

ISP-Aided Neighbor Selection for P2P Systems

Vinay Aggarwal



vinay@net.t-labs.tu-berlin.de

Anja Feldmann, Obi Akonjang, Christian Scheideler (TUM)

Deutsche Telekom Laboratories
TU Berlin, Germany

P2P traffic

- >50% of Internet traffic
 - Examples: Bittorrent, eDonkey, Skype, GoogleTalk...



Internet traffic distribution 2007 (Germany)
Source: ipoque GmbH (Nov 2007)

P2P from an ISPs view

- ❑ Good:
 - P2P applications fill a void
 - P2P applications are easy to develop and deploy
 - P2P applications spur broadband demand
- ❑ Bad:
 - P2P systems form overlays at application layer
 - Routing layer **functionality duplicated** at app layer
 - P2P topology agnostic of underlay → performance loss
 - Traffic engineering difficult with P2P traffic
- ❑ ISPs are in a **dilemma**

Solution: ISP-P2P cooperation

- Insight: **ISP knows its network**
 - Node: bandwidth, geographical location, service class
 - Routing: policy, OSPF/BGP metrics, distance to peers

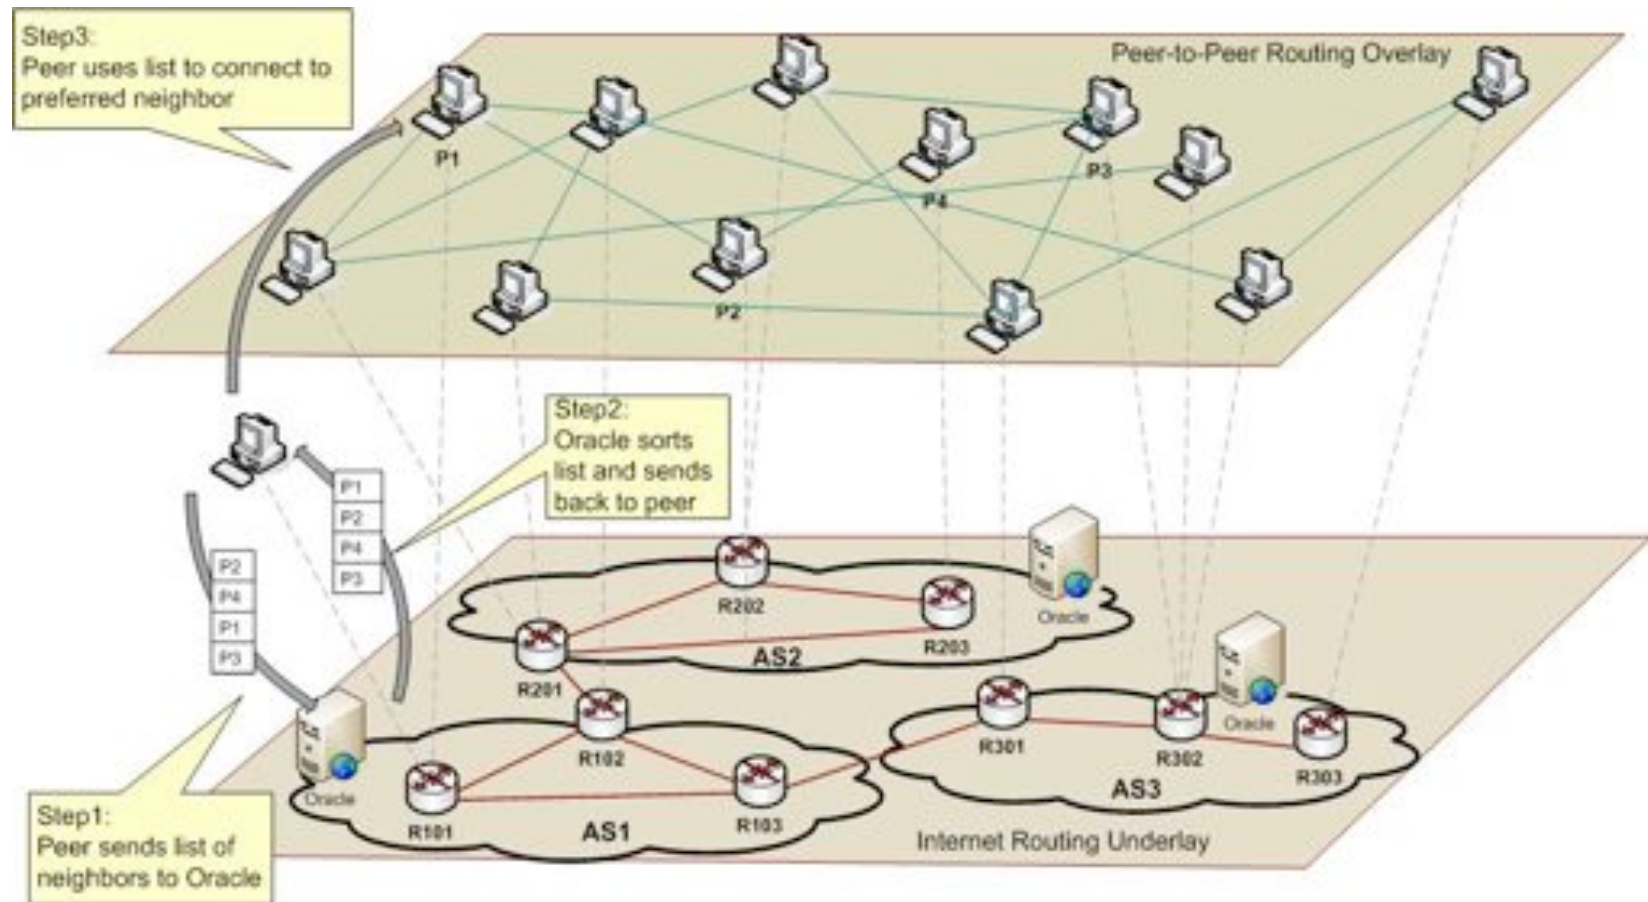
Solution: ISP-P2P cooperation

- ❑ Insight: **ISP knows its network**
 - Node: bandwidth, geographical location, service class
 - Routing: policy, OSPF/BGP metrics, distance to peers
- ❑ Our idea:
 - ISPs: offer oracle that provides network distance info
 - P2P: use oracle to build P2P neighborhoods

Solution: ISP-P2P cooperation

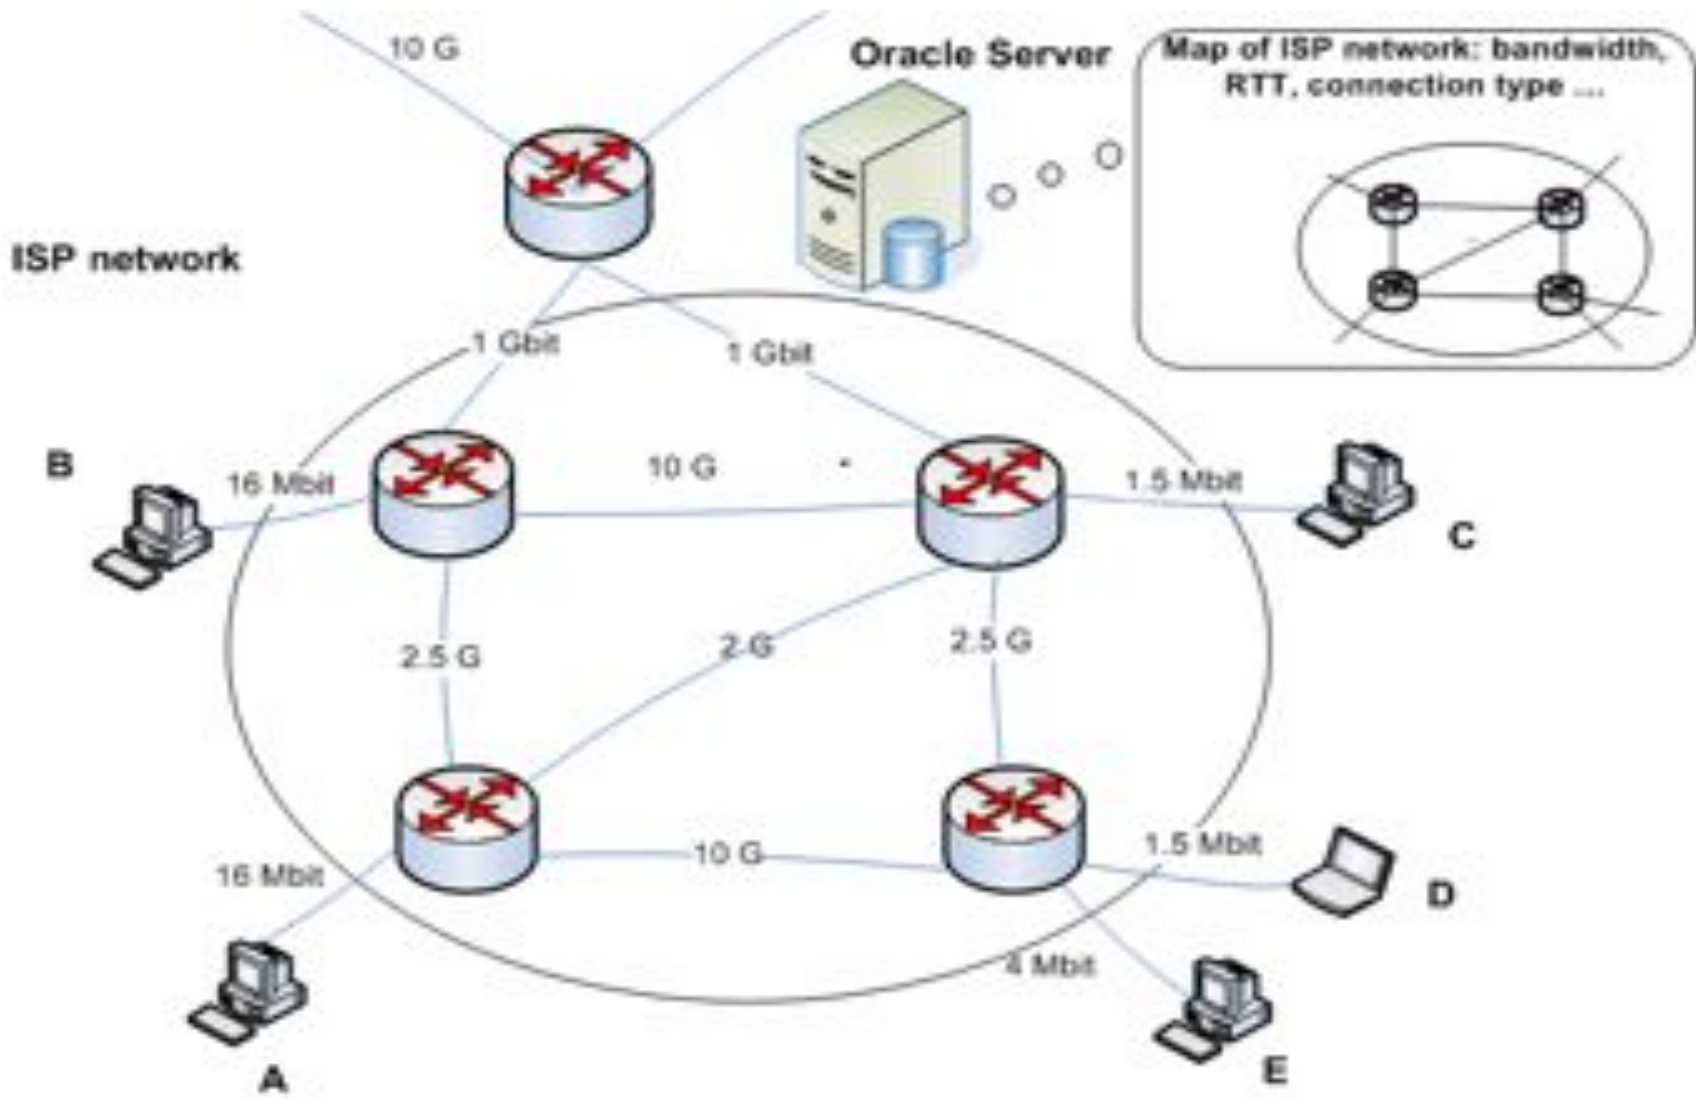
- ❑ Insight: **ISP knows its network**
 - Node: bandwidth, geographical location, service class
 - Routing: policy, OSPF/BGP metrics, distance to peers
- ❑ Our idea:
 - ISPs: offer oracle that provides network distance info
 - P2P: use oracle to build P2P neighborhoods
- ❑ Oracle concept
 - Service of AS / ISP
 - Input: list of possible dst IPs
 - Output: ranked list of dst IPs
 - E.g. according to distances between src IP and dst IPs

Oracle service




Oracle-based peer selection
→ for topology and content exchange

Oracle service (example)



Oracle service (2.)



Oracle-based peer selection
→ localizes topology and traffic

ISP-P2P cooperation

- ❑ **ISP-aided optimal P2P neighbour selection**
 - Simple and general solution, open for all overlays
 - Run as Web server or UDP service at known location
 - Similar to *bind* (DNS)?

- ❑ **Benefits: P2P**
 - No need to measure path characteristics
 - Easy to avoid bottlenecks => better performance

- ❑ **Benefits: ISPs**
 - Regains control over traffic
 - Less traffic leaving network => cost savings
 - Customer service improvement

ISP/P2P concerns

❑ Network info

- ISP ranking algorithm confidential
- Output is ranked list or classification, actual statistics not revealed
- List can be anonymized and dynamically altered
- Network info „revealed“ is already possible to reverse-engineer using available tools

❑ Legal Issues

- No caching of content at ISP
- Oracle is a peer mapping service => consulting oracle does not imply participation in file-sharing


Graph experiments

In ACM SIGCOMM CCR'07

- ❑ Overlay graph structure not affected
 - Node **degree**, overlay **path length** unchanged
 - Graph **connected**, **diameter** constant
- ❑ **Intra-AS peerings** increase heavily
- ❑ **Densely connected subgraphs** local to ISPs
 - P2P **topology correlated** with AS topology
- ❑ **Congestion analysis** using flow conductance
 - Traffic distribution with oracle **near-optimal**

- ❑ **Feasibility study** in testbed and Planetlab

Overlay-underlay topology correlation



Random vs. biased P2P topology


Packet-level simulations

In Global Internet 08

- ❑ Study the Impact of Topology and User-behavior patterns on end-user performance

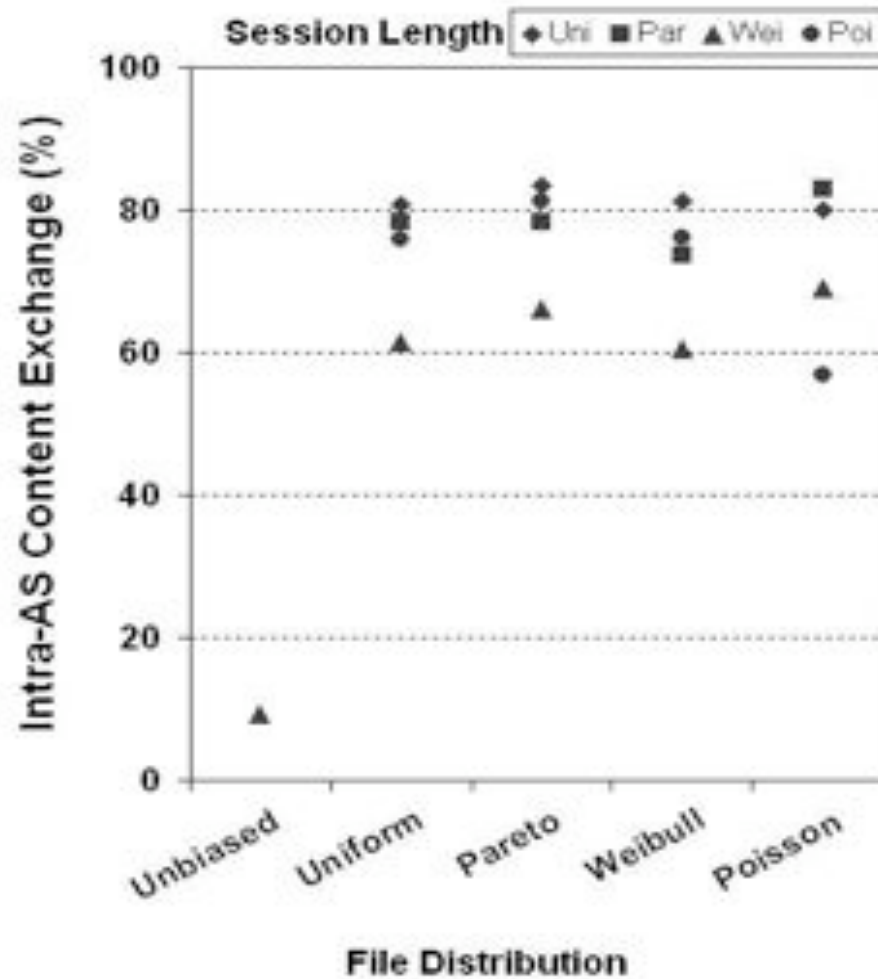
- ❑ Methodology
 - Sensitivity study
 - Use different ISP / P2P topologies
 - Use different user behavioral patterns
 - Content availability, churn, query patterns
 - Evaluate effects of on end-user experience

ISP experience: Intra-AS content




- Content stays within ISPs network
 - Without oracle 10 to 35%
 - With oracle 55 to 80%

ISP experience: Intra AS content (2.)



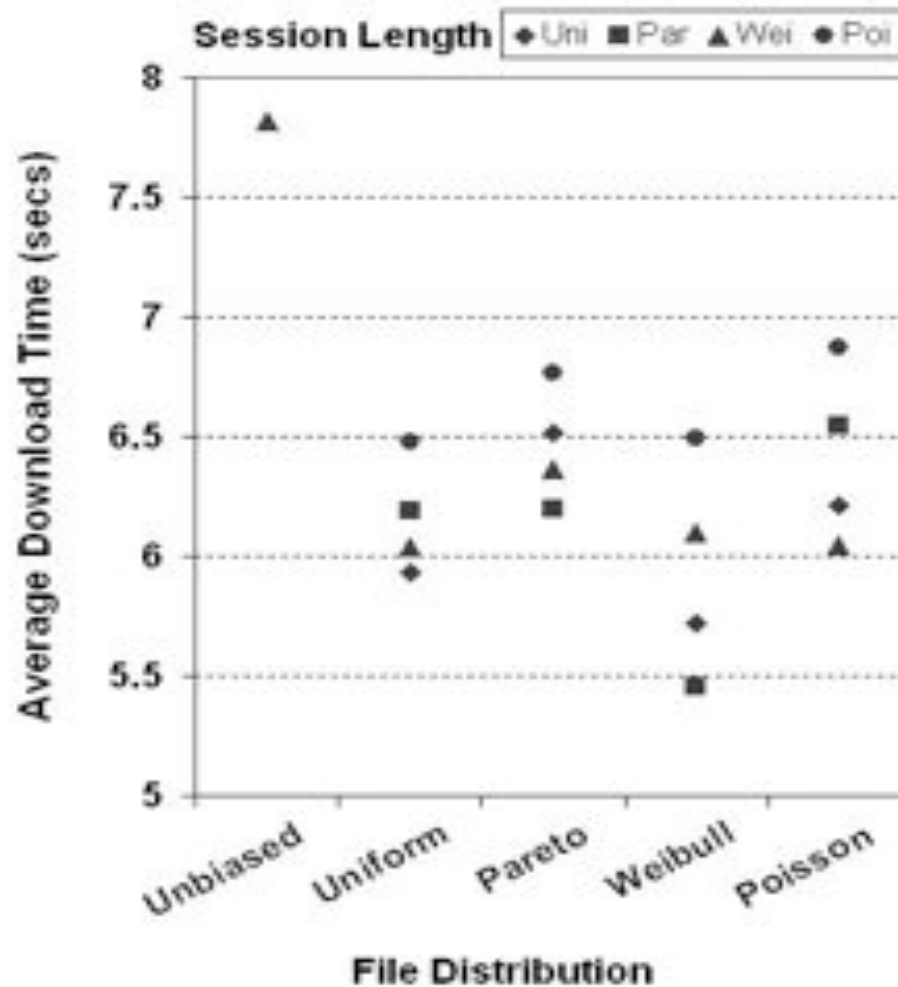
- Content stays within ISPs network

User experience: Download time



- ❑ Mean download time reduction: 1 – 3 secs (16 – 34%)
- ❑ Consistent across topologies

User experience: Download time (2.)




- Reduced mean download time

Summary

- ❑ Oracle
 - Simple and easy to implement
- ❑ Evaluation shows
 - Overlay graph structure not affected
 - Reduced AS distance
 - P2P topology correlated with AS topology
 - Traffic congestion analysis
 - Traffic distribution close to theoretical optimum
- ❑ Benefits
 - ISPs: regain control of network traffic
 - P2P network: sees performance improvements

Extension: Multiple ISP collaboration



Oracle based global coordinate system

- ❑ Question:
 - What bandwidth is available to IP address A
 - What is the delay to IP address B
- ❑ Insight: ISP knows
 - Backbone link capacities and current utilizations
 - Routing policy to neighbouring ISPs
 - Their IP address ranges
- ❑ **Combine oracles to build global coordinate system**
 - Use Oracle within AS
 - Oracle contacts other ISPs oracles for additional info
 - Provide summary information

Upcoming

- ❑ Oracle software release
 - Open source implementation will be available (Based on *bind*)
- ❑ Software patches for popular P2P clients
 - Gnutella
 - BitTorrent
 - eDonkey
 - P2P TV
- ❑ Project website:
<http://www.net.t-labs.tu-berlin.de/isp-p2p/>