

Resource Management and Congestion Control for In-Network Caching Environments

Ioannis Psaras, Lorenzo Saino

University College London
COMET/PURSUIT Workshop
30th October 2012
(i.psaras, l.saino)@ucl.ac.uk

Caching back on the agenda

Why has caching become a hot topic again (after 20 years)?


Or

Why is it interesting to invest on ICN caching?

- **Nature of traffic**
 - Content creation is not dominated by big content providers anymore: **user-generated content, interactive social networking and web applications** have changed radically the nature of the content to be stored, processed and delivered to the end-user.
 - Lots of content stays local
- **Centralised Content Storage and Distribution**
 - The CDN deployment spread has reached the state of deploying **one surrogate server farm per domain**. Novel and radical approaches to content storage and caching are needed, potentially based on **smaller scale and finer-grained CDN deployments**.
- **Inefficient Use of Resources**
 - **Location-centric** models limit the ability to fully utilise resources that are available along the route from the provider to the consumer(s), such as storage (e.g., for caching).

Centralised architecture has remained, but amount of content has grown massively

Overlay vs. In-Network Caching



- Important to understand:

“What are the differences between overlay and in-network caching?”

- Caching at the chunk-level NOT at the file-level (probably NOT at the packet level either)
 - As contents pass through router-caches they replace existing “old” contents
 - Caching can happen transparently into the network at random or predefined (rendezvous) points
- Therefore, *replacement happens at line-speed* – what does this imply?
 - Overlay caching depends on centralised (control-plane) co-ordination and management of caches (or de-centralised among very few nodes) – In-network caching *does not/cannot*
- Hence: **no book-keeping possible** (even LRU might be too costly!)
 - Impossible to signal/collaborate with other caches, or the control plane – the exact location of contents cannot be known
 - Caching operations happen transparently inside the network
 - Decentralised distribution and replacement of contents in caches

In-Network Caching

We need to identify the problems

- *Is the **cache location** problem still around?*
 - Are some routers more important than others, in order to equip them with caches?*
- *What is the **cacheable unit** (i.e., the size of a chunk)?*
 - Is it a universal setting, e.g., MTU?
- *What is the **size of a cache**?*
 - This was not an issue in overlay, proxy approaches to caching
- *How do we choose which contents to cache?*
 - Most popular ones? Reduce *traffic* redundancy
 - But also reduce *caching* redundancy! (was not an issue in coordinated environments, e.g., overlay caching)
- *What do we want to achieve by caching inside the network?*
 - **Native and transparent** content storage, caching and distribution functionality that fully utilises available resources

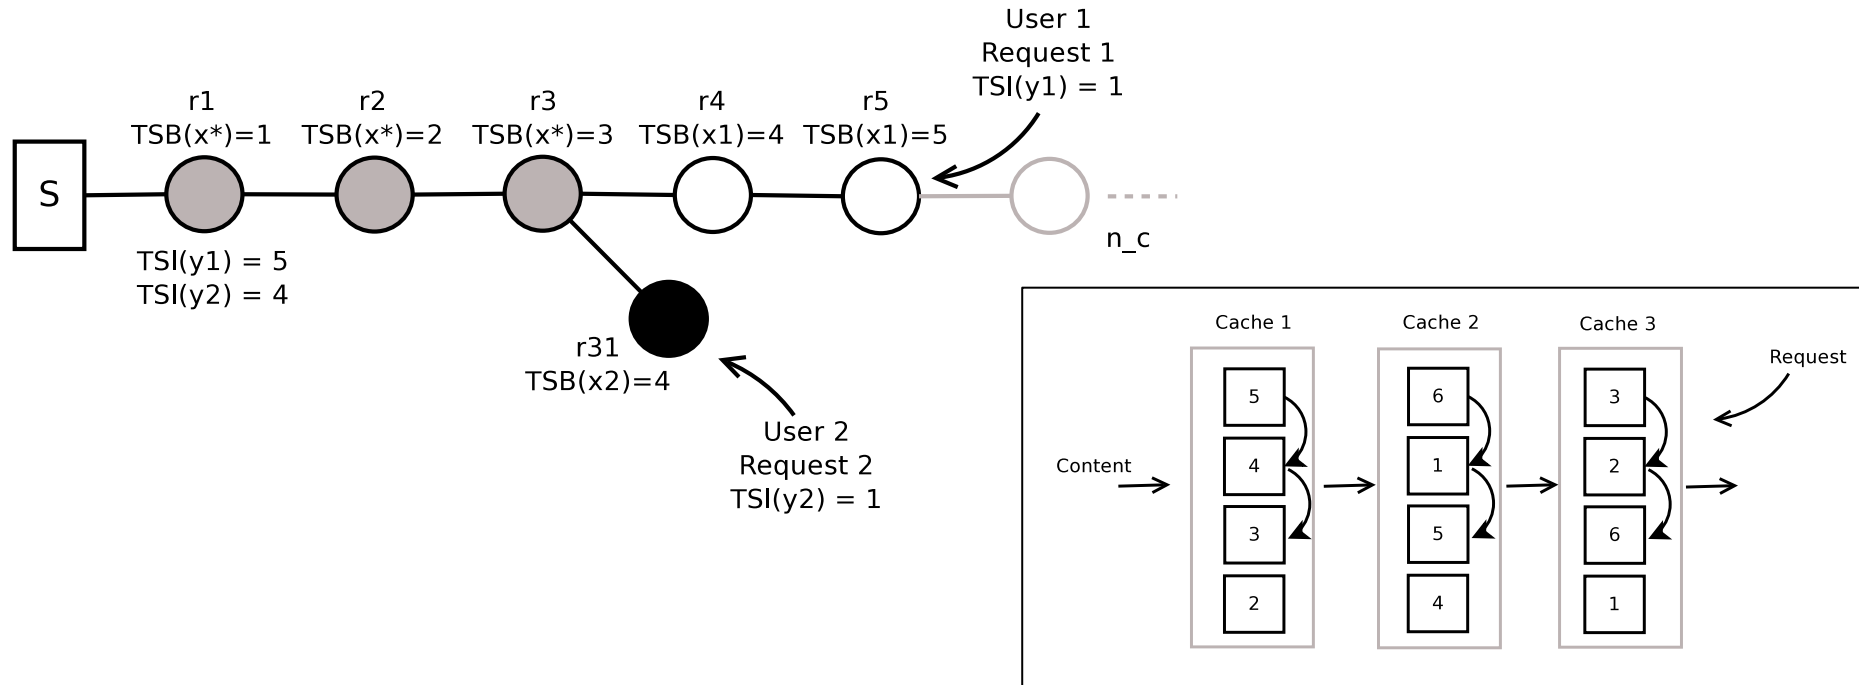
*Wei Koong Chai, Diliang He, Ioannis Psaras and George Pavlou

“Cache Less for More in Information-Centric Networks”, Best Paper Award, IFIP Networking 2012

Resource Management Considerations for In-Network Caches

- **Motivation:** decentralised, real-time resource management is crucial for in-network caching systems
 - High content replacement frequency may result in content being replaced before getting a hit.
- **Goal:** to study alternative in-network caching strategies that support *fair content multiplexing*.
 - We are concerned with *reducing caching redundancy* along a path of caches.
 - To reduce redundancy, we play with:
 - The *caching capability of a path of caches*.
 - The *population*, in terms of distinct content objects, that this path serves *per unit time*.

ProbCache: Probabilistic In-Network Caching for Fair Content Multiplexing



ProbCache: Probabilistic In-Network Caching

$$ProbCache(x) = \underbrace{\frac{\sum_{i=1}^{c-(x-1)} N_i}{T_{tw} N_x}}_{TimesIn} \times \underbrace{\frac{x}{c}}_{CacheWeight}$$

Caching **Capability** of a Path **Weight**-based Caching

Summary and Conclusions

- In-Network Caching:
 - Several differences between overlay and in-network caching
 - We have to identify the new problems and *ask the correct questions*
 - *Fair content multiplexing* is one of these open issues
- *ProbCache**: Probabilistic In-Network Caching:
 - we **argue against the necessity of a ubiquitous or deterministic caching** strategy in ICN environments
 - we investigate the possibility of caching selectively in order to achieve *fair content multiplexing* and in turn, higher performance gain.
 - *ProbCache* is an online algorithm, which is based on the concept of **path cache capability**.
 - *ProbCache* improves performance by approximately **8-15%**.

*“Probabilistic In-Network Caching for Information-Centric Networks”, ICN Sigcomm Workshop 2012