Simple Economic Management Approaches of Overlay Traffic in Heterogeneous Internet Topologies

European Seventh Framework STREP FP7-2007-ICT-216259

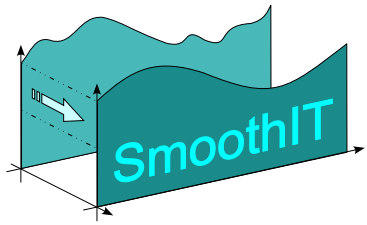

Initial Results on Simulations and Experiments with Economic Traffic Management (ETM)

UZH, DOCOMO, TUD, AUEB, PrimeTel, AGH, ICOM, UniWue, TID

Burkhard Stiller
(on behalf of SmoothIT partners)

COST IS0605 WG4
Ljubljana, Slovenia
May 21, 2010

Overview

- Introduction
  - Triple Win
  - SIS
  - ETM
- Simulations
  - BGP-Loc
  - IOP
  - Dyn-Loc
  - Modeling Locality Games
- External Trials
- Summary and Conclusions

Introduction

Basics and Motivation

- Use of economic mechanisms for controlling, managing network traffic of overlays at early stages:

  ⇨ **Economic Traffic Management (ETM)**

- SmoothIT results show already that such mechanisms do have the important property of scalability and effectiveness!
  - Situation-dependent ETMs lead to a more efficient network operation
  - ETMs generate a higher value (QoE) for its customers.

- In managing the traffic created and routed through their networks, today's ISPs are offered by SmoothIT methodologies suitable for modern traffic/service profiles
  - *E.g.*, peer-to-peer traffic is treated according to ETM approaches.
  - Applicable to traffic of different P2P applications

Triple Win

- Management of **overlay networks** based on a collaboration between the overlay provider and the network (underlay) provider in support of the user
  - Cost and investment recovery for operators
- Incentives for operators
  - Monetary: reduce overlay traffic and inter-domain traffic
  - Traffic management: less congested links, better performance
  - Reputation: keep customers, distinguish from other operators
- Incentives for overlay providers
  - Performance: Active role in traffic mgmt increases service quality
  - Reputation: increased user base due to better performing services
- Incentives for user
  - Performance: Increased service quality, *e.g.*, reliability, RTT, BW
  - Monetary: lower price for network access

SmoothIT's Technical Approach in a Nutshell

- Facilitate collaboration between network providers and P2P applications/users
- SIS Architecture
  - Client/Server architecture
  - Modular components
  - Suitable for all ETMs
  - Prototypical implementation evaluated in internal test-bed
- 13 ETM mechanisms defined 3 selected for investigations
  - Locality awareness (BGP-Loc)
  - ISP-owned overlay entities (IoP)
  - Highly Active Peers (HAP)

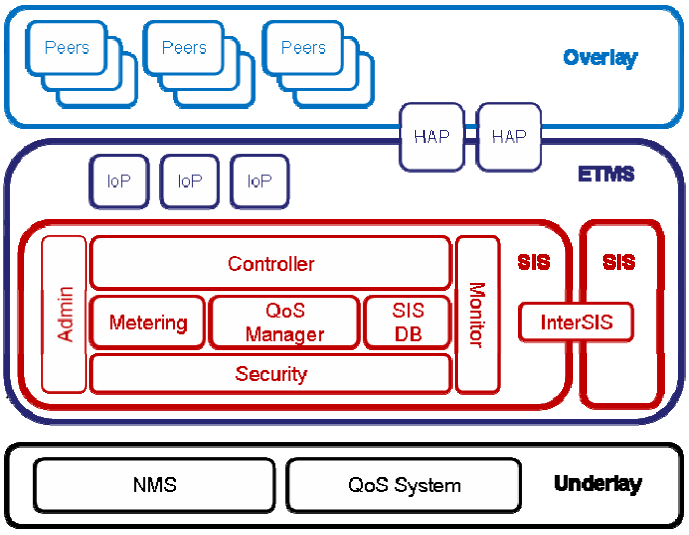

SmoothIT Information Service (SIS)

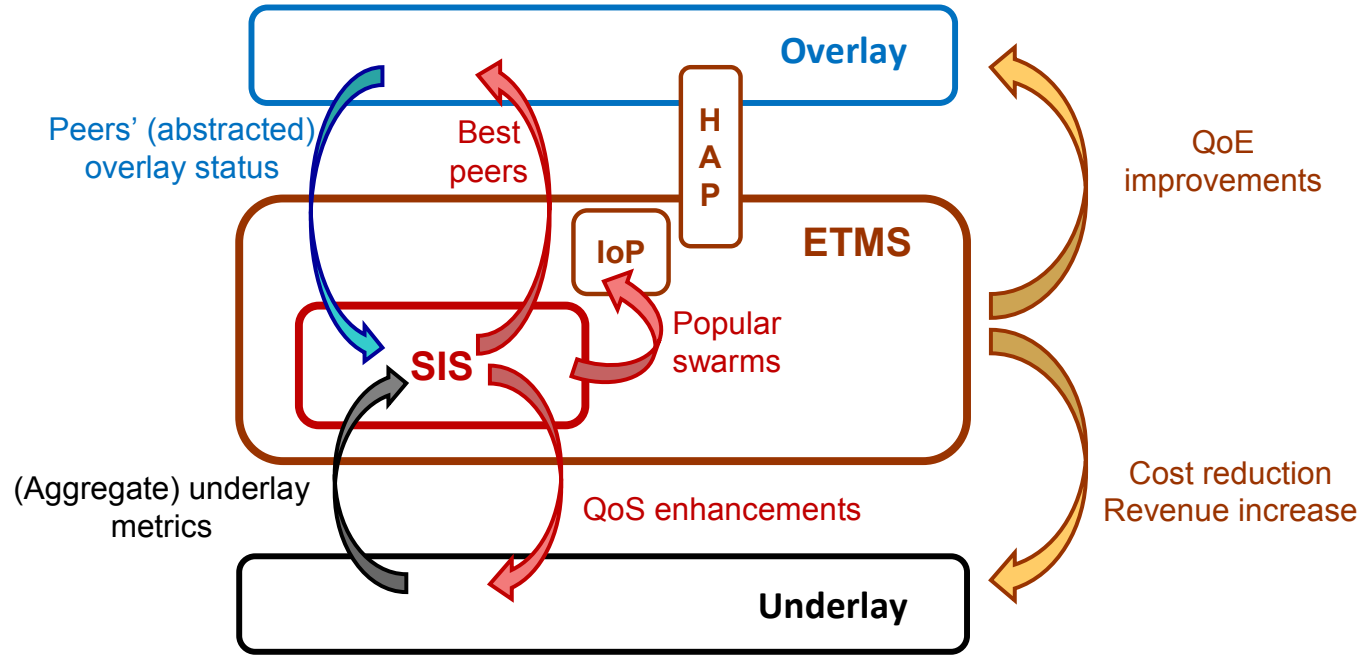

SIS is the core of the ETM System (ETMS), which can lead to *TripleWin*.

ETM Classification and Synergies

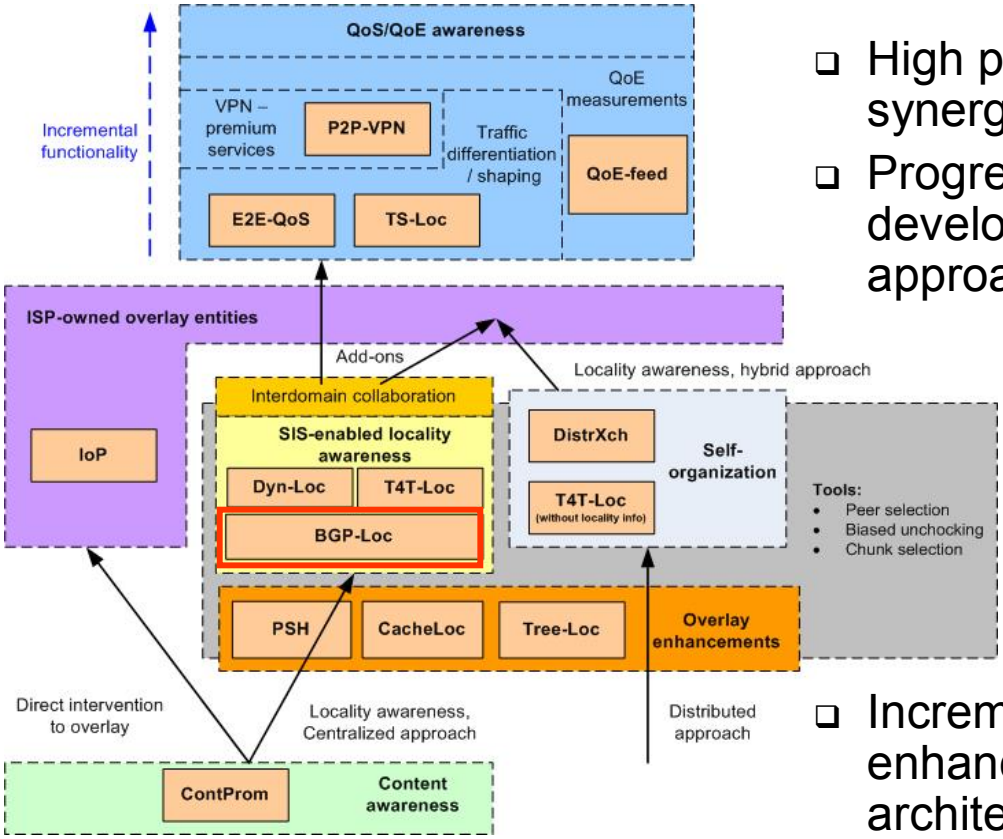

- High potential for synergies
- Progressive development of approaches possible
- Incremental enhancement of architecture possible

ETM Assessment Methodology

- Assessment aims at evaluating TripleWin

- Innovative assessment methodology for the ETM mechanisms is required:
  1. Different cost metric per player, no “total cost” optimization
  2. Wide variety of assessment scenarios
  3. Assessment of benefit of individual user (rather than on the average) with coupled simulations
  4. Game-type assessment, at ISP and user levels
  5. Study of dynamics, taking “Win” of users into account

Simulations

Family of Simulated Topologies

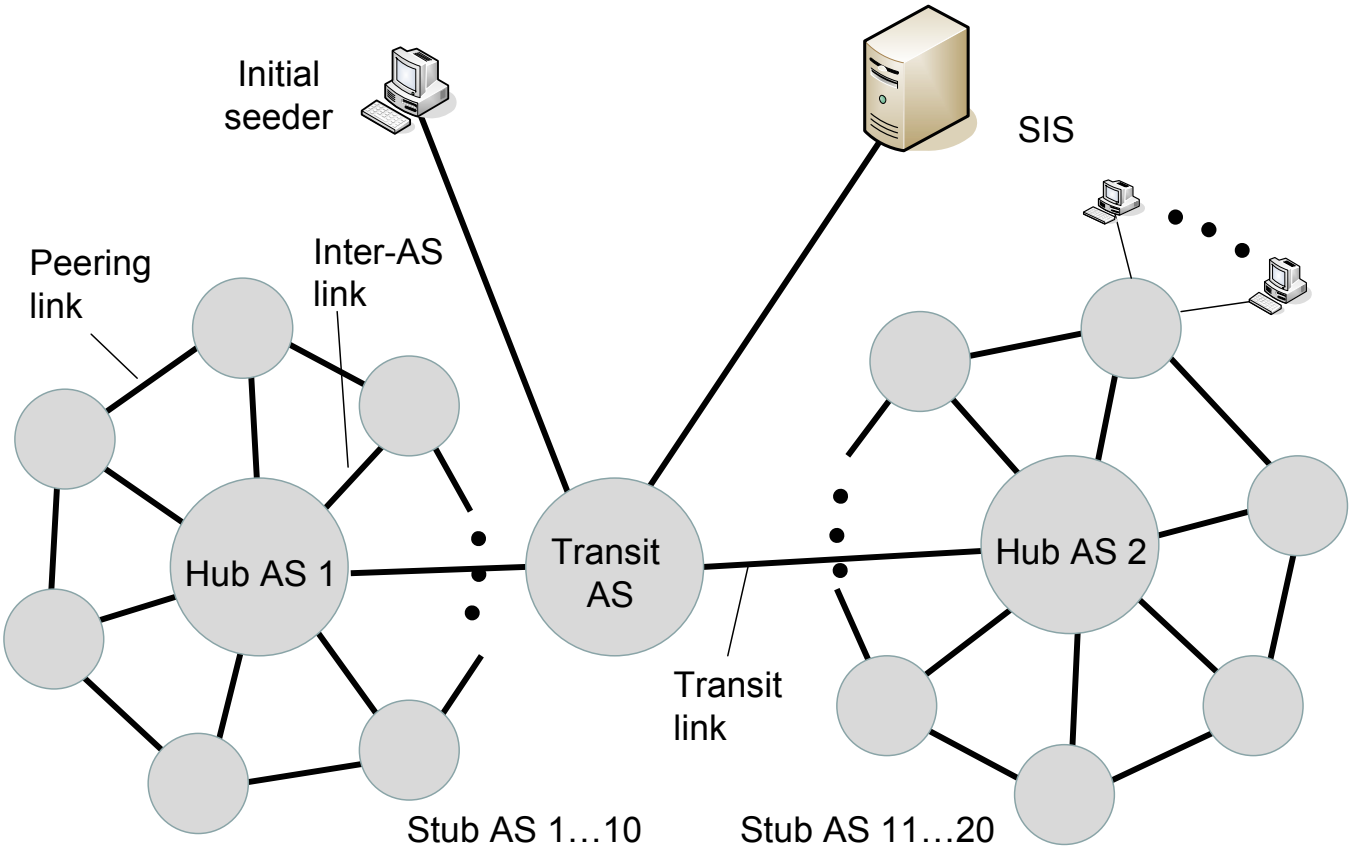

BGP-based Locality (BGP-Loc)

- *Motivation*: ISP has underlay information that could be used by peers in overlay decision making
- *Approach*: ISP provides ratings for peers based on BGP-related information;
  peers combine both overlay and underlay ratings
- *Impact*: Provide 'proximity' metric to peers in order to distinguish candidate overlay neighbors
- *Innovation*: General SIS rating framework, specific SIS and client intelligence

BGP-Loc: Specification

1. Peer send to SIS a list of IP addresses of candidate peers

2. The SIS rates the received peers according to
   - whether they belong to the local AS (highest ratings)
   - BGP metrics, employed for non-locals (lower ratings)
     - Local preference
     - AS path length
     - MED

3. The querying peer receives a sorted list of candidate peers and keeps as neighbors those with highest ratings from SIS

BGP-Loc: Client Intelligence

- Biased Neighbor Selection (BNS)
  - Top SIS-rated peers constitute $x\%$ of the neighbor set
  - Rest $(100-x)\%$ randomly selected

- Biased Optimistic Unchoking (BOU)
  - Unchoking probability is higher for highly-rated peers
  - Unchoking combines SIS (=underlay) and overlay ratings
  - *Innovative idea*

- Combination of BNS and BOU is the most effective approach

BGP-Loc: Simulation Results (1)

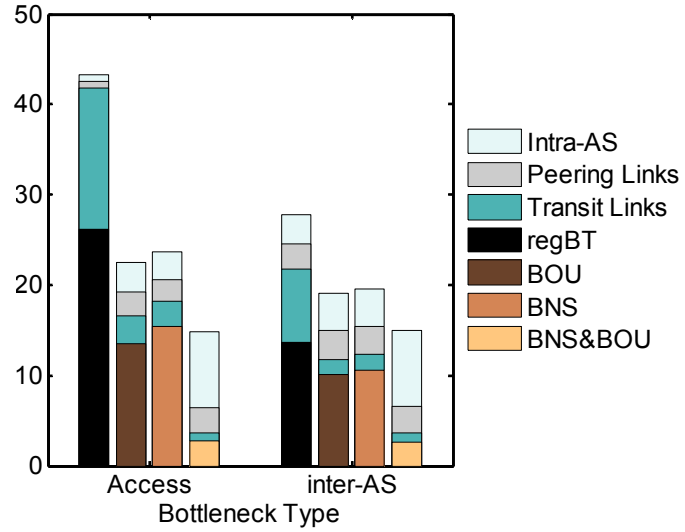

Mean traffic/sec (MB/sec)

BitTorrent: Bottleneck Types

Mean download times (min)

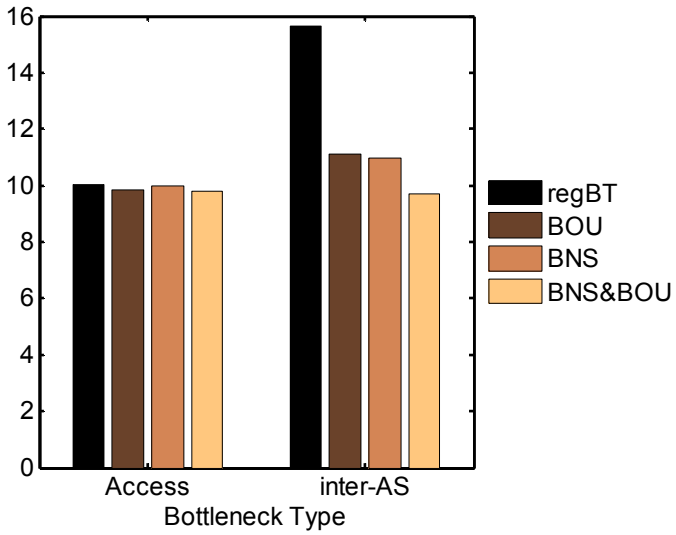

BGP-Loc: Simulation Results (2)

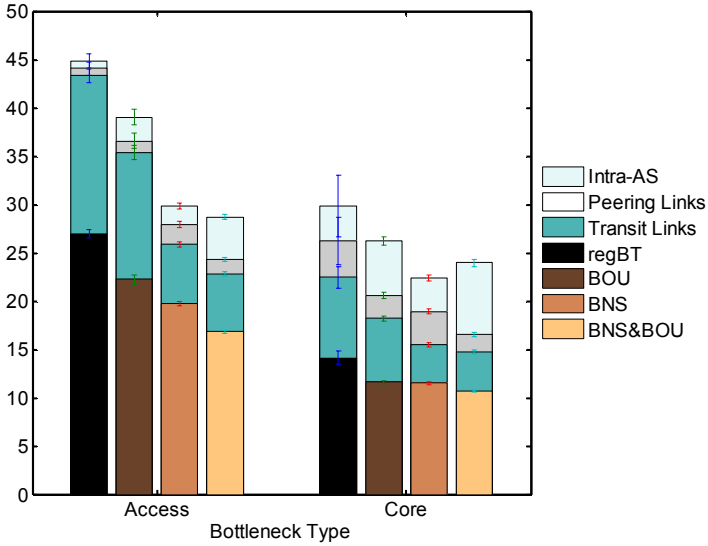

Mean traffic/sec (MB/sec))

Tribler: Bottleneck Types

Mean stalling times (sec)

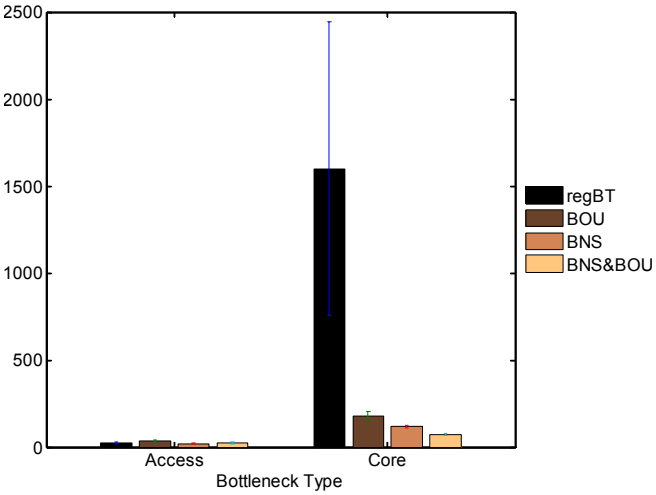

BGP-Loc: Simulation Results (3)

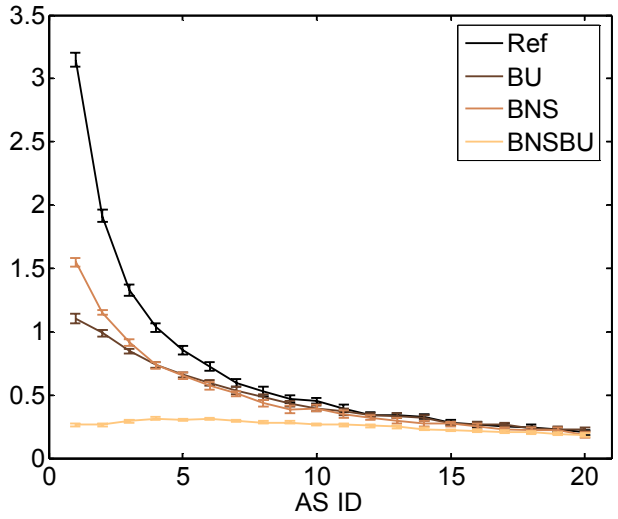

Mean upload traffic/sec for different ASes with different populations (MB/sec)

Heterogeneous peer distribution

Mean download times (min)

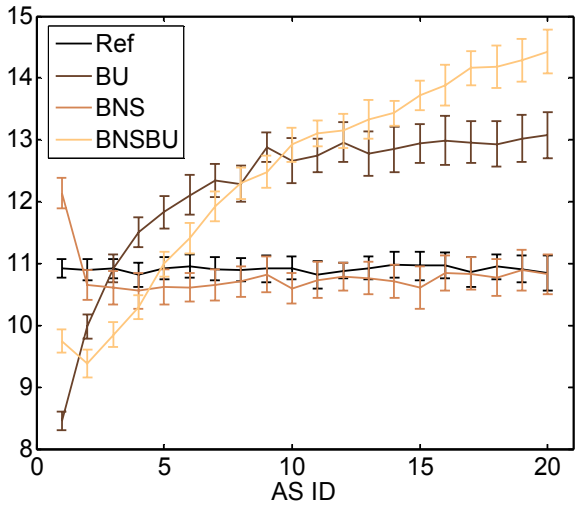

BGP-Loc: Evaluation Summary

- Bottleneck: access links → Win–No-lose
  - Download times remain unaffected
  - Savings in inter-AS bandwidth are achieved
- Bottleneck: inter-domain links → Win–Win
  - Still some inter-AS bandwidth can be saved
  - Download/stalling times can be improved significantly

- The efficiency of locality promotion is higher in ASes having larger fractions of the swarm
- The performance improvement further increases with the fraction of locality-promoting peers

Insertion of ISP-owned Peers (IoP)

- *Motivation*: Sole locality may not improve peers' performance. We can exploit overlay functioning to localize traffic and achieve Win-Win
- *Approach*: Insert an ISP-owned peer (overlay entity) provisioned with higher access capacity
- *Impact*: Improvement of peers' performance and reduction of inbound traffic
- *Innovation*: Transparency, no interception required. Variety of policies

Specification of IoP

- Identification of possible types of collaboration between ISP-owned Peer and SIS
- Three innovative modules:
  1. *SIS-enabled swarm selection*
     - Communication protocol between IoP and SIS
     - Communication protocol extension between peers and SIS
  2. *Bandwidth allocation among selected swarms*
     - Follows swarm selection rating
  3. *Unchoking policy*
     - IoP can unchoke only local peers in the seeding phase

Evaluation of IoP

- Plain IoP:
  - Significant improvement in users' performance
  - Inbound traffic reduction / Outbound traffic increase
    - Further impact under the unchoking policy

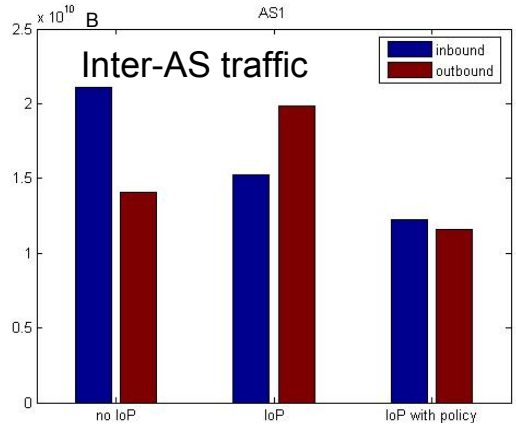

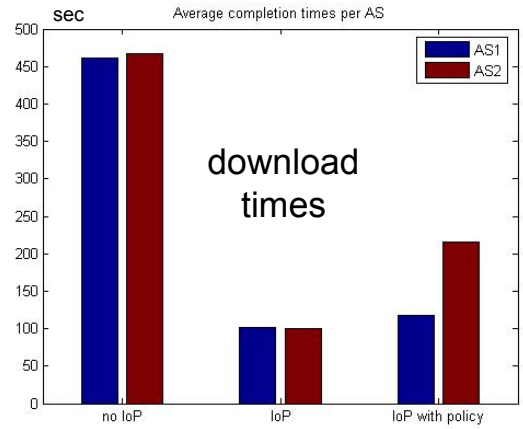

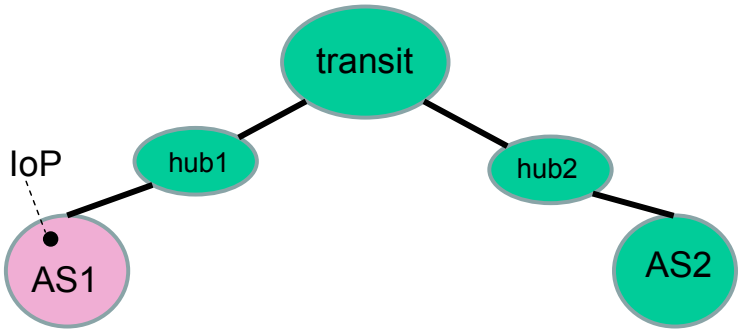

Dynamic Locality (Dyn-Loc)

- *Motivation*: Unconditional localization may be harmful, while continuous limiting of inter-domain is not the only way to reduce charges
- *Approach*: Monitor running 95th percentile and "switch" locality promotion ON/OFF in 5min slots
- *Impact*: Decrease inter-domain traffic for busiest $x\%$ slots, to reduce charges with less intervention
- *Innovation*: Dealing directly with the 95th percentile. Use of locality only when necessary

Dynamics

- *Assumption*: Users or ISPs can predict or verify "on-line" that an ETM mechanism is to their benefit
  - e.g. trial-and-error on-line bypassing of the mechanism
  - Players not benefiting can bypass the mechanism
- Investigate dynamics due to players that do bypass
  - Do the other players still benefit?
  - How would the others respond?
- *Approximation*: Assume perfectly informed players, that know a priori what is best to them

Dynamics with BGP-Loc

- With full locality, only peers in largest ASes benefit
- *Question*: Do the other ISPs have an incentive to employ BGP-Loc?
- *Experiment*: Only the top *K* ISPs promote locality, where *K* = *3,4,5,6,7,8,9*
- *Results*:
  - Even higher degree of unfairness w.r.t. download times
  - Only the largest ISPs save bandwidth
  - Smaller ISPs are forced to employ BGP-Loc
- *Conclusion*: Large user groups may be able to impose BGP-Loc to the swarm

Results (1)

Locality in the top 3 ASes

Upload bandwidth per AS (MB/sec)

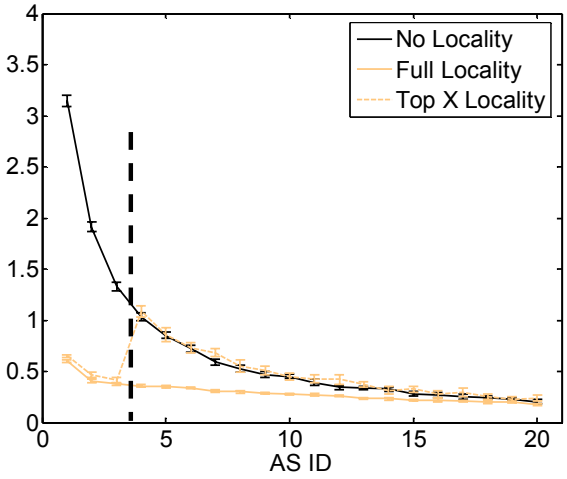

Download times per AS (min)

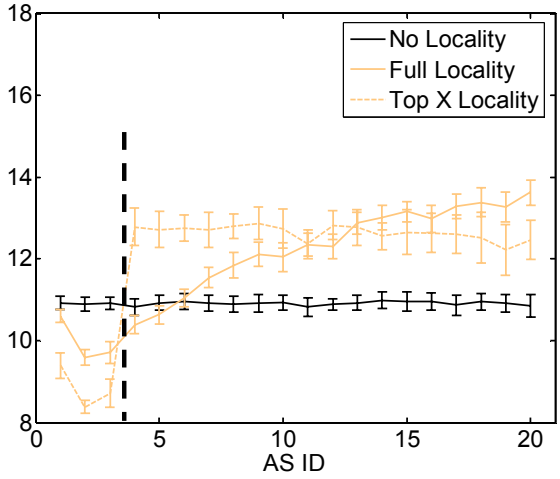

Results (2)

Locality in the top 4 ASes

Upload bandwidth per AS (MB/sec)

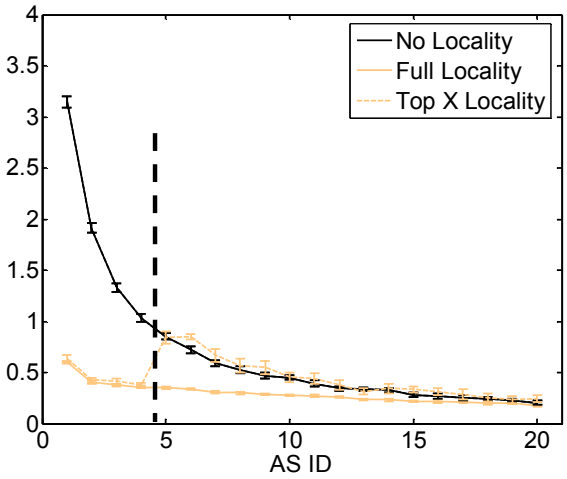

Download times per AS (min)

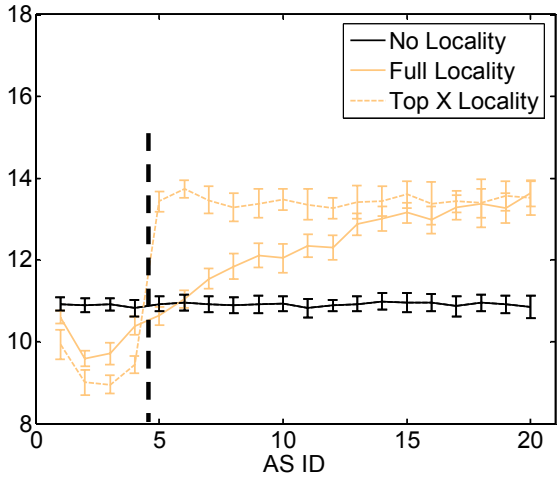

Results (3)

Locality in the top 5 ASes

Upload bandwidth per AS (MB/sec)

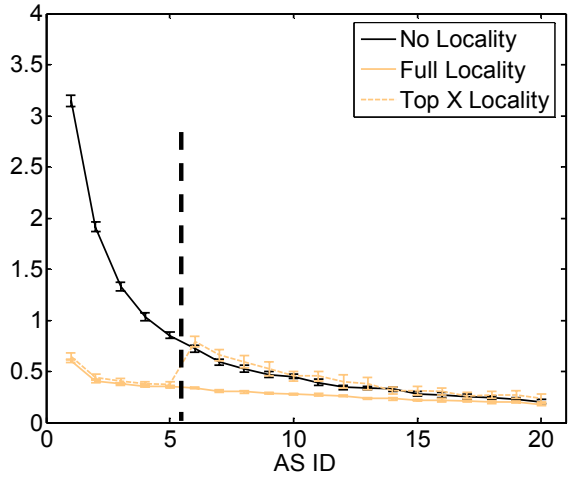

Download times per AS (min)

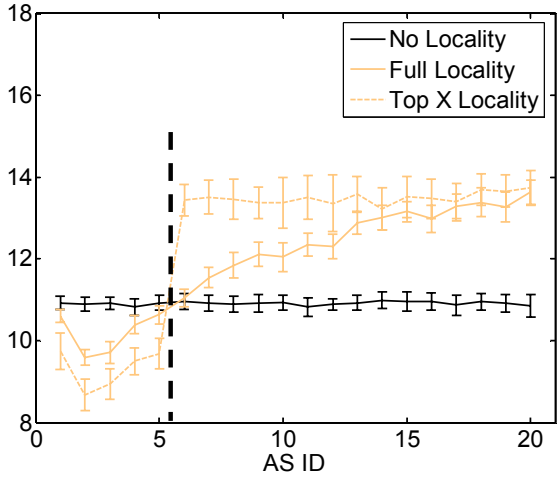

Results (4)

Locality in the top 6 ASes

Upload bandwidth per AS (MB/sec)

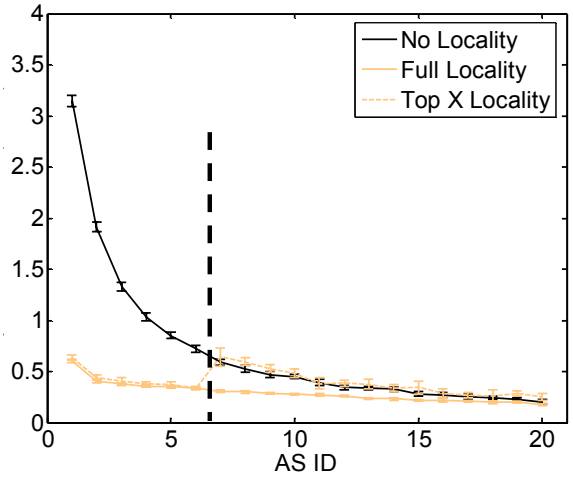

Download times per AS (min)

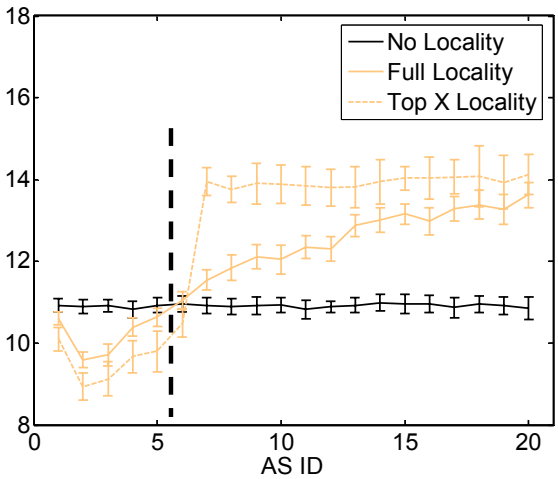

Results (5)

Locality in the top 7 ASes

Upload bandwidth per AS (MB/sec)

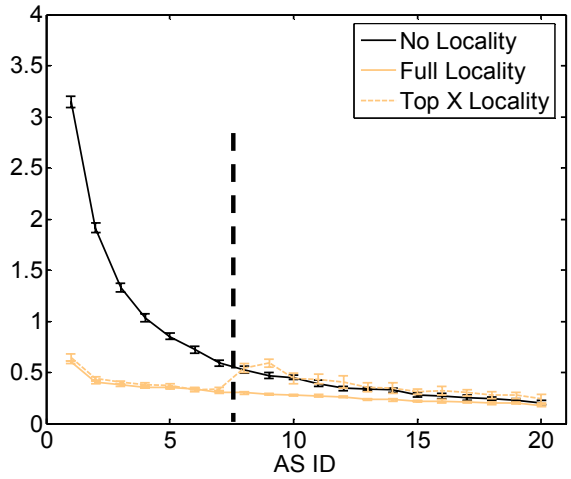

Download times per AS (min)

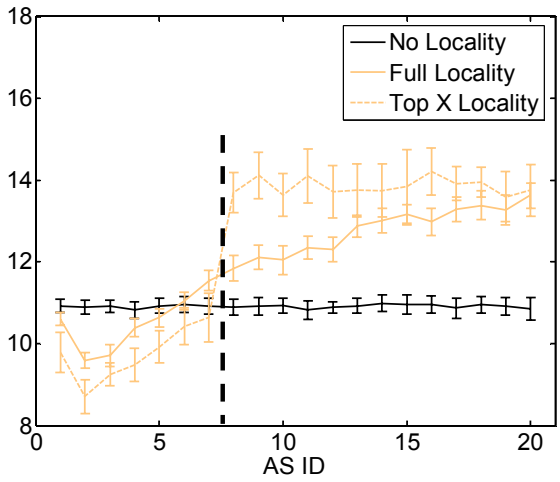

Results (6)

Locality in the top 8 ASes

Upload bandwidth per AS (MB/sec)

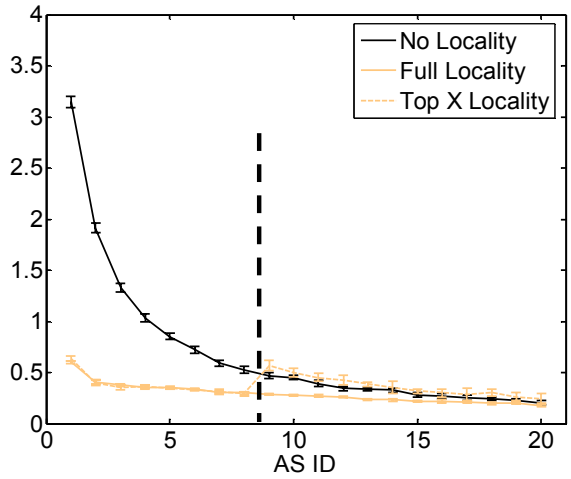

Download times per AS (min)

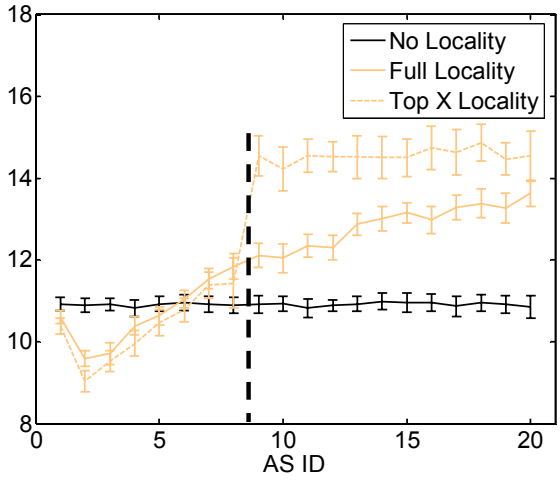

Results (7)

Locality in the top 9 ASes

Upload bandwidth per AS (MB/sec)

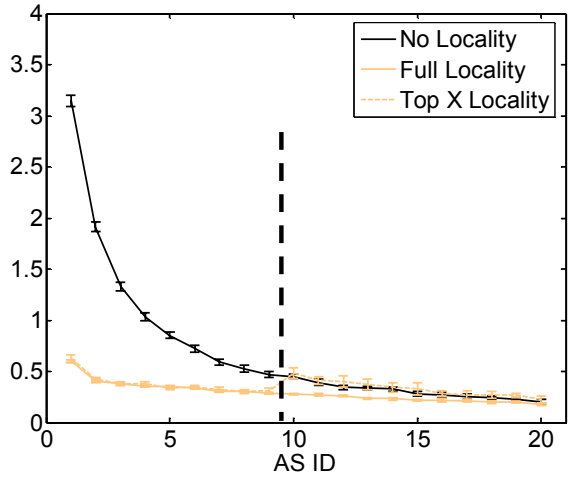

Download times per AS (min)

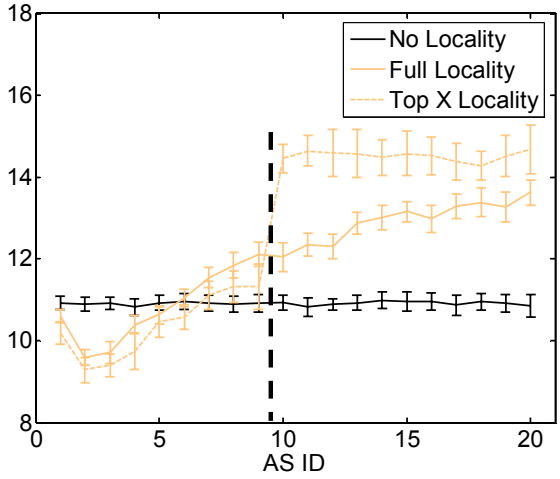

Modeling Locality Games

- What happens when two “neighbor” ISPs have each to decide to promote locality or not?
  - “Neighbor” ISPs: customers of the same higher-tier ISP
- How does an ISP respond to its neighbor’s actions?
  - Is locality promotion always the best reply?
- Assumptions:
  - Single swarm, Tit-4-Tat only
  - Demand (in chunks) always be fulfilled
  - Unchoking probability independent of “contention”
  - Consider only ISPs cannot control the effect of locality (realistic)

External Trials

Integration for External Trials

- Development of network-specific interfaces to Network Management System and routers
- Ensure appropriate user and system interfaces are in place from PrimeTel's side

- Adapt client
  - Prepackaged with correct settings
  - Fully automatic setup and configuration (Out-of-the-Box)
  - Automated customer feedback
- Prepare client with correct behavior simulation for GLab

External Trial — Integration

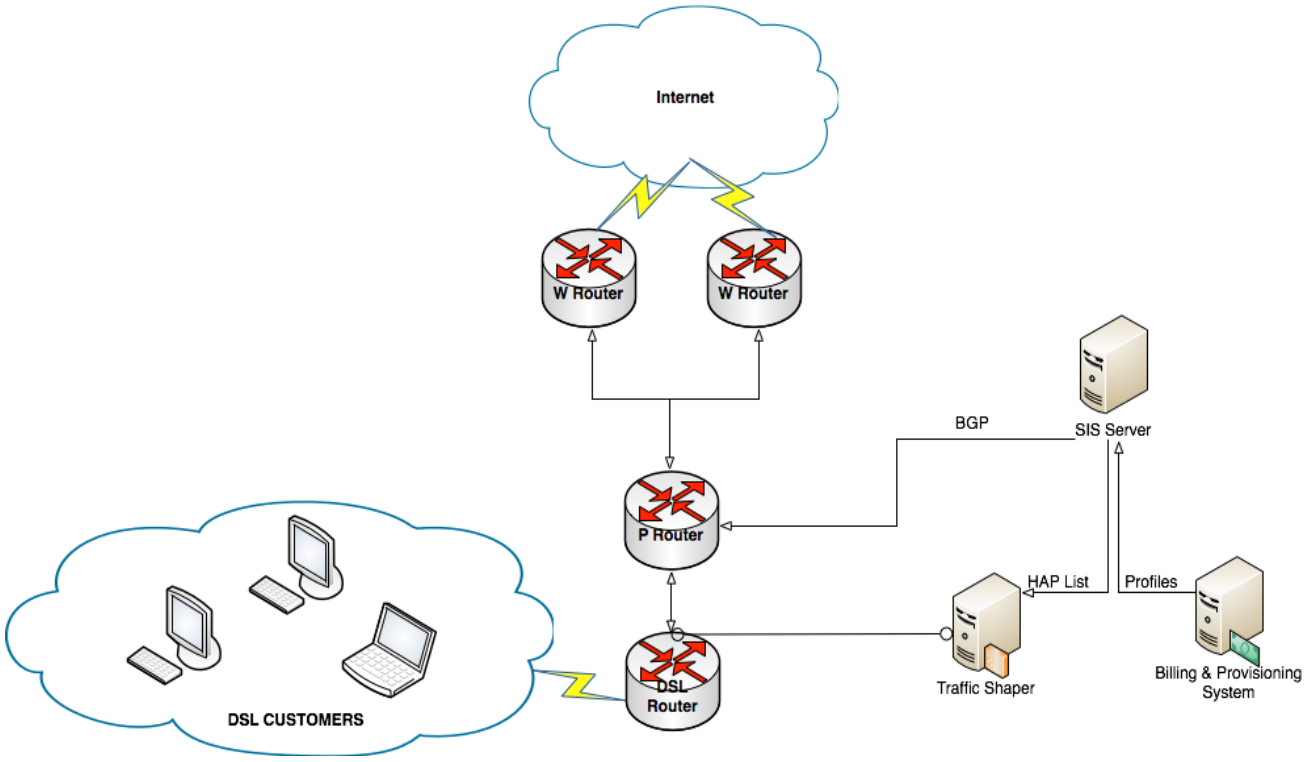

External Trial — Topology

- 2 ASes
- PrimeTel running SIS
- G-Lab with automated headless peers

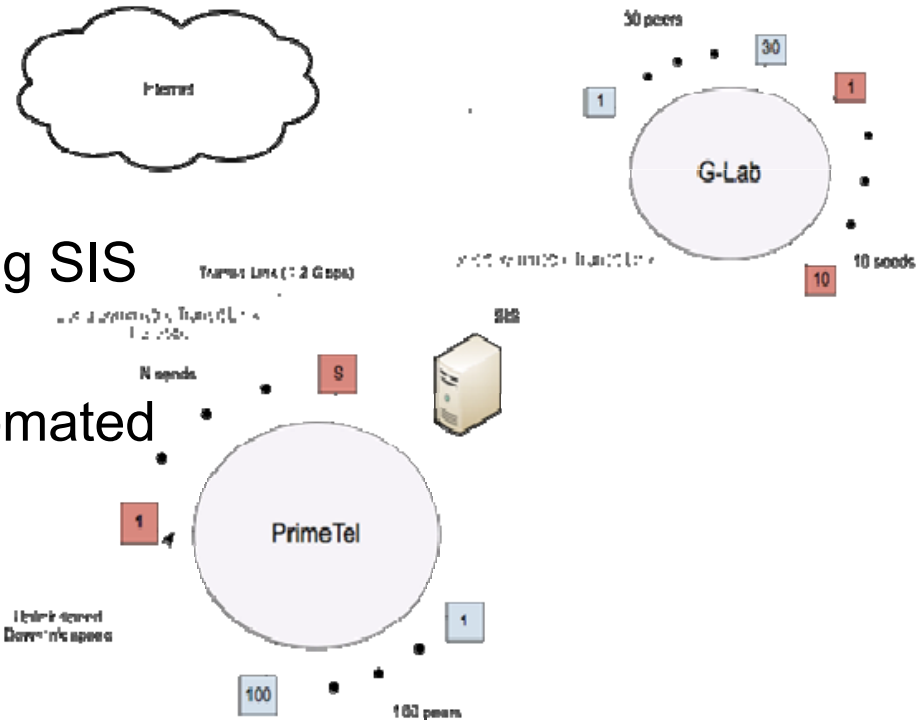

External Trial — Experiments

- Three experiments are being designed:
  - BGPLoc + IoP
  - BGPLoc + HAP
  - Reference
- Experiments to come in June-August 2010

Summary and Conclusions

Summary and Conclusions

- Dedicated management of overlay traffic is necessary
  - Due to smoothing large amounts of overlay traffic
  - Due to the minimization of high(er) costs for ISPs
- Detailed requirements analysis undertaken
- SmoothIT architectural design in progress
  - SmoothIT Information Service (SIS)
    - Deployed in networks of ISPs
    - Provides information to overlay applications
    - Optimizes traffic and achieves the Triple Win situation
- Prototypical SIS implementation undertaken
- First sets of simulative evaluations in place
- SmoothIT participates in ALTO (2 drafts) and covers socio-economic aspects of communications, too

Thank you for your attention!

Thanks to all SmoothIT's project partners:

UZH, DOCOMO, TUD, AUEB, PrimeTel, AGH, ICOM, UniWue, TID