


ISSUES WITH P2P

- P2P traffic comprises > 50% of Internet traffic
 - File-sharing (Bittorrent, eDonkey), VoIP (Skype, GoogleTalk), multimedia streaming
- P2P systems form overlays at application layer
- P2P routing system independent of Internet ⇒ functionality duplication, lack of coordination
- ISPs in a dilemma
 - P2P users source of broadband revenue, but
 - Traffic engineering difficult with P2P traffic
- P2P users want performance
 - Measure topology themselves using RTT ⇒ overhead
 - Build topology agnostic of underlay ⇒ performance loss

TYPICAL PEERINGS AT OVERLAY



- Random/RTT based peer selection → inefficient network resource usage
- Need to maximize intra-AS peerings, along more efficient routes

SOLUTION: COORDINATE SYSTEM

- Maps IP address to n-dim coordinate space
 - Distance in coordinate space reflects network distance
 - Defined as RTT propagation and transmission delay
- Cartesian, hyperbolic, spherical, toroidal coordinate systems proposed
- Reliance on RTT, an unstable metric dependant on network load and churn
- Accuracy, convergence and performance in real-world scenarios not convincing, susceptible to malicious nodes
- No bandwidth or path capacity estimates
- Do not address lack of coordination between routing↔application layers

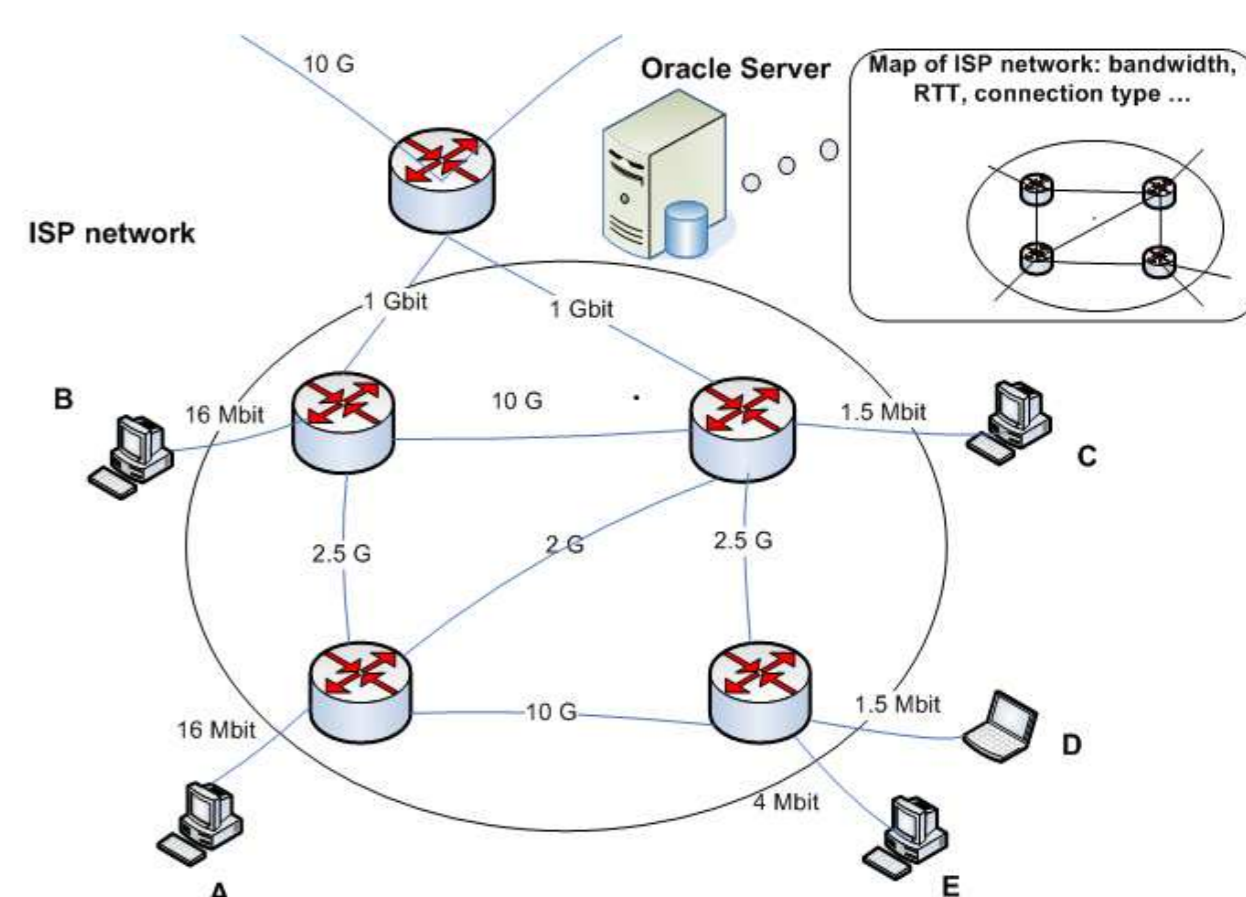
OUR PROPOSAL

- Global coordinate system based on ISP-P2P and ISP-ISP collaboration
- Concept: ISP knows its network
 - Node: last-hop bandwidth, estimated delay, geography, service class, load
 - Routing policy, OSPF/BGP metrics
- Given 2 IPs, easily determine path characteristics without measurement in own network
- Collaborate with other ISPs for global service
- ISP provides coordinate system service to applications, win-win for all
- Improves routing↔application layers and ISP↔P2P coordination

SINGLE-ISP COORD SYSTEM

- Based on ISP-P2P collaboration
- Each ISP offers Oracle service
- P2P nodes query it during neighbour selection, send list of potential peers
- Oracle ranks these by proximity
 - Inside network, last-hop bandwidth, geographical location (cityPoP), AS hops
- ISP-aided optimal P2P neighbour selection
 - Not just RTT, but combination of chosen metrics
- Simple, general, open service for all
- Run as Web server or UDP service, on single/replicated servers, operating on semi-static database

HOW THE ORACLE WORKS



- A consults oracle, peers with B instead of C (better bandwidth)
- E consults oracle, connects to D instead of C (same PoP)

ADVANTAGES FOR ISP/P2P

- ISPs regain control of network traffic
- Traffic across ISP boundaries reduced ⇒ immense cost savings
- Better QoS to other applications, improved service to customers
- P2P measurement overhead eliminated
- Avoid high-latency paths and bottlenecks at inter-ISP transit/peering links
 - Reduction in download time for users
- Search queries generate comparable response

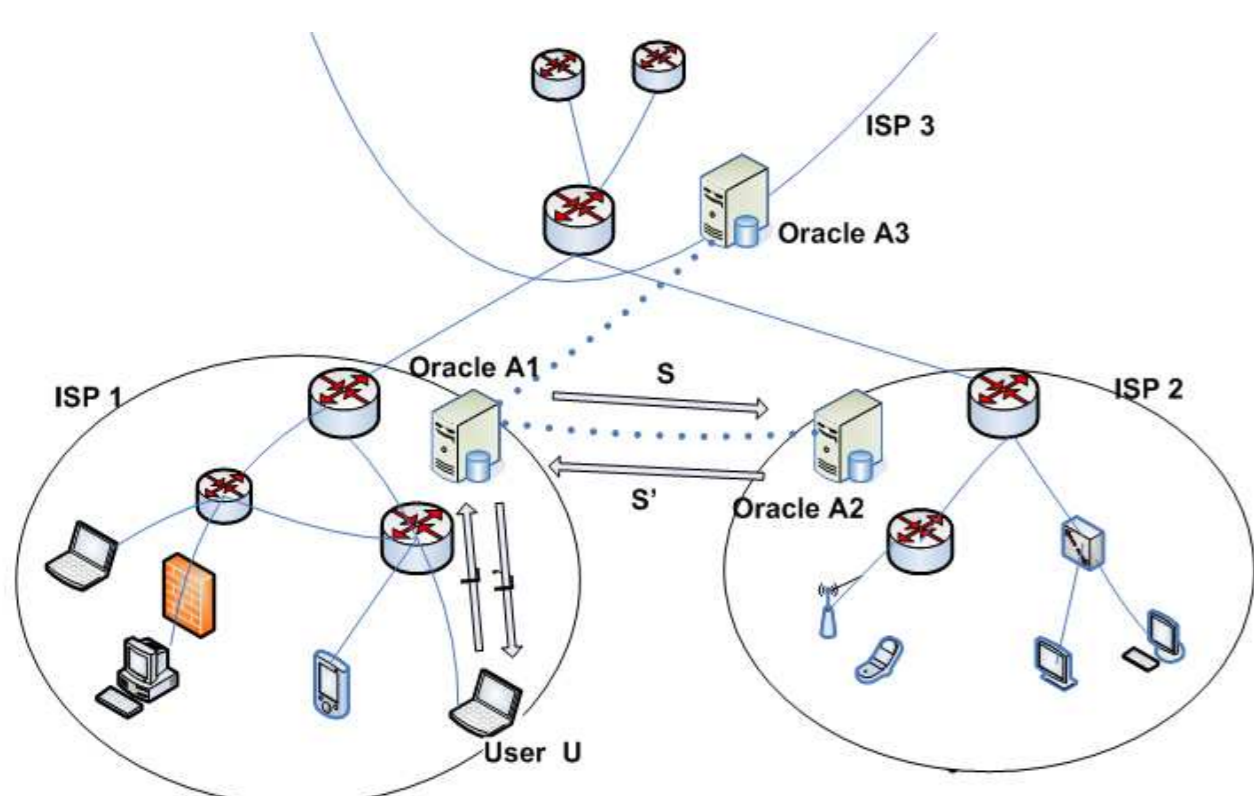
PRIVACY ISSUES

- Ranking algorithm confidential
- Output is ranked list or classification, actual stats not revealed
- List can be anonymized or dynamically altered for better QoS
- Amount of network info revealed can be reverse-engineered
- No caching of content ⇒ reprieve from legal issues
- Resistant to churn, improves scalability of P2P apps, less susceptible to malicious nodes

GLOBAL COORDINATE SYSTEM

- Based on multiple-ISP collaboration
- Concept: ISP knows
 - routing policy to neighbouring ISPs: customer, provider or peer AS
 - their IP address range
 - backbone link capacities in vicinity
- Rank nodes within own network (oracle), segregate nodes of other ISPs based on BGP metrics, routing policy, AS-hop distance
- Fine-tune list by contacting oracle from other ISPs, they sort their nodes and return sorted list

MULTIPLE ISP COMMUNICATION



- User U sends list L to oracle A1. A1 segregates subset S that belongs to ISP2, and sends to A2. A2 sorts S and returns S' to A1. A1 later combines all subsets, and returns fully-sorted L' to U.

SUMMARY

- Novel global coordinate system based on collaboration between
 - User applications and ISPs
 - Multiple ISPs
- Uses metrics like path capacity, available bandwidth, customer service class, AS relationships, routing policy (and also RTT)
- Scalable, resistant to churn, less susceptible to malicious nodes
- Removes need for large-scale Internet reverse-engineering
- Will prove very useful for the emerging breed of Internet applications, e.g., Web X.0

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: Locality-Aware P2P Query Search with ISP Collaboration. In Networks and Heterogeneous Media Journal, AIMS, Nov 2007.
- Aggarwal, Feldmann, Görke, Gaertler, Wagner. Estimating the Neighbour Selection Characteristics of P2P Systems Using Visualization. In Telenor Telektronikk Journal, 2007