

A hybrid overlay multicast and caching scheme for information-centric networking

Konstantinos Katsaros, George Xylomenos and George C. Polyzos

Mobile Multimedia Laboratory

Athens University of Economics and Business

Outline

- Motivation
- Design objectives
- MultiCache architecture
 - Functionality overview
- Protocols
 - Multicast
 - Caching
- Performance evaluation
- Conclusions & Future work


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. ICT FP7 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

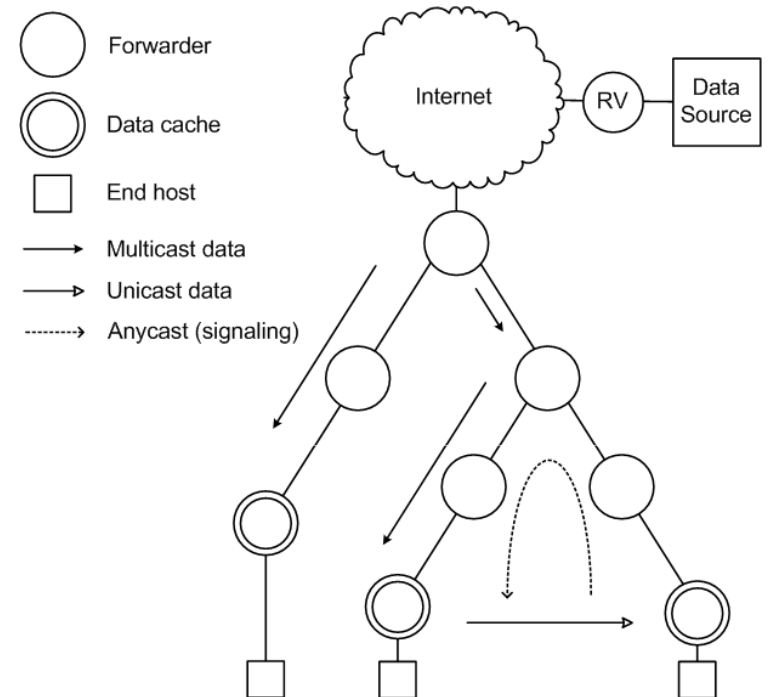
- **Overlay Access Routers (OARs)**
 - Deployed inside access networks
 - Gradual deployment is feasible
- Providing **overlay multicast**
 - Based on Scribe over Pastry
 - Scalable
 - Adaptive to physical topology
- Acting as **caches**
 - Multiple cache locations
 - Close to end-hosts
- Proxy-ing end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



MultiCache

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
 - Correlating forwarding state with data availability
 - Localizing traffic inside sub-trees
- Unicasting cached data
 - Reducing stretch...
- Content fragmentation
 - Piece level
 - Parallelizing transfers
 - Enabling partial caching
 - Block level
 - Facilitating cache provision...



MultiCache

Multicast


- Publisher advertizes data to RV point(s)
- OARs subscribe to data
 - Possibly aggregating end host requests
 - Scribe JOIN messages extended with joining OAR information:
 - IP address, listening port, 32-bit AS number
 - Later used for cache provision
- Subject to arrival time, requests served with:
 - Multicast from the source
 - Unicast from a near-by/local cache
- JOIN suppressing OARs: *meta-cache* OARs
 - Store joining OAR information
 - Keep track of forwarded data



MultiCache

Caching

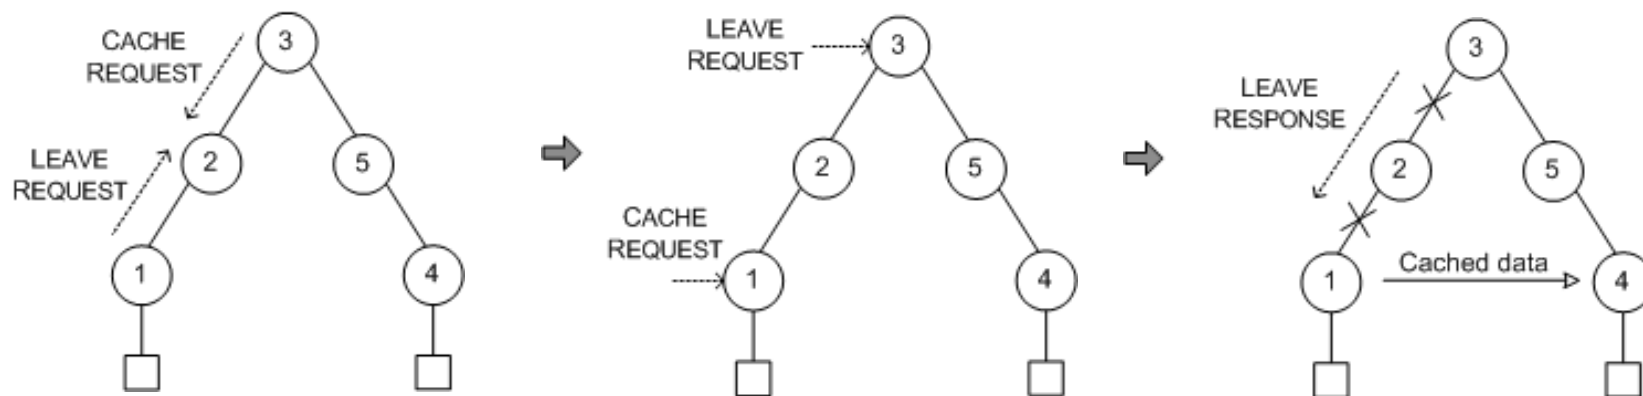
- Cache availability signaled to ancestors (CACHE UPDATE)
 - Suppressed at the first OAR already aware of another cache (or the RV)
- Cache request issued by meta-cache OARs (CACHE REQUEST)
 - Anycast towards rest of the children
 - Preferably selecting a child in the same AS with the new subscriber
- Cached data sent to requesting OAR via unicast
 - Also contributing to cache creation



MultiCache

Cache replacement: signaling

- Multicast tree participation correlated with data availability
 - Evicting forwarding state before data
 - Avoiding cache misses: either cached data or multicast
- Scribe's *Leave* procedure altered
 - Tree branches torn down on response from highest non leaving ancestor
 - Cache requests for items pending for eviction served normally
- Incoming requests buffered until cache space available



MultiCache

Cache replacement: policies

- Common policies
 - Least Recently Used (LRU)
 - Evicting the LRU fragment of the LRU file
- MultiCache specific policies
 - Most Recently Used - Intra Domain (MRU-Intra)
 - Evicting the fragment most recently delivered to an OAR k of the same domain
 - Increased probability of the fragment not evicted by OAR k
 - Most Frequently Used - Intra Domain (MFU-Intra)
 - Evicting the fragment most frequently delivered to OARs of the same domain
 - Increased probability for an alternative caching location to exist
- MRU/MFU – Intra enforced on fragments
 - Fragments not associated with files
 - No control signaling and state overhead

Performance evaluation

Workload and metrics

- Simulation based evaluation
- GT-ITM topologies
- BitTorrent-like workload
 - Mandelbrot-Zipf distribution of file popularity [1]
 - Exponential decay arrival process for file requests [2]
 - Fixed *file* arrival rate [2]
 - Trace sampled file sizes
- Metrics
 - Cache hit ratio (CHR) (%)
 - Intra-domain cache hit ratio (CHR - Intra) (%)
 - Distance to block source

- [1] M. Hefeeda and O. Saleh, "Traffic modeling and proportional partial caching for peer-to-peer systems," *IEEE/ACM Transactions on Networking*, vol. 16, no. 6, pp. 1447–1460, 2008.
- [2] L. Guo, S. Chen, Z. Xiao, E. Tan, X. Ding, and X. Zhang, "A performance study of BitTorrent-like peer-to-peer systems," *IEEE JSAC*, vol. 25, no. 1, pp. 155–169, 2007.

Performance evaluation




Important parameters

- Relative Cache Size, S_r
 - Percentage of “infinite cache size ”
I.e., minimum cache size to avoid replacements [1]
- Deployment density, $d \in (0, 1)$
 - Fraction of access routers enhanced with overlay functionality
- Localizability of traffic
 - Highly depends on item popularity inside a domain
 - High popularity favors cache availability
 - Localizability factor, $l \in (0, 1)$
 - $l = 0$, all nodes uniformly dispersed throughout the AS's
 - $l = 1$, all nodes inside a single AS

[1] L. Fan, P. Cao, J. Almeida, and A. Z. Broder, “Summary cache: a scalable wide-area web cache sharing protocol,” *IEEE/ACM Transactions on Networking*, vol. 8, no. 3, pp. 281-293, 2000.

Results: cache replacement policies




$l = 0$, $d = 0.25$, 8 MB Piece, LRU vs. MFU-Intra vs. MRU-Intra




- No significant difference between policies
 - Same performance with simpler, file-oblivious MFU/MRU-Intra policies
- High cache hit ratios
 - Taking advantage of cache multiplicity
 - Overlay multicast minimized
- Localizing traffic
 - High CHR-Intra values for $S_r \geq 2.5\%$
 - Reduced travelled distances

Results: incremental deployment

$I = 0.5$, 8 MB Piece, MFU-Intra




- Higher CHR for denser deployments
 - Taking advantage of the increased caching space in the entire network
- CHR-Intra decreases with deployment density
 - Increased overlay size, increased overlay (stretched) routing
 - Less *direct cache hits*
 - Increasing interdomain cache provision
- Slight reduction of travelled distances
 - More caching locations, decreased distances
 - Less direct hits
- Sparse deployments yield good performance!
 - Low investment cost



Results: Localizability

$d = 0.25$, 8 MB Piece, MFU-Intra

- Higher localizability imposes greater stress on caches
 - Reducing CHR
 - Multicast gradually takes over
- Taking advantage of localized request patterns
 - Increasing CHR-Intra
 - Up to the cache size limit



Conclusions & Future Work

- Resource sharing enabled by information awareness
- Overlay character facilitating deployment
- Exchanging traffic with storage
 - Building on cache multiplicity
- Avoiding (stretched) overlay multicast as possible
 - Building though on its forwarding scalability
- Sparse deployments enough to rip the benefits
 - Low investment cost
- Next Steps:
 - End user experience
 - Direct comparison with BitTorrent
 - Gaining control of inter-domain cache provision
 - *Canonical* Pastry

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/>