

MOTIVATION

- P2P traffic comprises > 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, but
 - Traffic engineering difficult with P2P traffic
- P2P users want performance
 - Measure topology themselves, or
 - Build topology agnostic of underlay

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
 - Last hop bandwidth, delay, network congestion
 - Same city/PoP, customer service class
- P2P node chooses optimal neighbor with help from ISP
- Benefits both P2P users and ISP
- Open service available to all applications

OUR APPROACH

- Multi-stage experiments
- Overlay and underlay topology structural 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

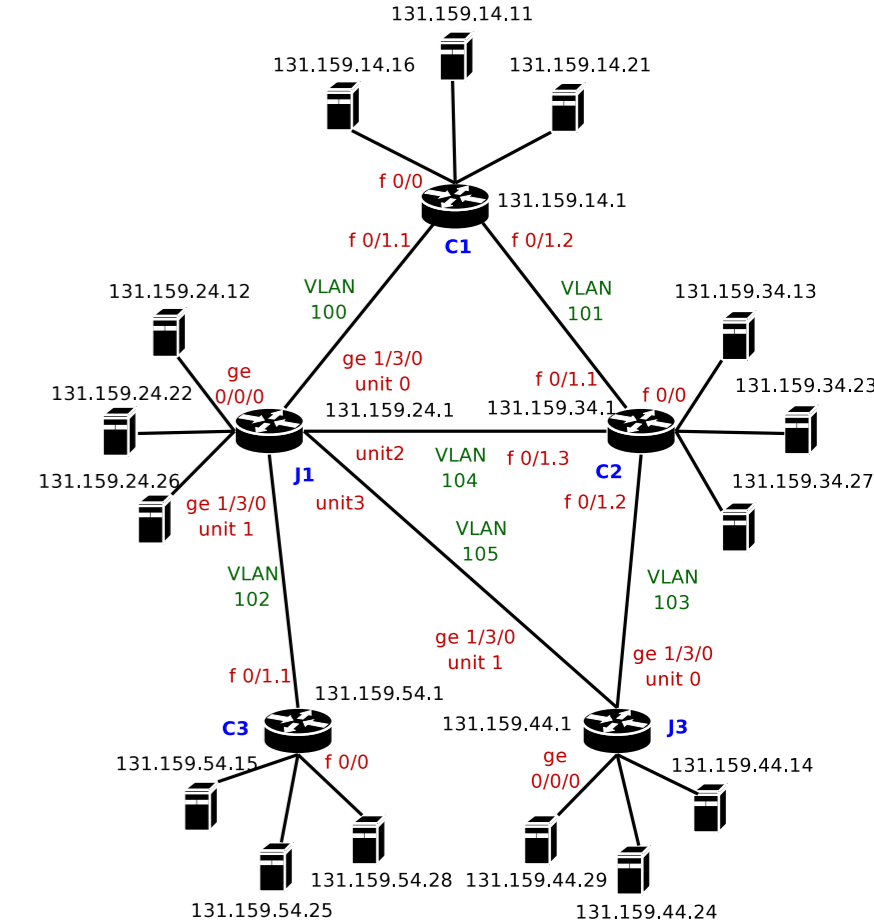
GRAPH RESULTS

- 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
- Flow conductance - a measure of ability to route traffic with low congestion - suffers by factor of at most 2


GNUTELLA SSFNET RESULTS

- 25 ASes, 1000 P2P nodes topology
- Compare unmodified Gnutella with Oracle-biased Gnutella
- Scalability of biased Gnutella network improves significantly
- Negotiation msgs reduced to half, search msgs reduced to a third
- Gnutella topology correlated to Internet AS topology
- Nodes within AS form dense cluster, only few inter-AS peerings
 - Should result in lesser delay, better bandwidth, lesser load on Internet bottlenecks
- Consulting oracle during file-exchange as well as bootstrapping leads to factor 4 increased locality in data transfers


A TESTLAB TOPOLOGY



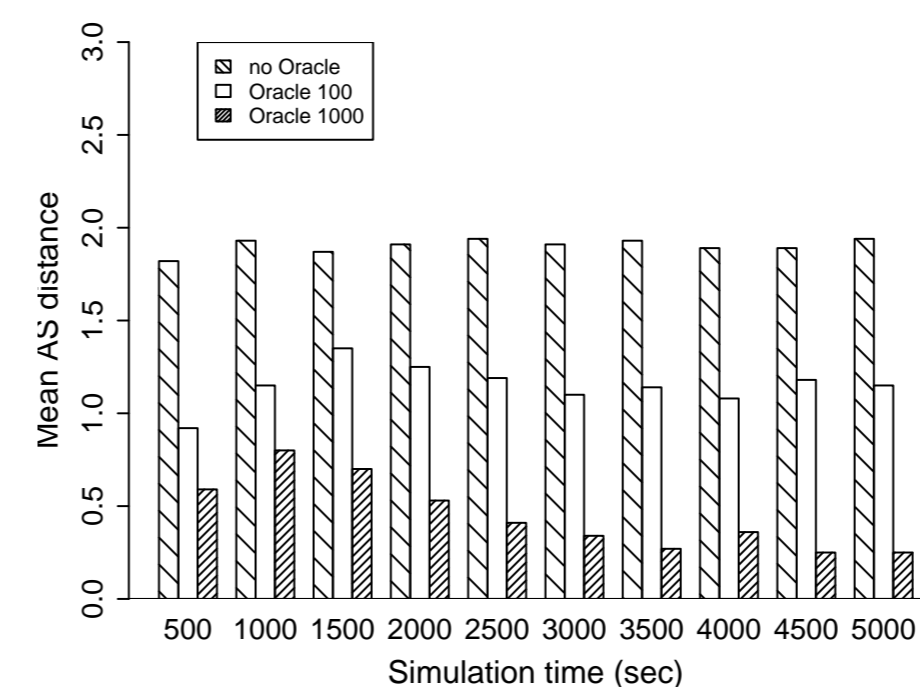
P2P NODE DEGREE




OVERLAY PATH LENGTH




UNDERLAY AS DISTANCE




INTRA AS PEERINGS



P2P TOPOLOGY UNMODIFIED



P2P TOPOLOGY WITH ORACLE



TESTLAB EXPERIMENTS

- Experiments in Testlab with real devices
 - 6 routers, 6 switches, 25 computers, 45 Gnutella servers
- Use VLANs, VTP and router configuration to simulate a 6 AS topology
- Implement the Oracle on a central machine in the lab
- Gnutella nodes consult Oracle with HostCache contents to find nearest node
- Employ uniform and variable file sharing with real content
- All search Querys satisfied in unbiased Gnutella are also satisfied with Oracle
- Negotiation traffic reduced by 50%, improved locality for file exchange

A GLOBAL COORDINATE SYSTEM

- Provides network path property estimates to neighbours/servers within/outside AS
- Concept: ISP-P2P collaboration (oracle) as well as multiple ISP-ISP collaboration
- Uses not only RTT (like existing systems), but also path capacity, available bandwidth, customer service class, AS relationships and routing policies
- No need for network-intensive large scale measurements or complex cartesian coordinate space
- Resistant to churn and malicious nodes, scalable
- Usable by all applications
- Very useful for emerging Internet applications

ONGOING RESEARCH

- Calibrate effect of user behaviour patterns (churn, file sharing, search strings, content type) and AS/P2P topologies on overlay locality using Oracle
- Improved metrics for user experience: download time, delay faced, link bandwidth, quality of search results
- Planetlab experiments for Gnutella (query search analysis) and Bittorrent (locality in content exchange)
- Analysis of congestion within AS using flow conductance, reliability and robustness of overlays
- Compare Oracle with P2P Web-caching and latency-based coordinate systems

PUBLICATIONS (SUBSET)

- Aggarwal, Feldmann, Scheideler: Can ISPs and P2P Users Cooperate for Improved Performance? In ACM SIGCOMM Computer Communication Review, 37(3), July 2007
- Aggarwal, Feldmann, Karrer. An Internet Coordinate System to Enable Collaboration between ISPs and P2P Systems. In ICIN 2007, Bordeaux
- Aggarwal, Feldmann: ISP-aided Biased Query Search for P2P Systems in a Testlab. In ECCS 2007, Dresden
- Aggarwal, Feldmann, Görke, Gaertler, Wagner. Estimating the Neighbour Selection Characteristics of P2P Systems Using Visualization. In Elektronik Journal, 2007
- Aggarwal, Bender, Feldmann, Wichmann. Methodology for Estimating Network Distances of Gnutella Neighbors. In GI INFORMATIK 2004, Ulm