ISP-Aided Neighbor Selection for P2P Systems

Anja Feldmann
<anja at net.t-labs.tu-berlin.de>
Vinay Aggarwal, Obi Akonjang, Christian Scheideler (TUM)

Deutsche Telekom Laboratories
TU-Berlin

P2P traffic

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

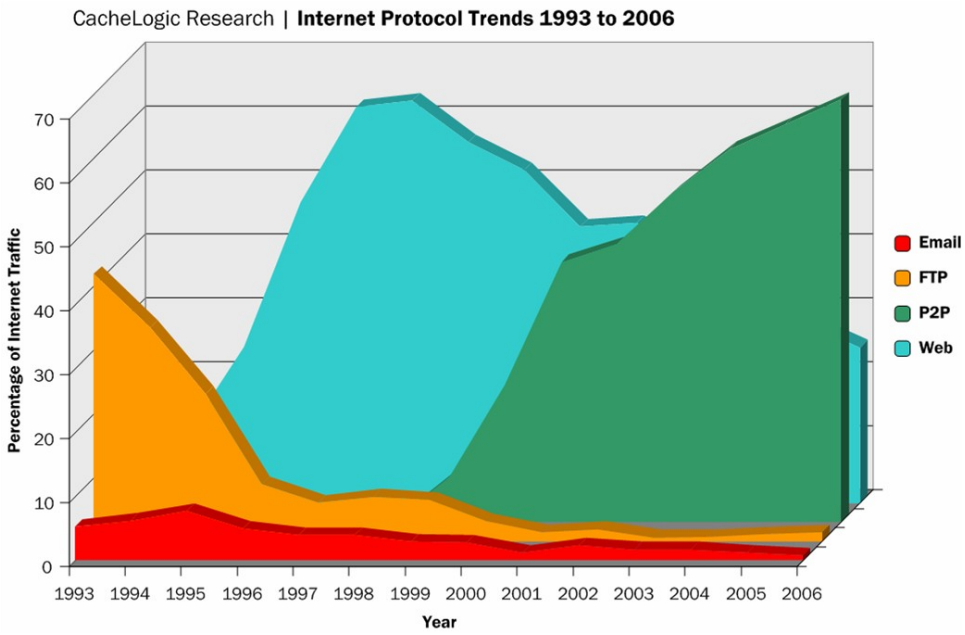

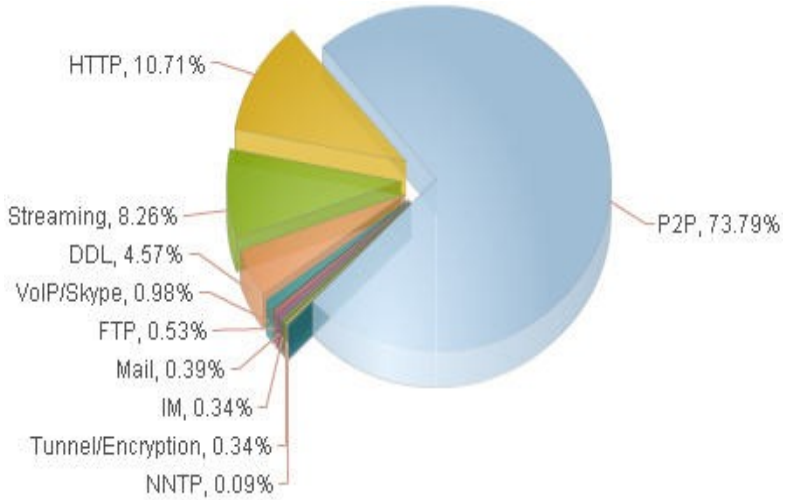

Internet traffic distribution 2007 (Germany)
Source: ipoque GmbH (Nov 2007)

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

ISP dilemma

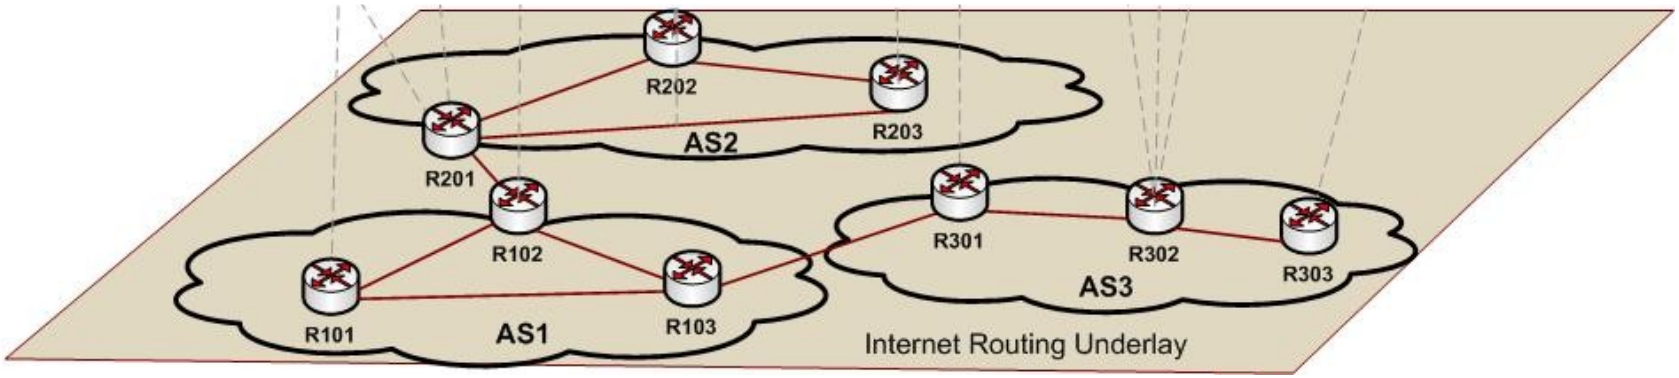

Random/RTT-based peer selection
→ inefficient network resource usage

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

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 neighborhood

- P4P
  - Provide interfaces for applications and networks to communicate regarding
  - Example: Modified iTracker for BitTorrent

Oracle service

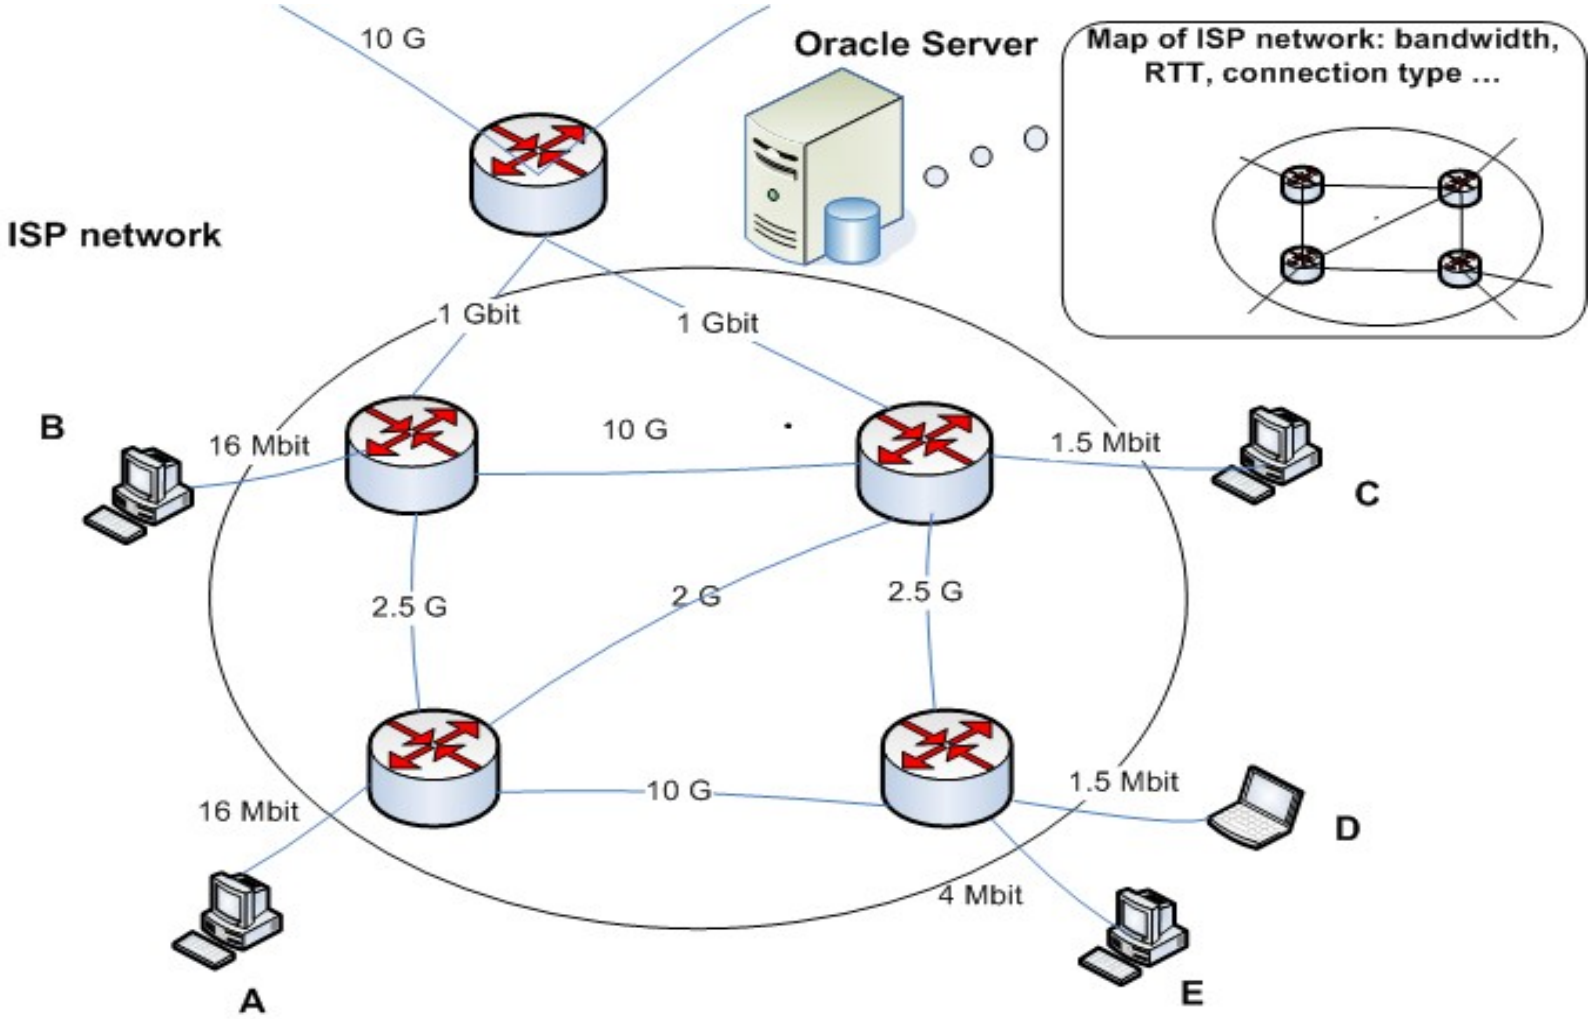

Solution: ISP-P2P cooperation

- Insight: ISP knows its network
  - Node: bandwidth, geographical location, service class
  - Routing: policy, OSPF/BGP metrics, distance to peers
- Oracle concept
  - Service of AS / ISP
  - Input: list of possible dst IPs
  - Ouput: ranked list of dst IPs
    - E.g. according to distances between src IP and dst IPs

Oracle service (2.)

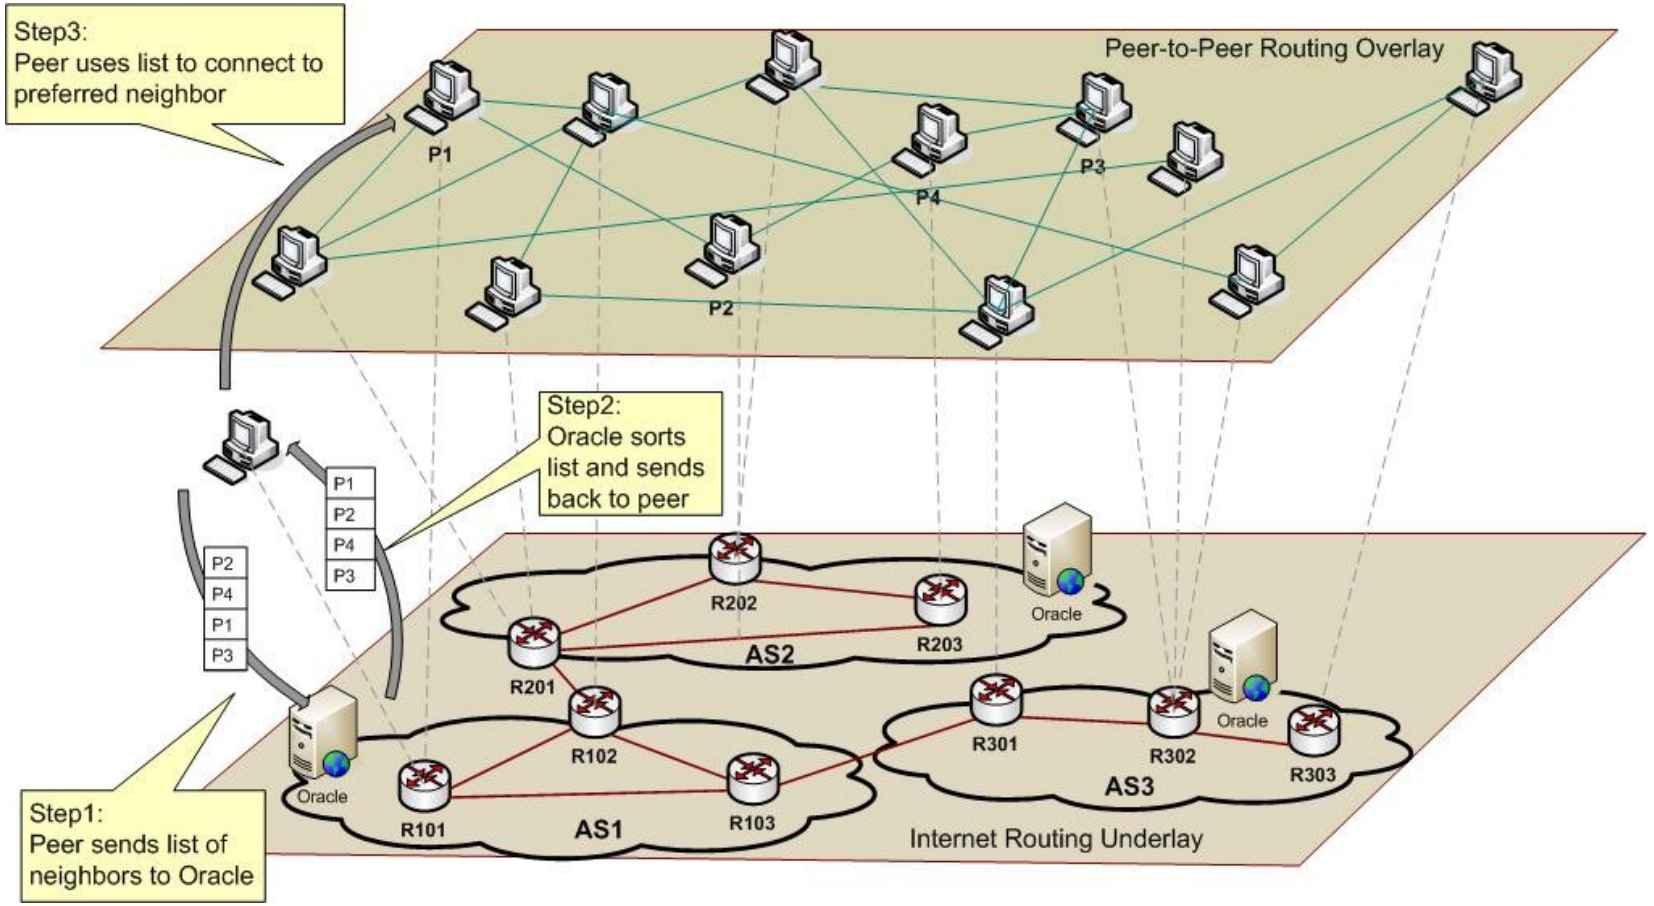

Oracle-based peer selection

→ for topology and content exchange

Oracle service (3.)

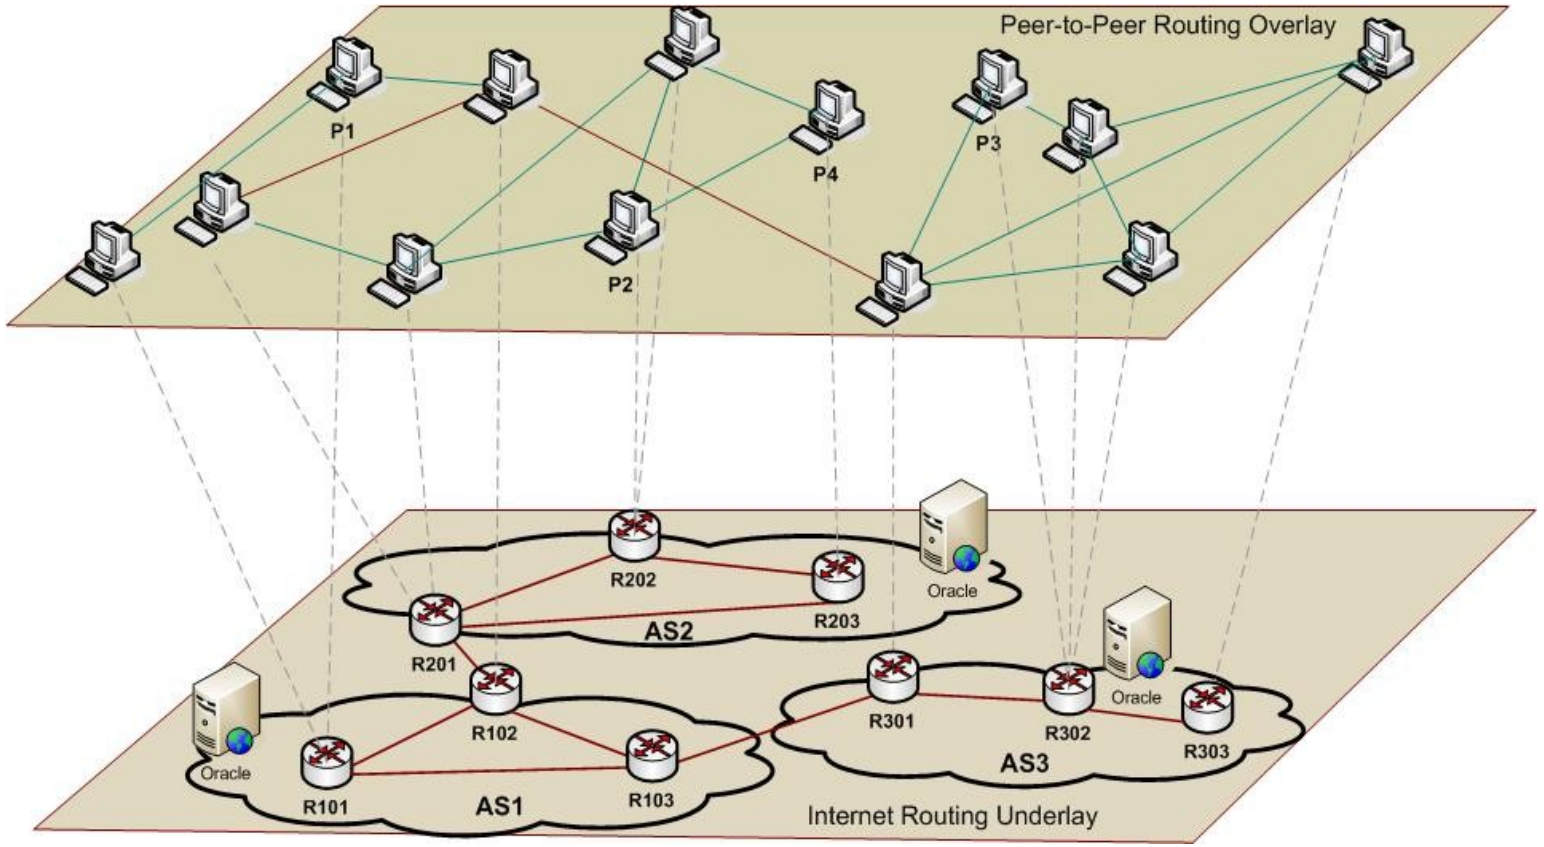

Oracle-based peer selection

→ 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

- Benefits: P2P
  - No need to measure path characteristics
  - Easy to avoid bottlenecks => better performance
- Benefits: ISPs
  - Regains control over traffic
  - Cost savings
  - No legal issues (as no content is cached)

Evaluation

- Impact
  - Topology
  - Congestion
  - 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

Overlay-underlay topology correlation

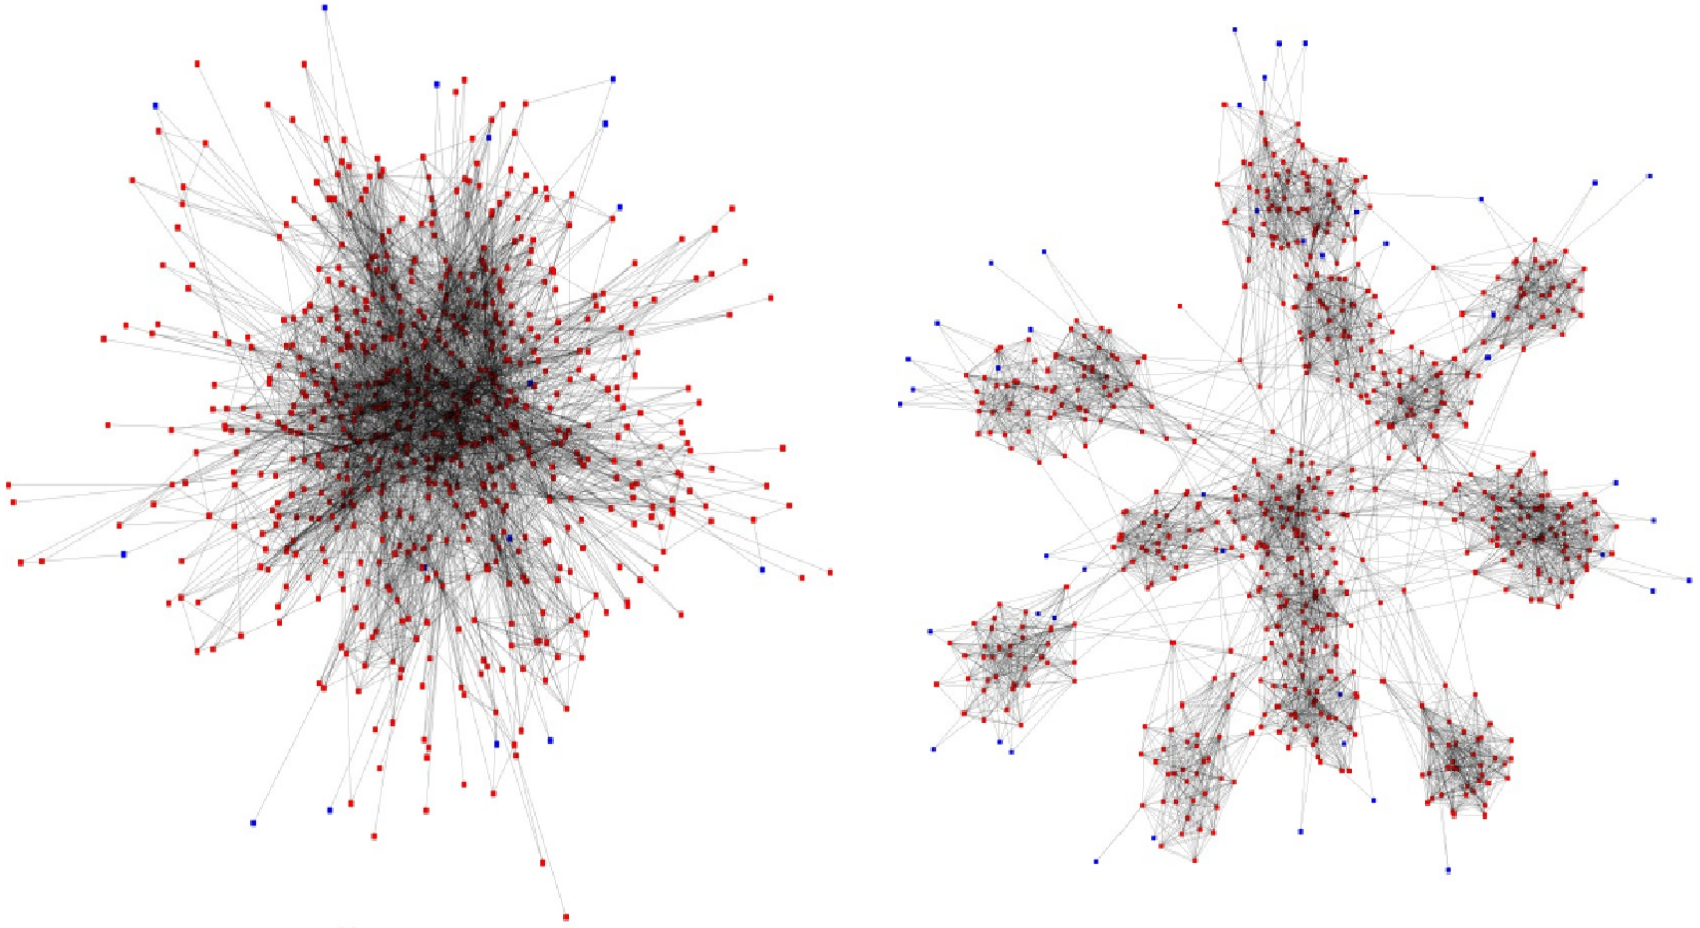

Random vs. biased P2P topology

End-user performance evaluation

- Packet-level simulations
  - Scalable Simulation Framework (SSFNet)
  - Models for IP, TCP, HTTP, BGP, OSPF, etc.
  - Limited to about 700 overlay peers (memory constraints)
- Gnutella-based P2P system
  - Content search via flooding
  - Content exchange via HTTP
- Topologies: several
- User behavioral patterns: several

Topologies: ISP vs. P2P

- Germany
  - 12 ISP's (subset derived from published measurements)
  - 700 peers distributed according to ISP-published customer numbers
- USA
  - 25 Major ISP's (from Rocketfuel)
  - 700 peers distributed in AS's according to city population
- World topologies
  - Sub-sample of measured Internet AS-Topologies: 16 AS's, 700 peers

	Tier1 (# AS / # peers)	**Tier2 (# AS / # peers)**	**Tier3 (# AS / # peers)**
World1	1 / 10	5 / 46	10 / 46
World2	1 / 355	5 / 23	10 / 23
World3	1 / 50	5 / 46	10 / 42

P2P user behavior

- Churn: online/offline duration
  - Pareto and Weibull – close to observed behavior
  - Uniform – base comparison
  - Poisson – reflects worst-case scenario
- Content: type, availability and distribution
  - Constant size (512kB)
  - Pareto and Weibull – typical (many free-riders)
  - Uniform – base comparison
  - Poisson – hypothetical case (most peers sharing)

ISP experience: Intra-AS content

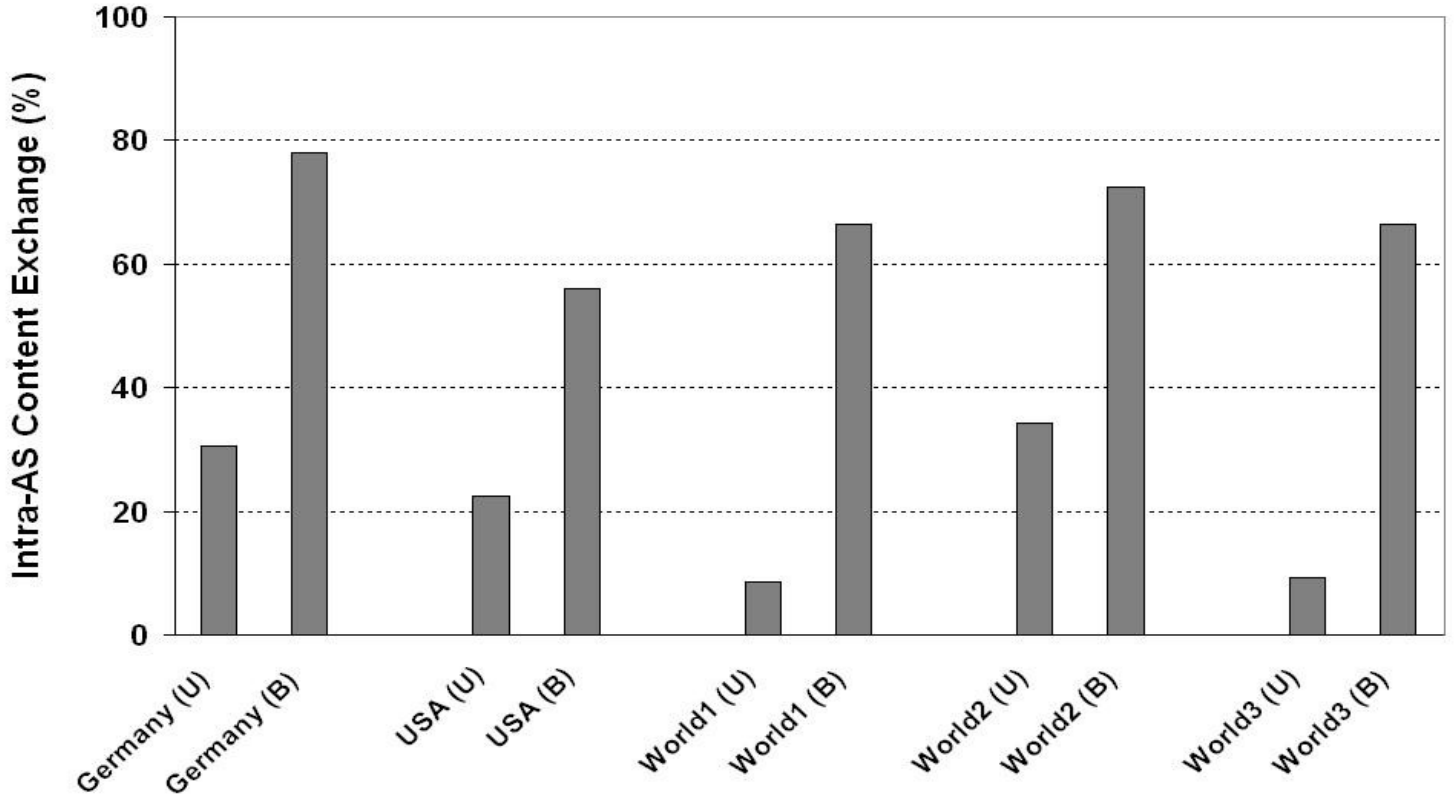

- Content stays within ISPs network
  - Without oracle 10 to 35%
  - With oracle 55 to 80%
- Consistent with Telefonica field trial results for BBC

ISP experience: Intra AS content (2.)

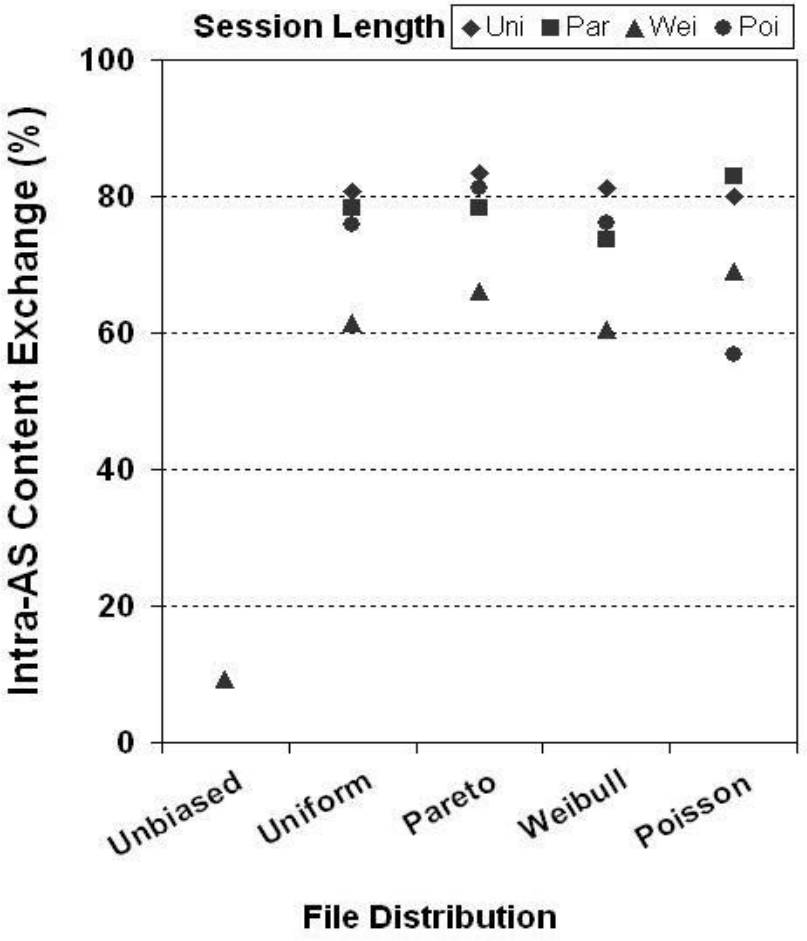

- Content stays within ISPs network

User experience: Download time

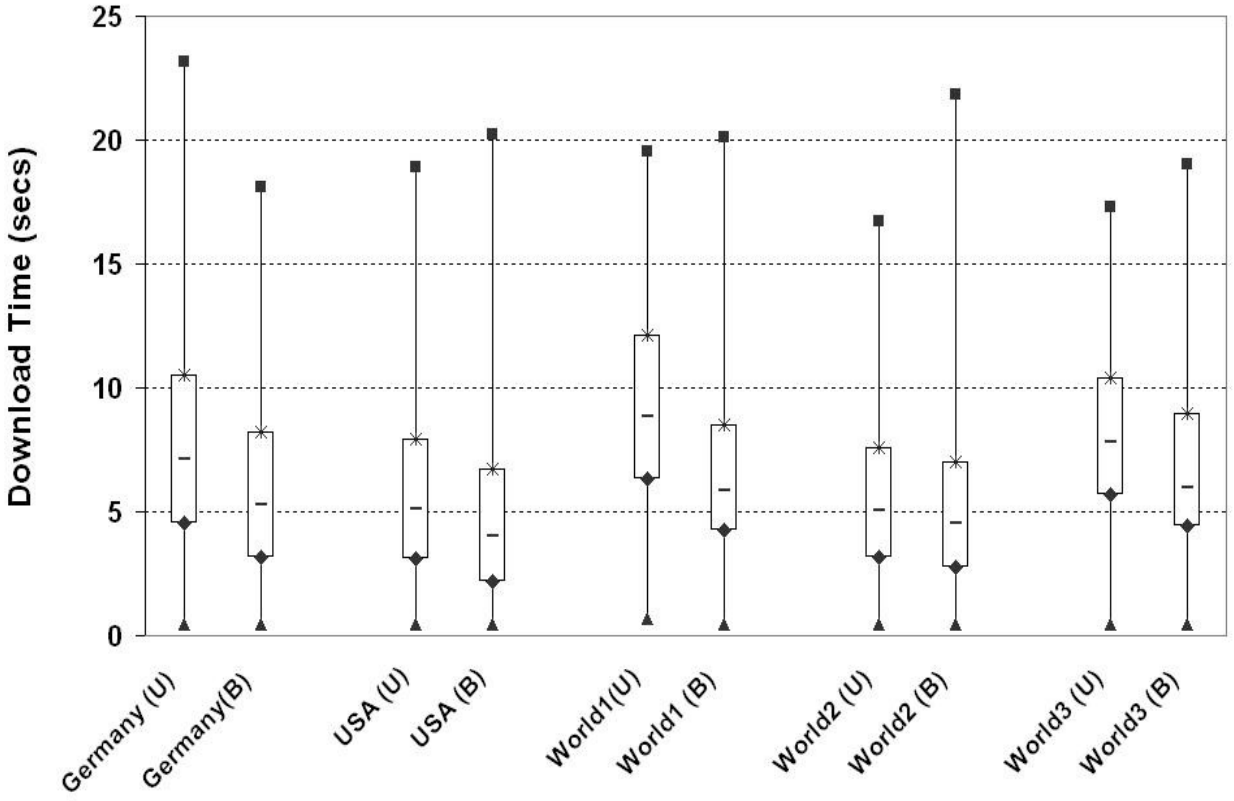

- Mean download time reduction: 1 – 3 secs (16 – 34%)
- Consistent across topologies

User experience: Download time (2.)

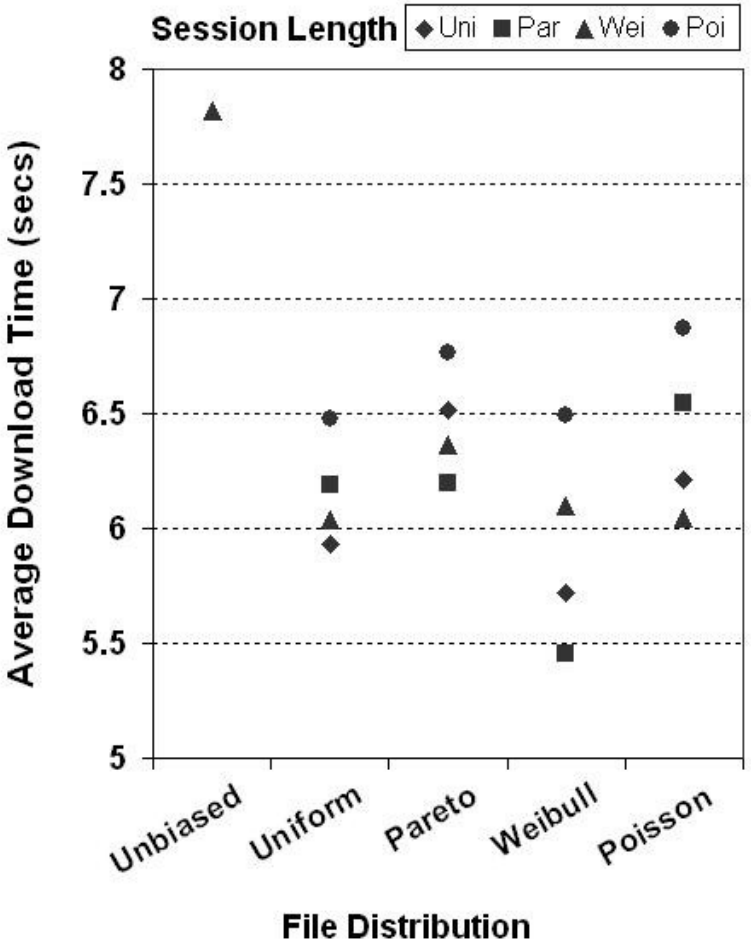

- Reduced mean download time

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
    - Reduces inter-AS traffic => load and costs
    - Traffic distribution close to theoretical optimum
- Benefits
  - ISPs: regain control of network traffic
  - P2P network: sees performance improvements

Upcoming

- Oracle software release
  - Open source implementation will be available (Based on *bind)*
- Software patches for popular P2P clients
  - Gnutella
  - BitTorrent
  - eDonkey
  - P2P TV
- http://www.net.t-labs.tu-berlin.de/isp-p2p/
- Upcoming IETF workshop (May 28th)
  - P2P infrastructure workshop