ALTO Problem Statement

draft-marocco-alto-problem-statement-02

Enrico Marocco
Vijay Gurbani

72nd IETF Meeting

Outline

- History
- The problem
- Main issues
- Use cases
- The cache location “sub-problem”

Internet Applications

197x	Email
198x	File transfer Usenet
199x	Web browsing
200x	Peer-to-peer
2008	

Source: mostly Wikipedia

Internet Applications

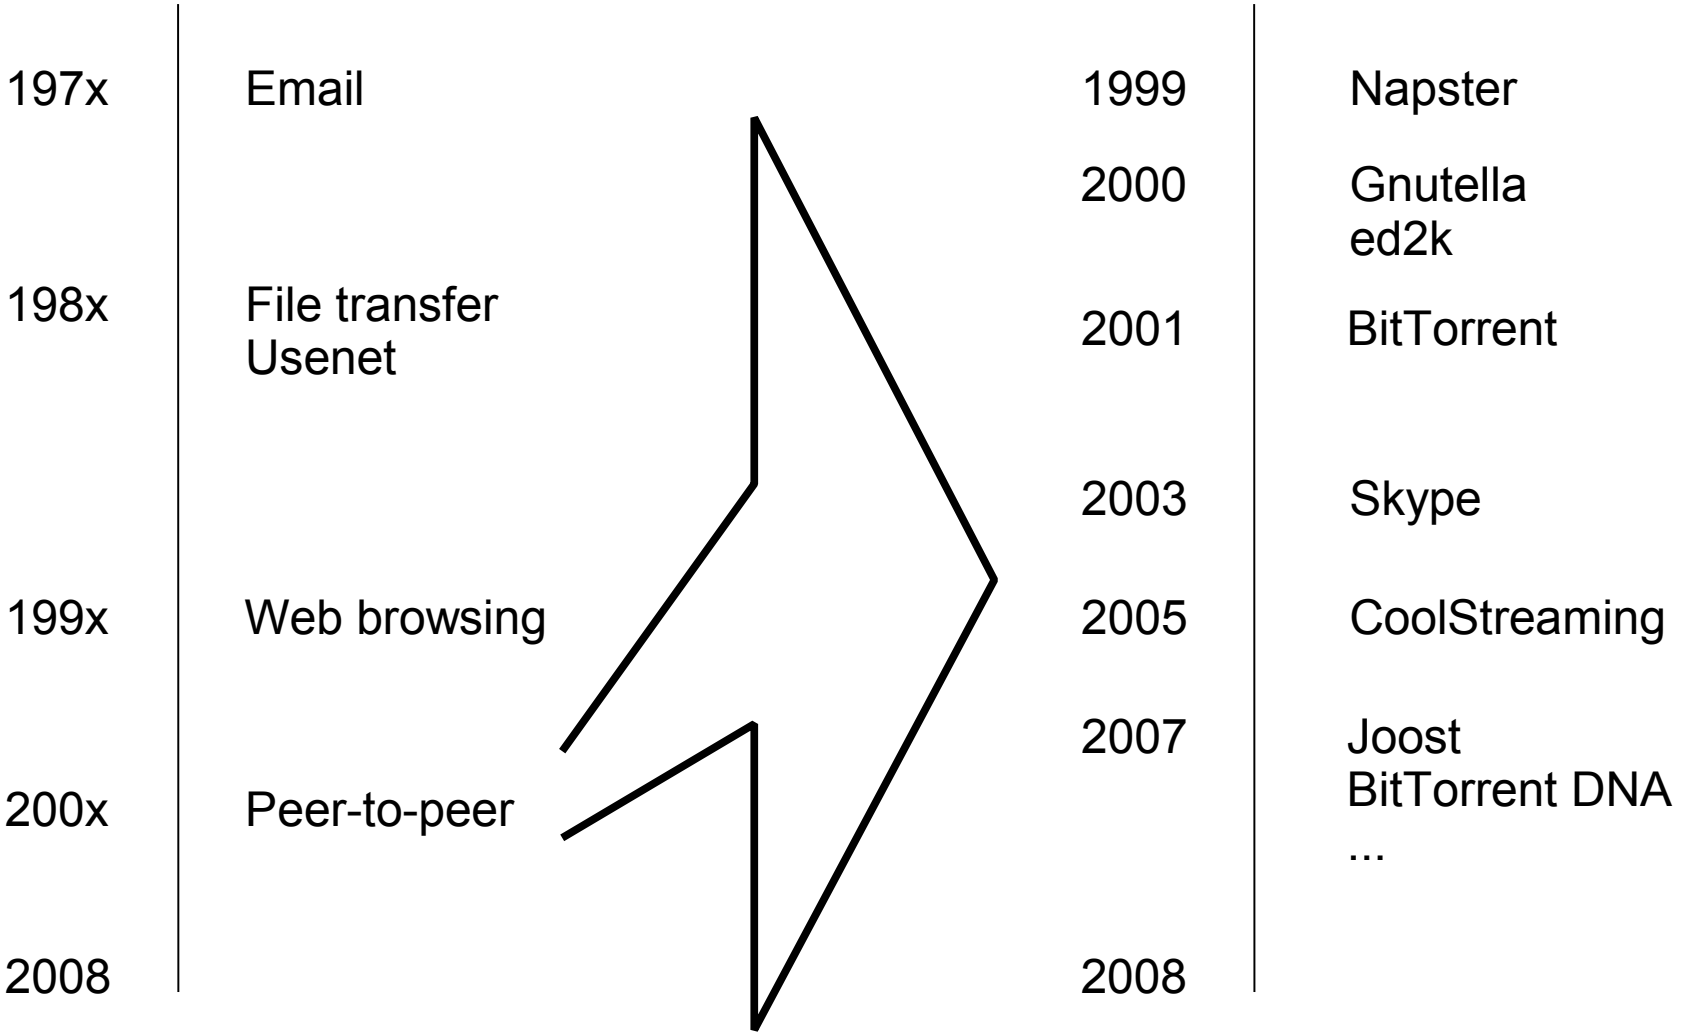

Source: mostly Wikipedia

Peer-to-peer Traffic

- 50% - 85% of total traffic
- Upstream as well as downstream
- Bandwidth-greedy
- Interferes with real-time traffic
- Unpredictable
- ...

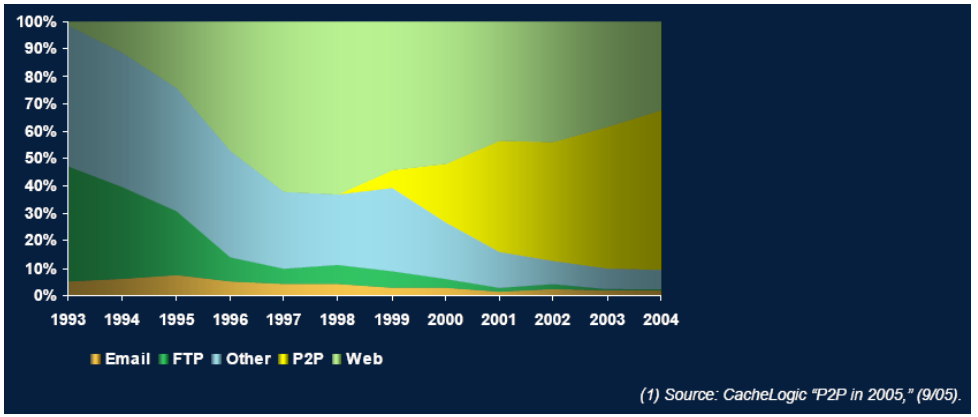

Incoming traffic

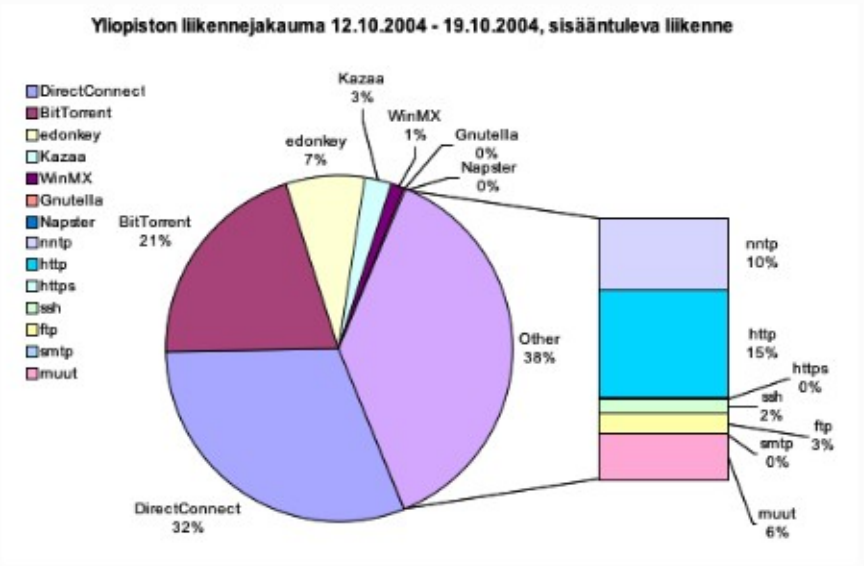

Outgoing traffic

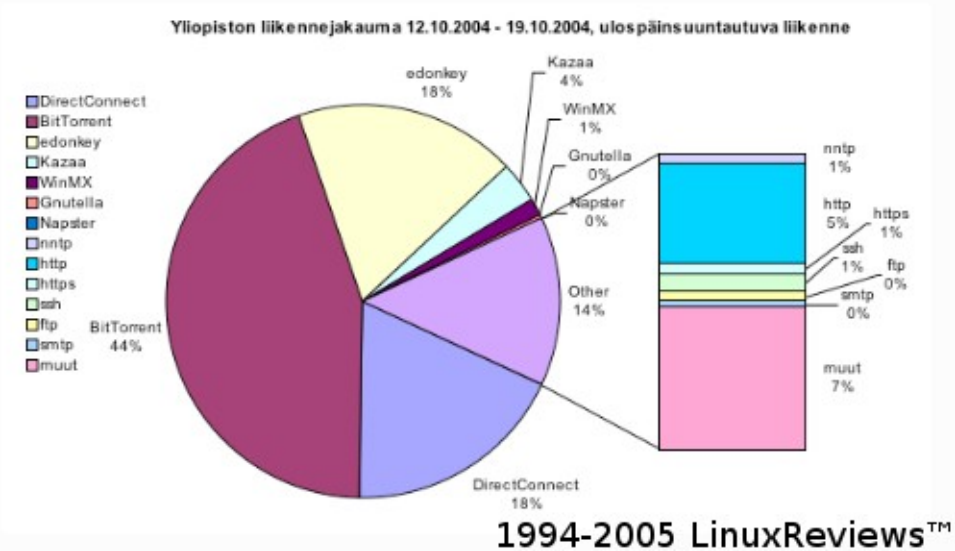

P2P Traffic in the News

- “Comcast Throttles BitTorrent Traffic. Seeding Impossible”[1]
- “ISPs Fear iPlayer Overload”[2]
- “Comcast and BitTorrent Agree to Collaborate”[3]
- “Verizon Reports P4P Can Slash P2P's Impact on ISPs”[4]
- “New Software Allows ISPs & P2P to Get Along Without Getting too Cozy”[5]

References

1. August 2007, http://torrentfreak.com/comcast-throttles-bittorrent-traffic-seeding-impossible.
2. August 2007, http://www.bnvillage.co.uk/games-village/91455-isps-fear-iplayer-overload.html.
3. March 2008, http://news.cnet.com/8301-10784_3-9904494-7.html.
4. March 2008, http://www.newsfactor.com/story.xhtml?story_id=032002XVIJS0.
5. May 2008, http://esciencenews.com/articles/2008/05/05/new.software.allows.isps.and.p2p.users.get.along.without.getting.too.cozy.

IETF P2P Infrastructure Workshop

- Boston, May 29, 2008
- Organized by RAI ADs
- Discuss problems related to P2P traffic
- Identify a reasonable solution space
- Three different (complementary) approaches:
  - Localization and caches
  - New approaches to congestion
  - Quality of service

IETF P2P Infrastructures Workshop

- Boston, May 29, 2008
- Organized by RAI ADs
- Discuss problems related to P2P traffic
- Identify a reasonable solution space
- Three different (complementary) approaches:
  - **Localization and caches (RAI/APP)**
  - New approaches to congestion (TSV)
  - Quality of service (TSV)

What's New in Network Applications

- Client/Server
  - Target is a host (one or few IPs)
  - Traffic optimization consists of finding the best network path
  - GeoDNS, DiffServ, MPLS...
- Peer-to-peer
  - Target is a resource (usually shared by many peers)
  - Traffic optimization consists of selecting the “best” peer(s)
  - Vivaldi, iPlane, Ono, P4P, IDIPS...

The ALTO Problem

- Peers have no knowledge of the network topology
  - Common case in file-sharing: a peer in Dublin downloads a chunk from a peer in Tokyo when the same chunk is available in London
- No optimization causes congestion (bad for ISPs **and** bad for P2P)
- Endpoints are in the worst position for selecting the “best” peer(s)
  - Typically hundreds/thousands of possible peers
  - Measurements either too poor or too expensive

Addressing the ALTO Problem

- Defining an interface for a peer selection optimization service
  - Request: I am peer P and have to exchange *n* Mb of real-time/bulk data with anyone among X, Y, Z
  - Response:
    - Choose X!
    - You are in AS_1, X is in AS_1, Y is in AS_2 and Z is in AS_3
    - Bit-cost from P is: *j* to X, *k* to Y and Z
    - X is located at (39.3° N 76.6° W), Y at ...
    - ...
    - Any reasonable combination of the above

Architecture

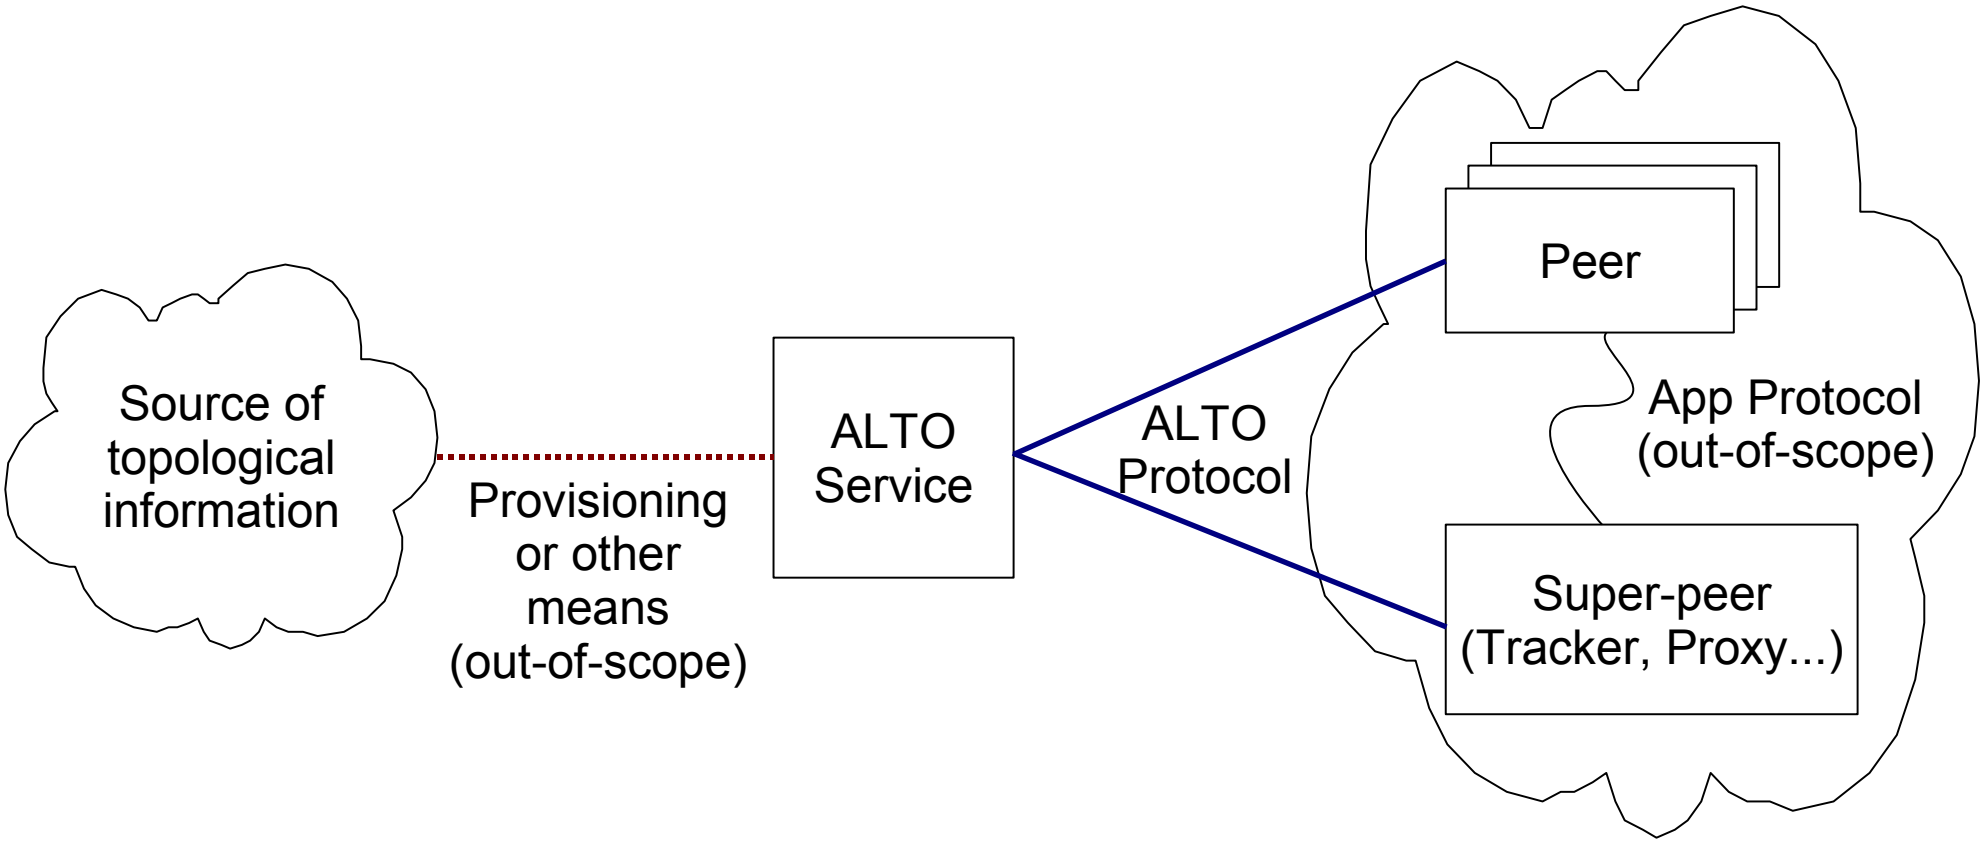

ALTO Service Providers

- Network operators
  - Know the network topology and the peering policies
- Communities
  - Running distributed algorithms (Internet coordinate systems, distributed path evaluation algorithms...)
- Third-parties aware of the network topology
  - E.g. exploiting redirections from distributed services (e.g. Ono & Akamai)
  - On behalf of ISPs

“The (desired) ALTO Effect”

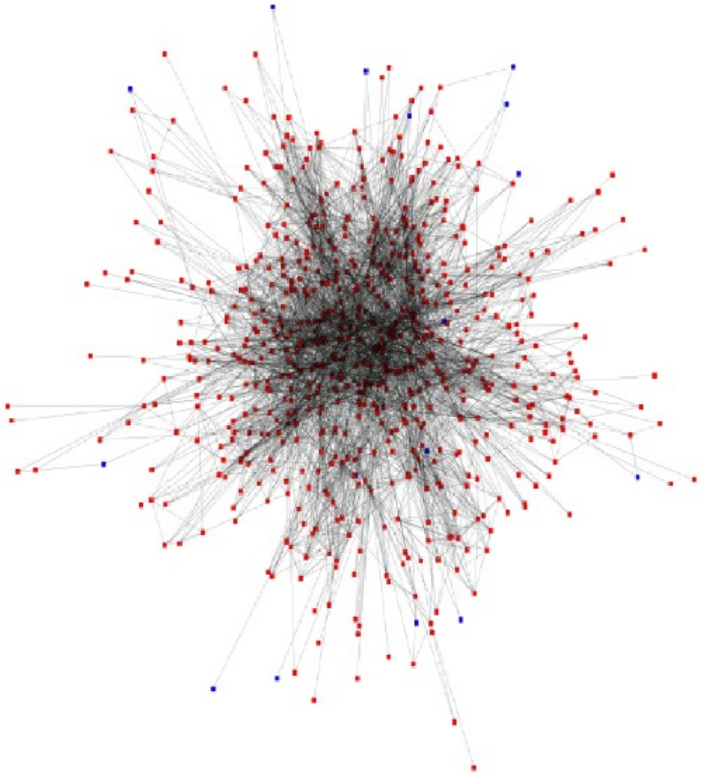

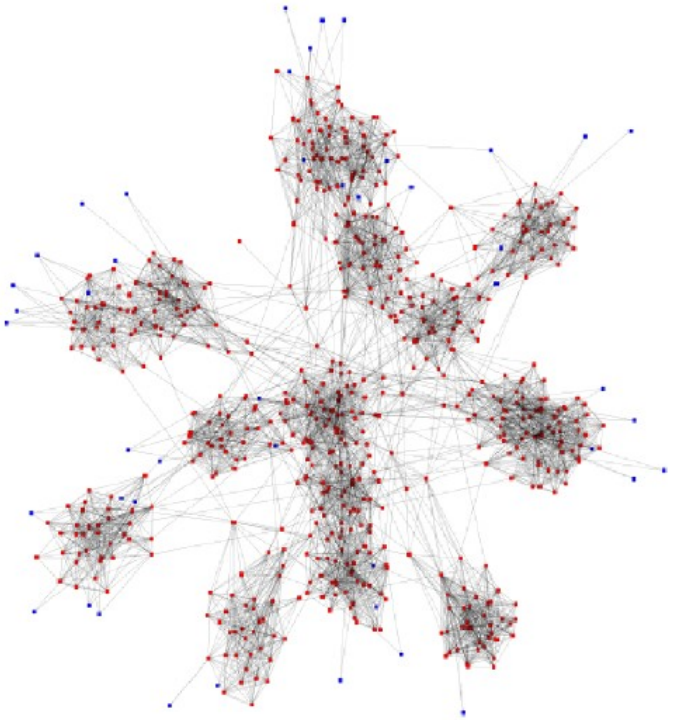

(Gnutella simulations)

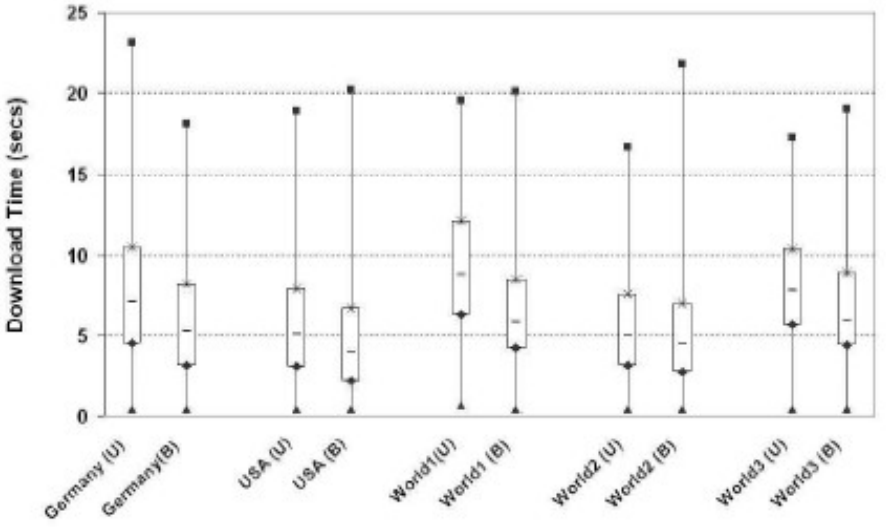

V. Aggarwal, A. Feldmann, C. Scheideler. **Can ISPs and P2P systems co-operate for improved performance?**

V. Aggarwal, O. Akonjang, A. Feldmann. **Improving User and ISP Experience through ISP-aided P2P Locality**

“The (desired) ALTO Effect”

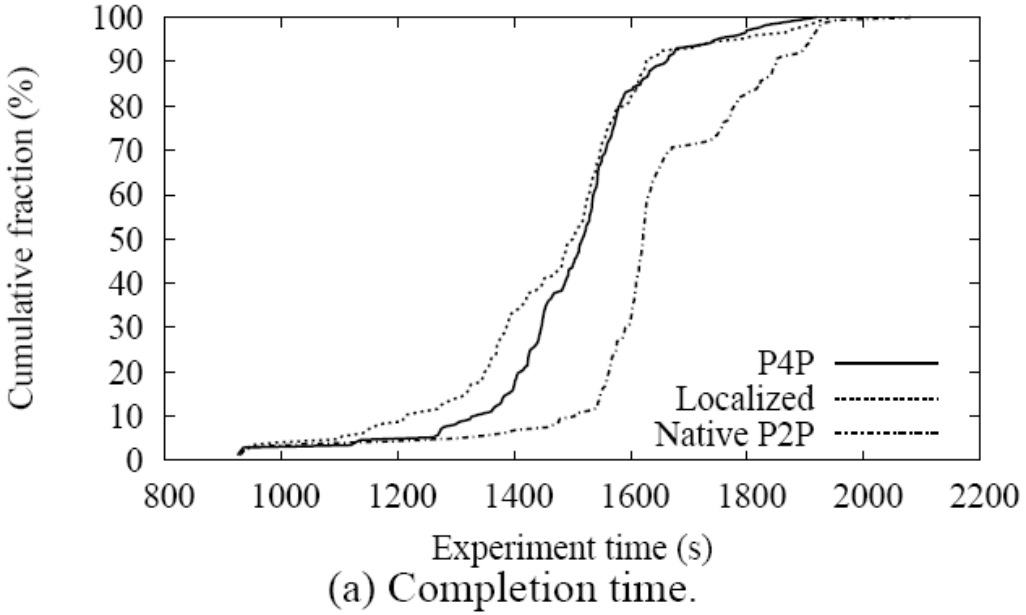

(a) Completion time.

(BitTorrent experiments)

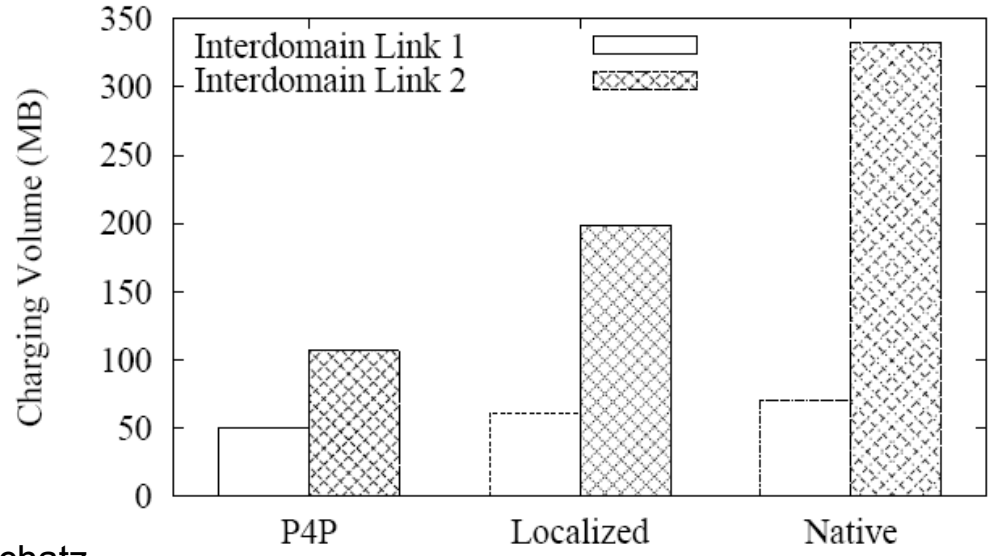

(b) Charging volumes.

H. Xie, Y. R. Yang, A. Krishnamurthy, Y. Liu, and A. Silberschatz.
P4P: Provider Portal for Applications

Issues: Topology Hiding

- As a matter of fact, ISPs consider their networks' internals as reserved information
- Goal: to be able to provide network topology information without revealing network topology
  - Provide arbitrary priority values (e.g. IDIPS)
  - Use opaque identifiers and return perturbed distance values (e.g. P4P)

Issues: Locating the Oracle

- Unlikely to have a centralized service
- An oracle could be virtually everywhere, but...
  - Most relevant information concerns the querying peer's network (i.e. the best oracle may be the closest)
  - It may be useful to get topology information about the networks of the peers under evaluation

Issues: Trust

- What prevents an ALTO service to mis-behave and:
  - Redirect querying peers to corrupted mediators
  - Collect information to track P2P connections
  - Apply sub-optimal policies (i.e. to second economic factors other than network efficiency)
- Hint: ALTO is optional

Core Blocks of an ALTO Solution

- Discovery mechanism for locating the oracle
  - “What ALTO server should I query from my location?”
- Query/Response protocol for querying the oracle
  - “I can connect to X, Y, Z; who should I choose?”

Use Cases: File-sharing

- Shared files/chunks are often available from multiple sources
  1) First selection is usually random (from ~10^3 to ~10)
  2) Then selection based on goodput, tit-for-tat...
- ALTO may be useful for (1) above
  - In P2P clients
  - In trackers, where available

Use Cases: RT Communications

- Selection of the closest media relay for NAT traversal
- Especially useful in highly distributed services (e.g. Skype, P2PSIP)
  - Any client is potentially a media relay

Use Cases: P2P Streaming

- Selection of the “best” peer(s) to send/receive a stream to/from

Use Cases: Mirror Selection

- Providers of popular content (e.g. media and software repositories) resort to geographically distributed mirrors
  - Manual selection
  - Automatic selection through Geographical DNS Load Balancing
- ALTO may be adopted both client-side and server-side

Use Cases: DHTs

- Some DHTs use proximity information for populating peers' routing tables
  - E.g. Pastry, Bamboo, CAN
  - Usually based on RTT estimation
- ALTO could provide additional information

Peer Selection and Cache Location

- In theory, caches could be transparently handled as if they were peers
  - Caches are nothing but powerful and selfless peers
  - If an ALTO server recognizes caches' addresses in the request, it can simply put them on the top of the list
- But, for example...
  - A cache may not be involved in a swarm
  - Chances that caches involved in a swarm are not passed to the client may be very high
    - E.g. if the tracker limits the number of peers passed to the client

Peer Selection and Cache Location

- Peers may be interested in locating caches
  - Offline – through an application specific cache discovery mechanism
  - Within the ALTO transaction
    - Useful if the ALTO service is aware of caches
    - Requires the querying peer to pass additional information (application-id, content-id...)
- Cache location is a good fit for ALTO, but MUST be optional
  - Many (most of?) potential adopters will not want to disclose sensible information