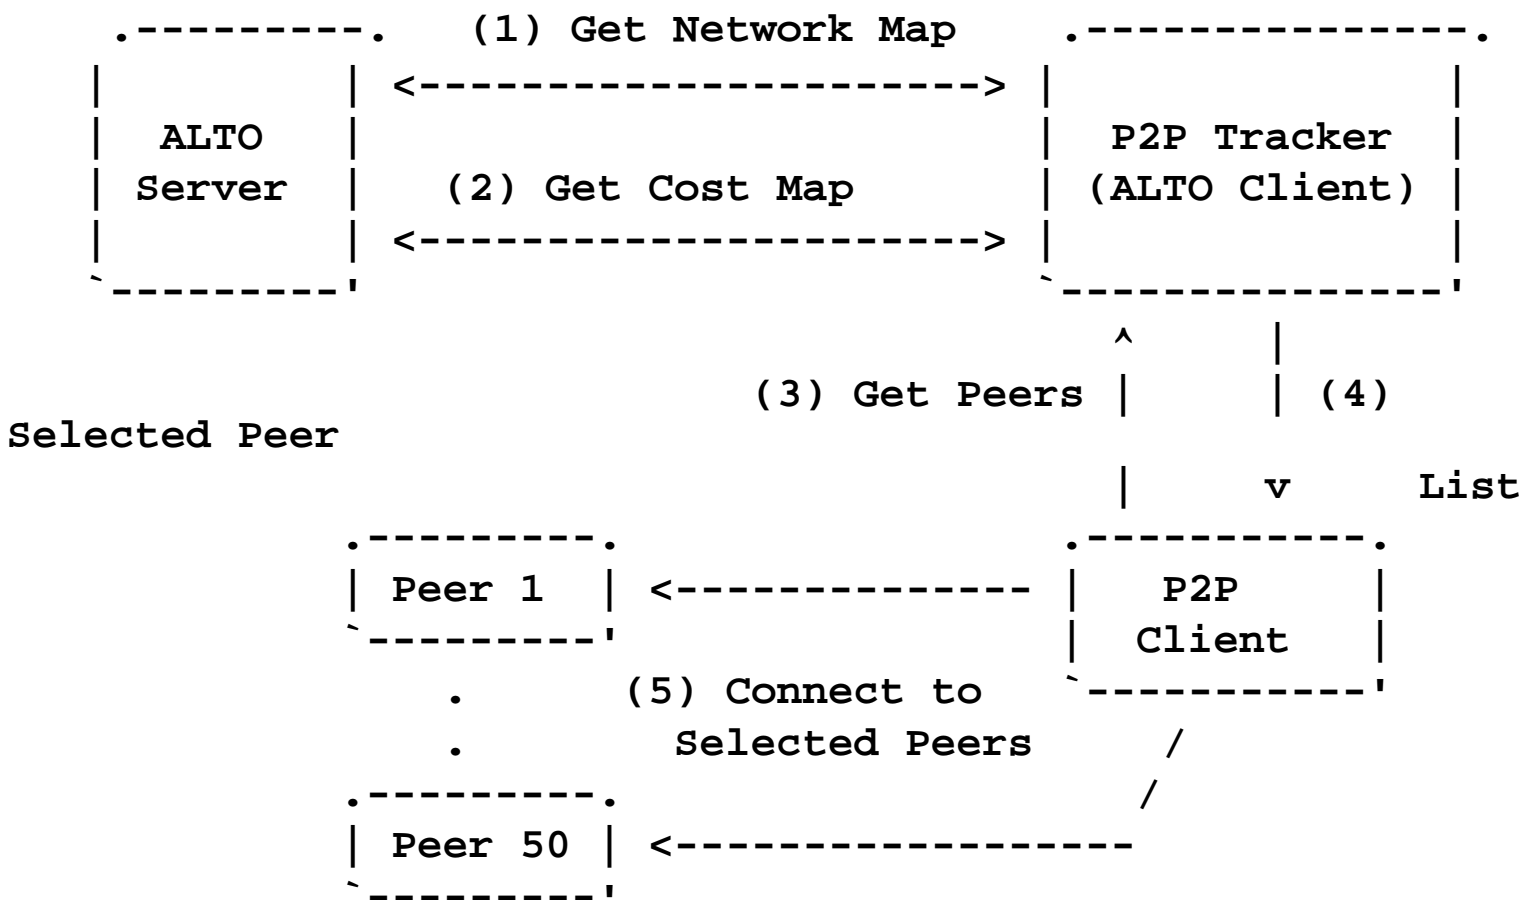
```

•      GET /map/core/pid/cost HTTP/1.1
•      Host: alto.example.com:6671

•      HTTP/1.1 200 OK
•      Content-Length: [TODO]
•      Content-Type: application/alto
•      {
•          "meta" : {
•              "version" : 1,
•              "status" : {
•                  "code" : 1
•              }
•          },
•          "type" : "cost-map",
•          "data" : {
•              "map-vtag" : "1266506139",
•              "cost-type" : "routingcost",
•              "cost-mode" : "numerical",
•              "map" : {
•                  "PID1": { "PID1": 1, "PID2": 5, "PID3": 10 },
•                  "PID2": { "PID1": 5, "PID2": 1, "PID3": 15 },
•                  "PID3": { "PID1": 20, "PID2": 15, "PID3": 1 }
•              }
•          }
•      }
    
```

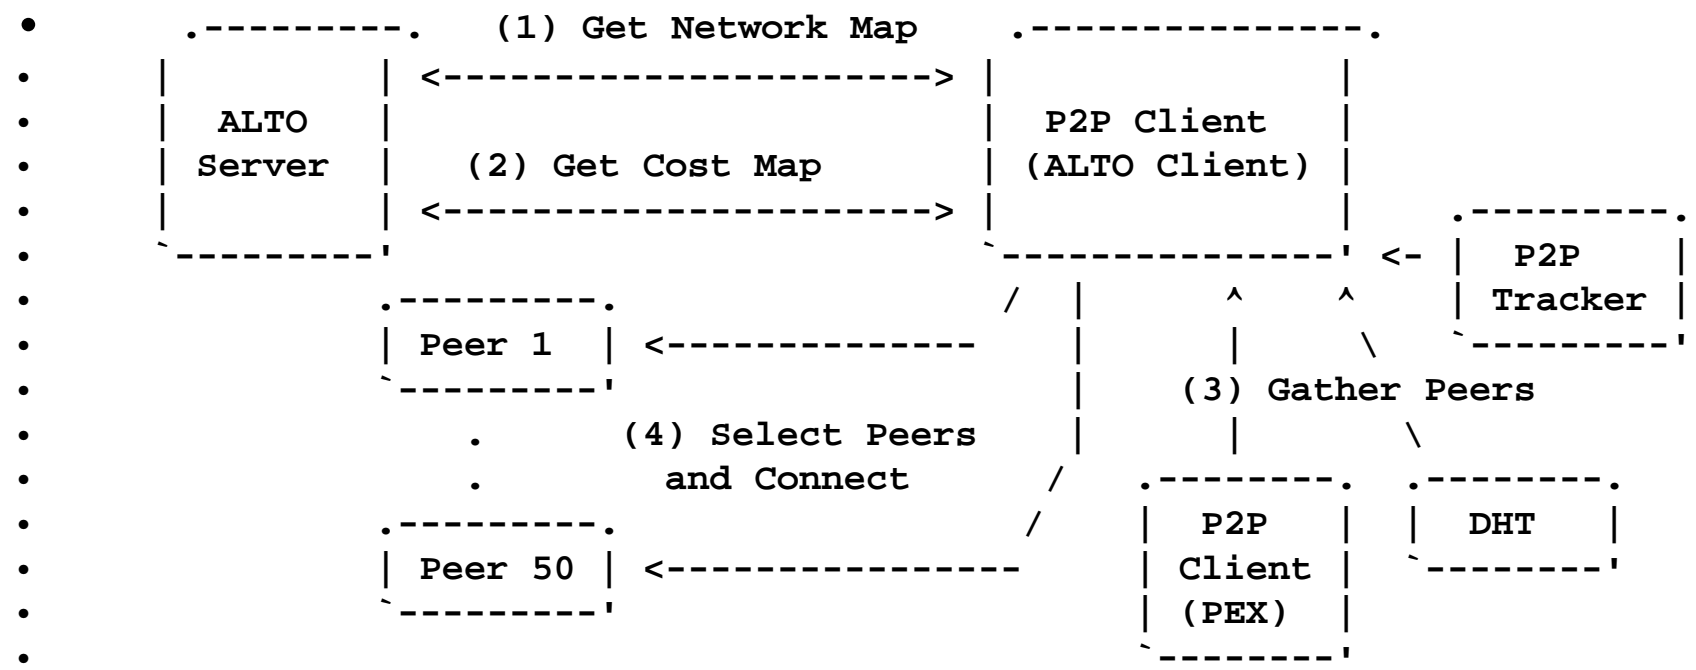
Application-Layer Traffic Optimization (ALTO)

- ALTO Client Embedded in P2P Tracker (P2P System-Based)



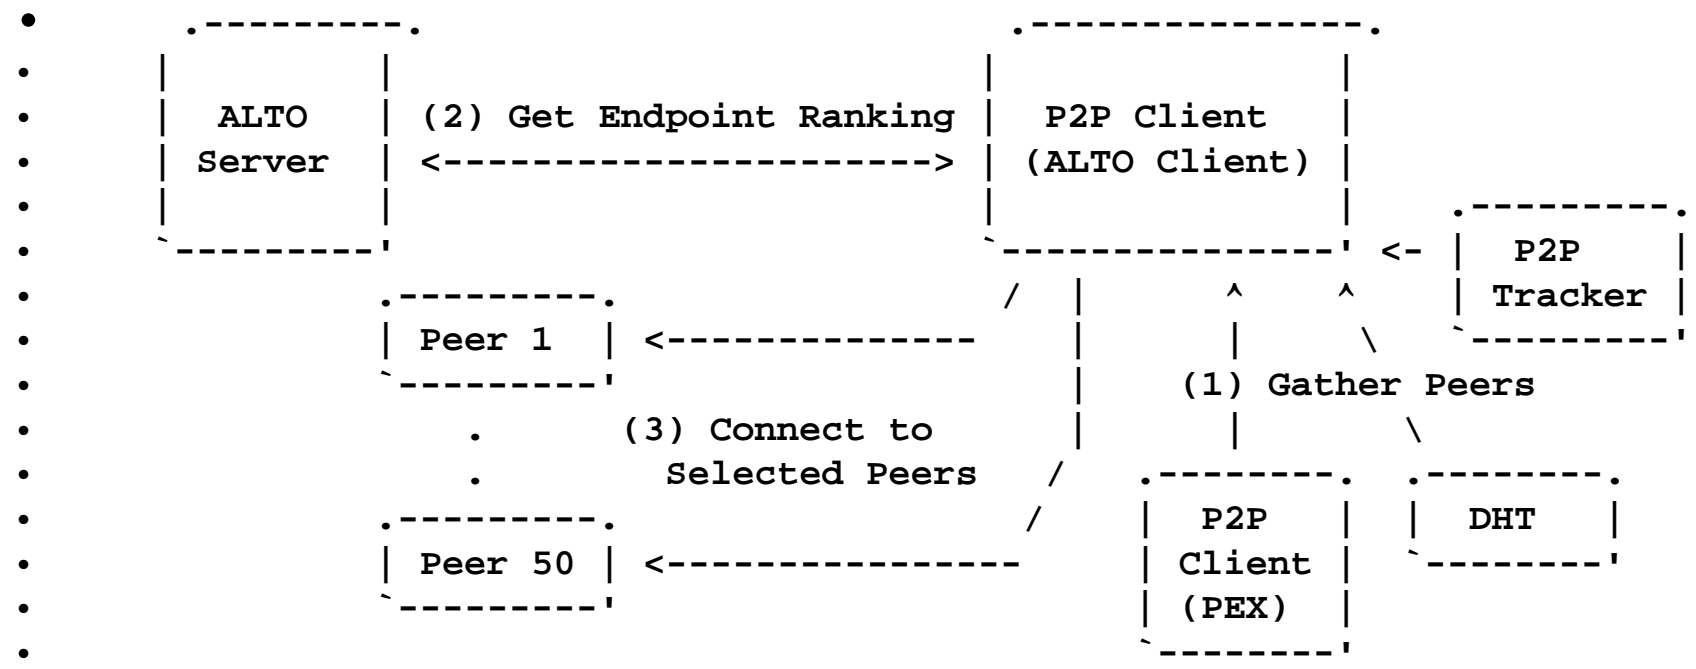
Application-Layer Traffic Optimization (ALTO)

- ALTO Client Embedded in P2P Client (Client-Based)



Application-Layer Traffic Optimization (ALTO)

• ALTO Client Embedded in P2P Client: Ranking (Client-Based)

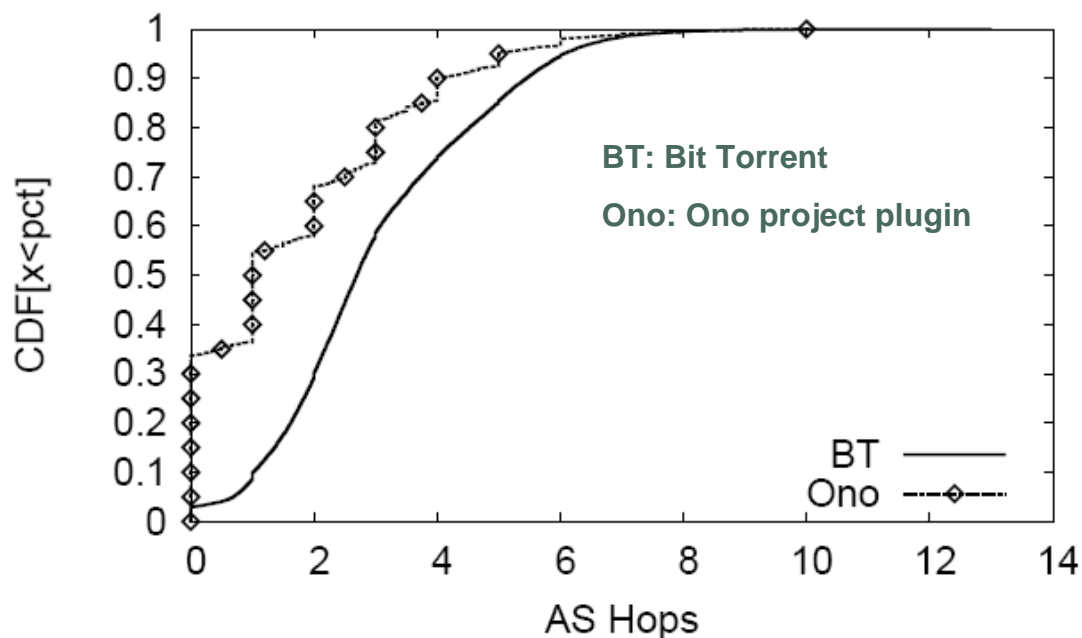


Application-Layer Traffic Optimization (ALTO)

- **ALTO Milestones**
- Done Working Group Last Call for problem statement
- Done Submit problem statement to IESG as Informational
- Jan 2011 Working Group Last Call for requirements document
- Jan 2011 Working Group Last Call for request/response protocol
- Mar 2011 Submit request/response protocol to IESG as Proposed Standard
- Mar 2011 Submit requirements document to IESG as Informational
- May 2011 Working Group Last Call of deployment considerations document
- Aug 2011 Submit deployment considerations document to IESG as Informational
- Nov 2011 Working Group Last Call of discovery mechanism
- Feb 2012 Submit discovery mechanism to IESG as Proposed Standard
- Mar 2012 Dissolve or re-charter

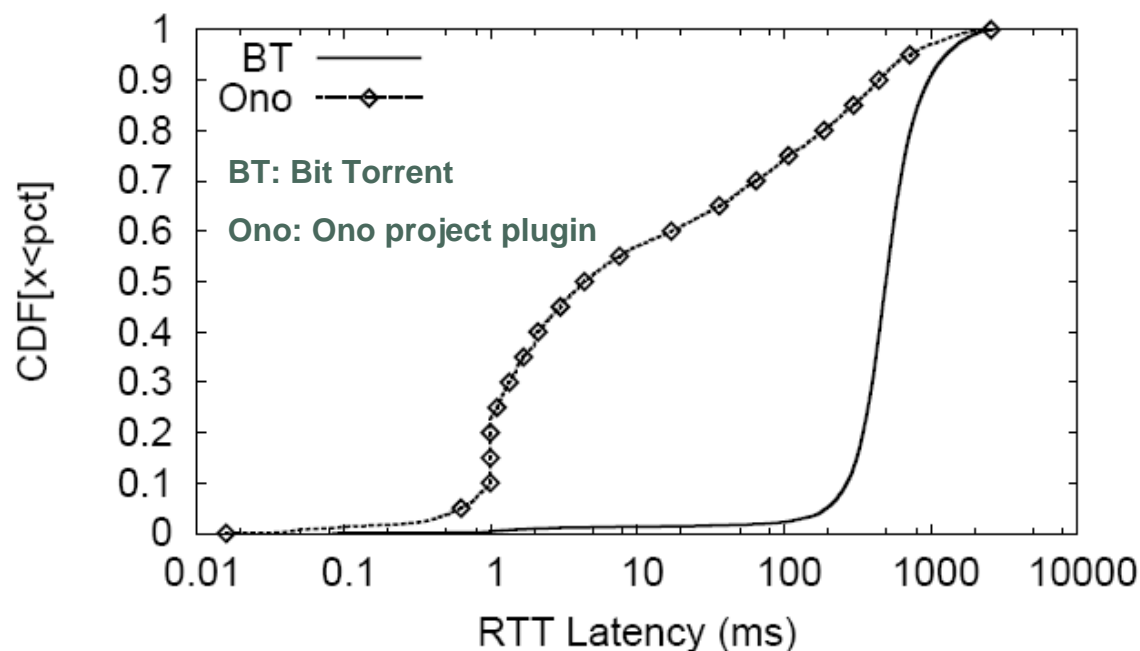
P4P Performance

- **Ono Project** – Performance
- → 33% paths within a single autonomous system (AS)



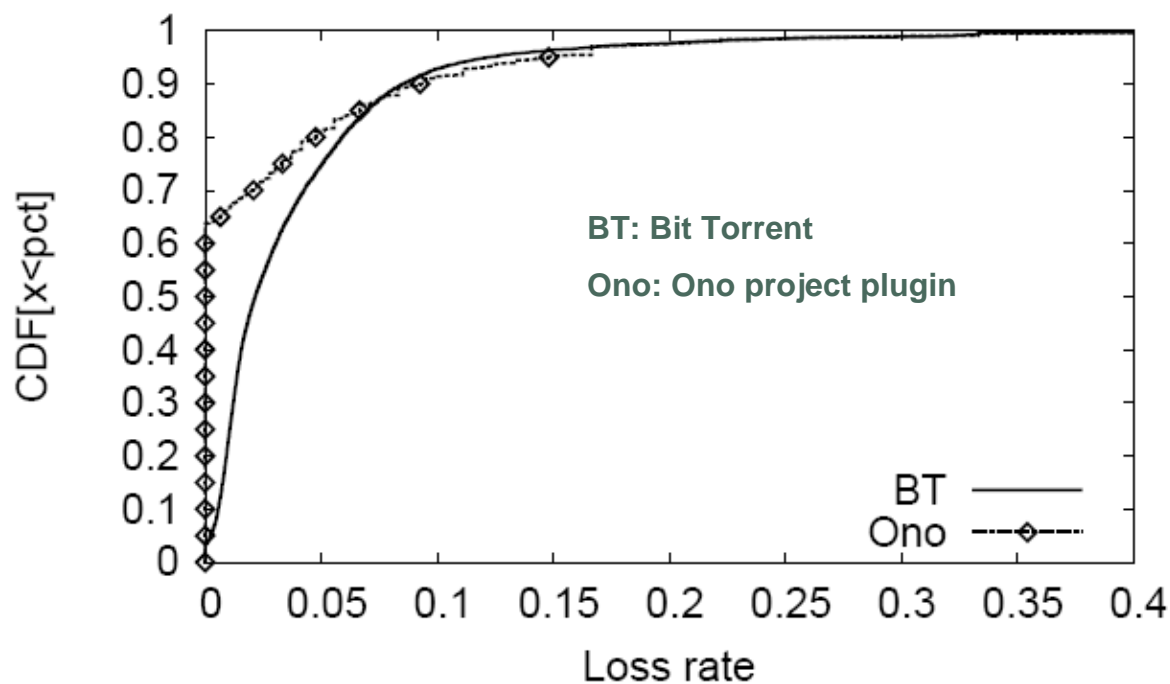
P4P Performance

- **Ono Project** – Performance
- → two orders of magnitude lower latency



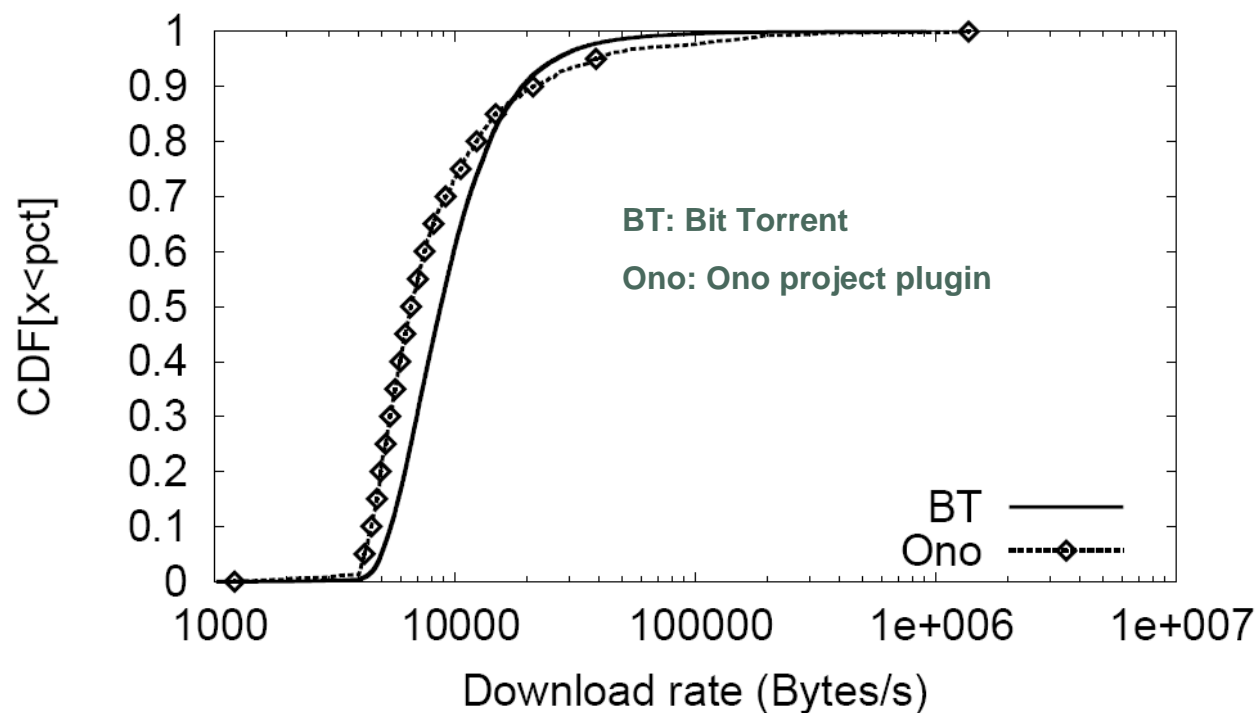
P4P Performance

- **Ono Project** – Performance
- → 30% lower loss rates



P4P Performance

- **Ono Project** – Performance
- → 31% average download rate improvement



(a) Download rates.

P4P Performance

- Ono Project – Performance
- Networks which offers higher transfer rates inside the ISP, allows users to obtain significant performance gains by reducing cross-ISP traffic
- Networks which does not offer different transfer rates for in-network traffic, exhibits negligible performance differences for connections with different AS-path lengths
- Performance from recommended peers will not be significantly different than those picked at random
- Rather than blocking BitTorrent flows, ISPs could change their bandwidth allocations so that it is more favorable to connect to peers inside the ISP than to those outside

Conclusions

- P4P solutions, in order to be fully operational requires the cooperation between providers and P2P users
- Cooperation is feasible if benefits are multi-lateral
- Cooperation by users can be achieved only by providing them significant advantages, e.g. higher download rates
- Privacy and censorship issues are particularly relevant for many P2P users' communities