Technische Universität Berlin · Fakultät für Elektrotechnik und Informatik

Can ISPs and P2P users co-operate for improved performance?

Vinay Aggarwal · Anja Feldmann · Christian Scheideler (TU Munich)

Motivation

- P2P traffic is more than 50% of Internet traffic
- Neighbor selection in P2P systems arbitrary
- P2P routing system independent of Internet
- ISPs in a dilemma
  - P2P users source of revenue
  - Traffic engineering difficult with P2P traffic
- P2P users want performance
  - Measure topology themselves
  - Build topology agnostic of underlay

Gnutella Measurements

- A combination of active and passive techniques to explore and crawl the Gnutella network
- Identify P2P edges at TTL 2, and map them to ASes using BGP tables
- Most P2P traffic crosses AS boundaries, mostly 3-5 AS hops for each P2P session
- Content is often available in proximity, due to language and geographical regions of interest
- Can we introduce locality in P2P topology? How? Will it really benefit?

Solution: Oracle

- Each ISP offers Oracle service
- P2P nodes query it during neighbor selection phase, send list of potential neighbors
- Oracle ranks these by proximity
  - inside/outside of AS, #AS hops in BGP path
  - same PoP/city, expected delay/bandwidth
- P2P node chooses optimal neighbor with help from ISP
- Benefits both P2P users and ISP
- Open service available to all applications

Stages of experiments

- Multi-stage experiments
- Overlay and underlay topology graph experiments
- Simulations with a P2P application on a simulation framework
  - Gnutella protocol implemented in SSFNet
  - P2P protocol weaved at application layer on top of SSFNet routing layers
- Experiments in a Testlab with real routers, switches, computing machines
- Feasibility study in the real Internet using Planetlab

Evaluation Methodology

- Propose metrics for evaluating structural properties of Overlay topology graphs
  - Degree of Overlay nodes
  - Hop count diameter of Overlay
  - AS diameter of Underlay
  - Flow Conductance: measure of congestion in network
- Use realistic Internet topologies for simulations
  - Sample actual AS-level Internet topology derived from BGP data

Preliminary Results

- 300 ASes, 4372 P2P nodes in topology
- Unbiased vs Biased approach (varying Oracle list lengths)
- Overlay graph remains connected
- Degree and mean Path Length in Overlay remain largely unchanged
- Overlay Hop count diameter (=7 hops) and AS diameter (=5 hops) constant
- Large improvement in AS distance and #intra-AS peerings
  - should result in lesser delay, better bandwidth, lesser load on Internet bottlenecks

Degree

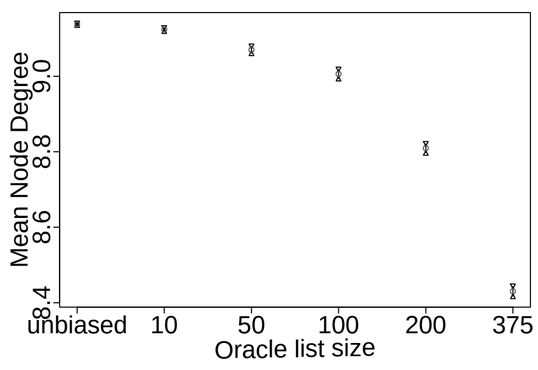

Path Length

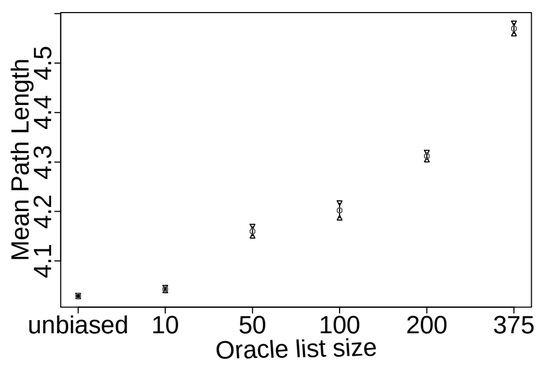

AS Distance

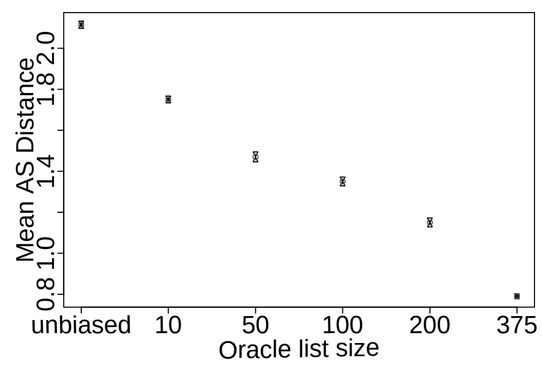

Intra AS peerings

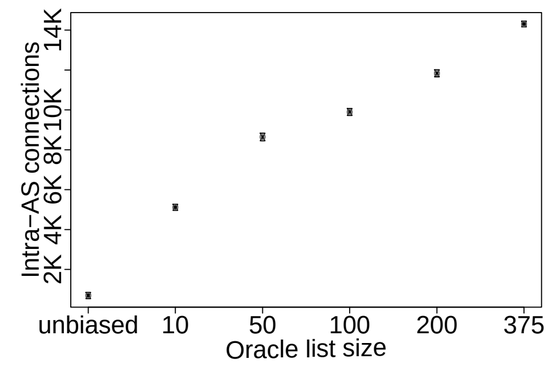

Gnutella SSFNet Results

- 25 ASes, 1000 P2P nodes topology
- Compare unmodified Gnutella with Oracle-biased Gnutella
- Scalability of biased Gnutella network improves significantly
- Ping/Pong msgs reduced to half, Query msgs reduced to a third
- Gnutella topology correlated to Internet AS topology
- Nodes within AS form dense cluster, only few inter-AS peerings
- Overlay network still connected

Gnutella unmodified

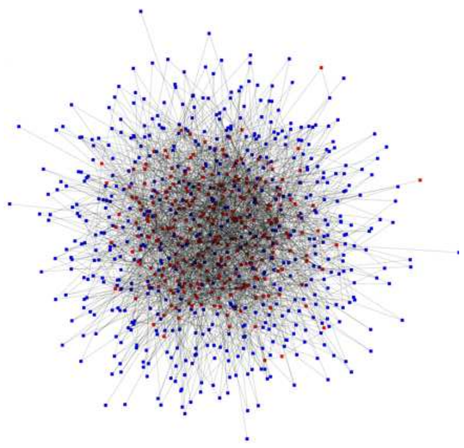

Gnutella with Oracle

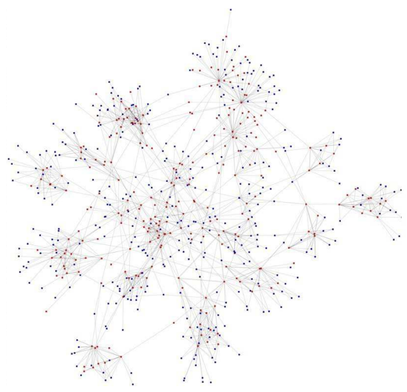

Testlab Experiments

- Experiments in Testlab with real devices
  - 6 routers, 6 switches, 25 computers, 45 Gnutella servents
- Use VLANs and router configuration to simulate a 6 AS topology
- Implement the Oracle on a central machine in the lab
- Gnutella servents query the Oracle with HostCache contents during bootstrapping, then connect to nearest node
- All search Querys satisfied in unbiased Gnutella are also satisfied with Oracle

Future Work

- Larger topologies in experiments at all stages
- Improvement in computation time for Conductance algorithm
- Compare average latency per bit for file transfers in SSFNet
- Use more complex topologies and delay patterns in Testlab
- Reliability and robustness of emerging overlay graphs
- Implementation of Oracle in Planetlab with Gnutella and Bittorrent nodes
- Comparison of Oracle performance with other proximity-based schemes