MultiCache: an incrementally deployable overlay architecture for information-centric networking

Konstantinos Katsaros, George Xylomenos and George C. Polyzos
Mobile Multimedia Laboratory
Athens University of Economics and Business

IEEE INFOCOM Work-in-Progress (WiP), San Diego , CA, USA, March 2010

Outline

- Motivation
- Design objectives
- MultiCache architecture
- Functionality overview
- Performance evaluation
- Ongoing and Future work
- Conclusions

Motivation

- Internet model: end-to-end principle
  - Need to resolve a specific end-host to retrieve data
- Internet use: information-centric
  - “Anyone” that can provide the required data is fine
    - E.g. P2P, cloud computing, etc.
- Arbitrary overlay content delivery structures, ignoring:
  - Network topology
  - Data location
  - Data popularity
- Inefficient use of network resources
  - E.g. 70% percent of an AS ingress traffic could be avoided in BitTorrent[1]

[1] T. Karagiannis, P. Rodriguez, and K. Papagiannaki, “Should internet service providers fear peer-assisted content distribution?” in Proc. Of the Internet Measurement Conference, 2005, pp. 63–76.

Design objectives

- Efficient use of network resources
  - Resource sharing mechanisms: multicast, caching
- Scalability
  - Unlimited size of the information domain
- Usage model simplification
  - End hosts not engaging in translating what to where
- Facilitated deployment of new functionality
  - Clean-slate requires replacing existing functionality
    - E.g. PSIRP Project
  - Network layer available solutions (e.g., IP Multicast)
    - Practically not available
    - Not easy to deploy gradually
    - Difficult group management
  - Targeting at an overlay architecture…

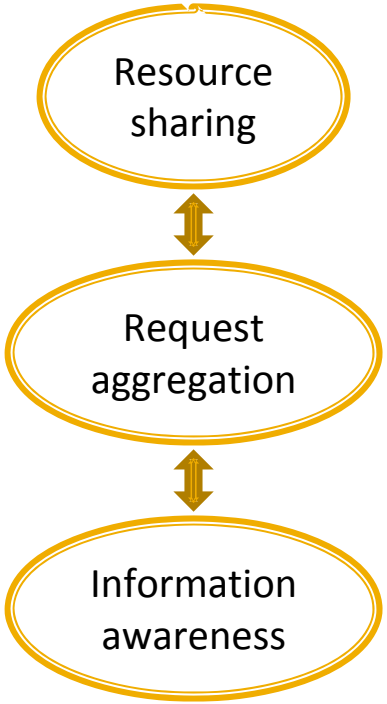

MultiCache architecture

- Deploying Overlay Access Routers (OARs) inside access networks
  - Gradual deployment is feasible
- Providing overlay multicast
  - Based on Scribe over Pastry
    - Scalable
    - Adaptive to physical topology
- Acting as caches
  - Multiple cache locations
  - Close to end-hosts
- Proxy-ing end host access to the overlay
  - Facilitating group management
  - Proxy OAR designated during network attachment

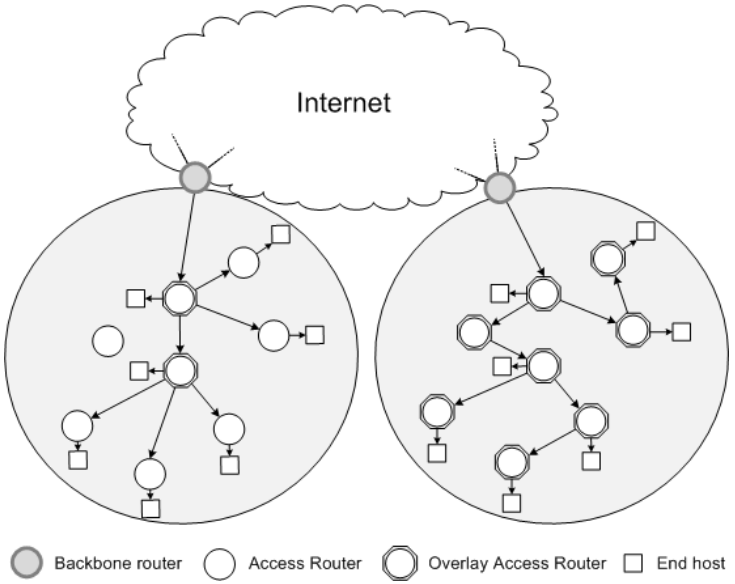

Functionality overview

- Currently focusing on content distribution
- Overlay multicast brings content from its origin
- Caching
  - Data @ proxy OARs, i.e., multicast tree leafs
  - Forwarding state @ Forwarding OARs
- Anycasting cache requests
  - Localizing traffic inside sub-trees
  - Taking advantage of Pastry's locality properties
- Unicasting cached data
  - Reducing stretch…
- Content fragmentation
  - Parallelizing transfers
  - Enabling partial caching

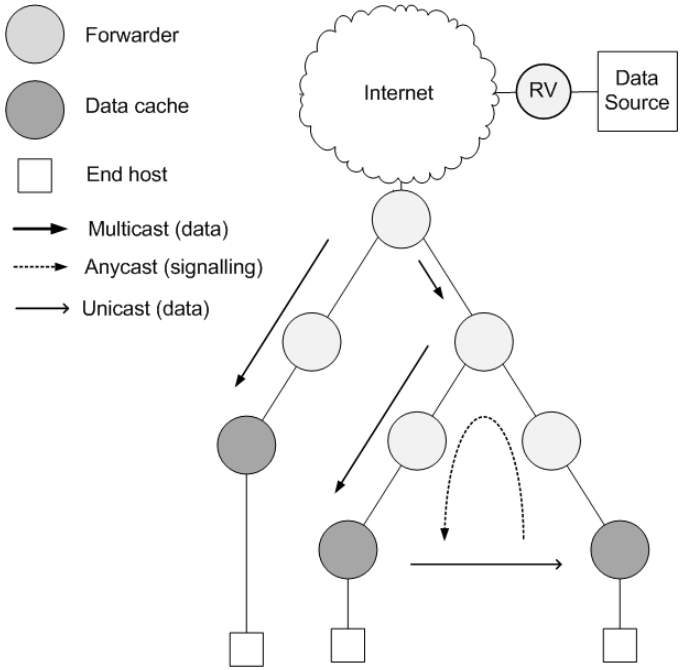

Performance evaluation

- Cache replacement scheme not available at that time
  - Infinite cache sizes assumed
  - Upper bound on potential benefits
- Simulation based evaluation
  - OMNeT++, OverSim
- Comparing against BitTorrent [1]
- Scenario
  - Single 256MB file
  - 100 end hosts
  - GT-ITM topology
    - 1200 access routers in 25 AS's
- Metrics:
  - Download Time (sec)
  - Egress Interdomain traffic (MB)
  - Intra-domain link stress

[1] K. Katsaros, V. P. Kemerlis, C. Stais and G. Xylomenos, "A BitTorrent Module for the OMNeT++ Simulator," Proc. 17th Annual Meeting of the IEEE International Symposium on Modeling, Analysis and Simulation of Computer and Telecommunication Systems (MASCOTS), London, UK, September 2009, pp. 361-370

MultiCache vs. BitTorrent: traffic

Infinite cache size, 100 users, single 256MB file, 25 AS's, 40 Access Routers / AS, 100% deployment

- Egress **inter-domain** traffic reduction: ↓60%
- Same gains for **intra-domain** traffic
- Traffic localized due to cache deployment
  - Forwarding mechanisms favors the discovery of near-by cache locations
  - Average block hop count:
    - BitTorrent: 8.86 hops
    - MultiCache: 4.61 hops
- Reducing operational costs for network operators
- Exchanging transmission with storage

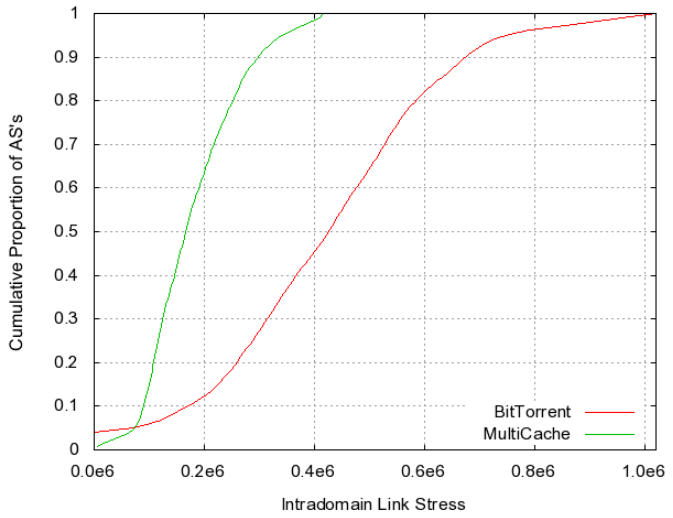

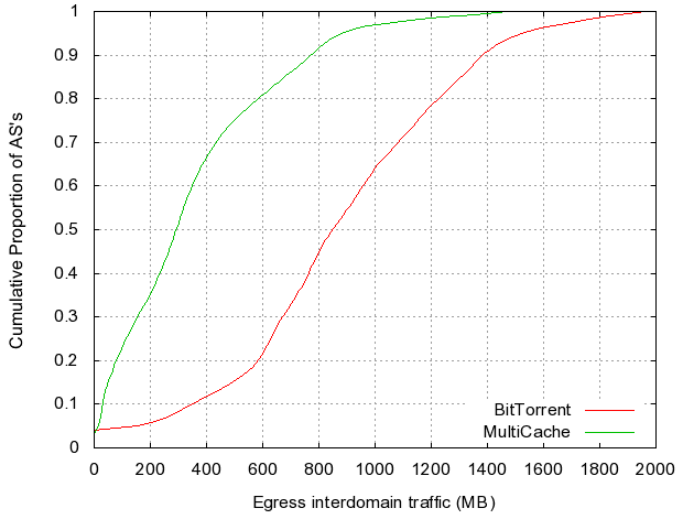

MultiCache vs. BitTorrent: download time

Infinite cache size, 100 users, single 256MB file, 25 AS's, 40 Access Routers / AS, 100% deployment

- Huge download time reduction:
  ↓ 90%
- This is only an upper limit
  - Infinite cache sizes guarantee cache availability
- But:
  - Localized traffic favors faster downloads
  - End hosts do not search for content
    - Direct consequence of the information-centric model
  - No peer uplink bottlenecks

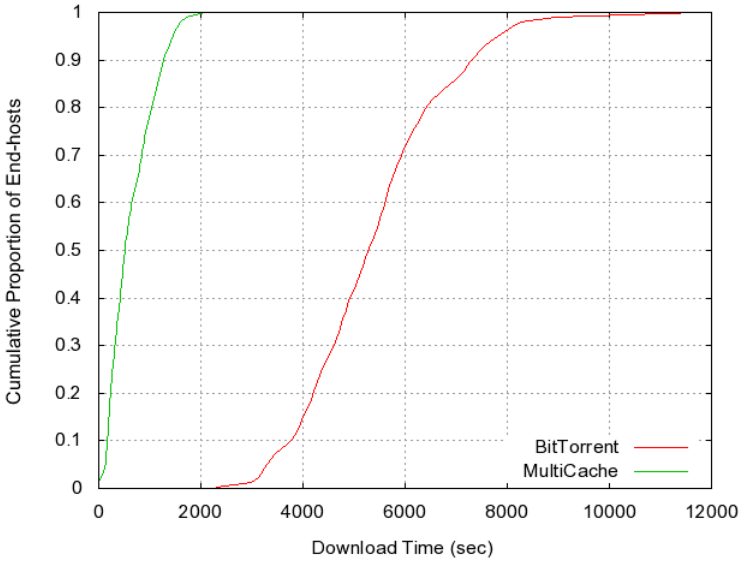

Ongoing & Future Work

- Caching scheme completed
  - Global Internet Symposium, this Friday, March 19th
    - MultiCache takes advantage of multiple cache locations
- Ongoing work:
  - Comparing full fledged MultiCache with BitTorrent
    - Multi-torrent scenarios
- Future work:
  - Gaining control of inter-domain cache provision
    - Establishment of peering relationships between domains
      - Expressing them in a new “proximity” metric for Pastry
      - Employing *Canonical* Pastry

Conclusions

- Resource sharing, Request aggregation, Information awareness
- Combined use of multicast & caching inside access networks
- Overlay approach facilitates deployment
- Potential benefits:
  - Localizing traffic at sub-trees
    - Exchanging network traffic with storage
    - Reducing load on content provider
  - Potentials for download reduction
    - No search for data provider
    - No uplink bottlenecks

Thank you!
Questions?

Konstantinos Katsaros

Mobile Multimedia Laboratory

Department of Informatics/Computer Science
Athens University of Economics and Business

ntinos@aueb.gr
http://mm.aueb.gr/