CoCaching in Named Data Networking

Azzam Alsudais (student) Eric Keller
University of Colorado at Boulder
{Azzam.Alsudais,Eric.Keller}@colorado.edu

Named Data Networking (NDN) is a proposed future Internet architecture that is designed around a content-centric model, where content is requested by name and cached within the network. A recent proposal [1] demonstrated that an architecture centered around caching at edge routers is more deployable, but also has nearly the same amount of performance gains. As shown in Figure 1, we note, however, that there is still a large gap between the edge-based caching and the ideal (where caches can perfectly predict what content will be requested).

In this work, we are looking to close the gap. To do so, we propose CoCaching, a cooperative caching for NDN. That is, where router nodes can leverage the cache content of nearby nodes, but not necessarily nodes on the path to the content origin. The key research challenge is how should we share caches? Below we discuss two high-level approaches that we are exploring and the tradeoffs that factor in.

A **Reactive** approach can be applied to share cache content with nearby routers. That is, when an edge router receives an interest in some content, it reactively expresses its interest in that content to other nearby routers, and similarly, routers nearby should respond with the content has it been cached before. However, there are some challenges that should be taken into account when implementing a reactive approach. For example, a router should choose whether to request content from nearby routers or just forward the request towards the content origin. Also, a router should determine what routers to consider as nearby routers - perhaps by comparing latency to the content origin with latency to other routers.

On the other hand, a **Proactive** approach can be applied towards cooperative caching, where nearby routers proactively share cache content with one another prior to receiving interests in that content. When implementing the proactive approach with a high accuracy rate (where routers can accurately predict what cache content other routers might be interested in), requests latency can be minimized which closes the gap even further between the edge-based caching and the ideal case. But, of course, there are some tradeoffs that should be considered. There is a tradeoff between storage space and quantity. In other words, cache sizes are limited; thus, the choice of how much and what content to share is critical. Another key principle is object placement in the cache. Upon receiving unsolicited content packets, do routers treat them differently? Or do they just add them to the cache using the default replacement policy, e.g. *LRU* or *LFU*? In addition, both approaches, reactive and proactive, impose a tradeoff between network traffic and performance. That is, impulsively sharing cache content with other nearby routers can result in high network traffic which could negatively affect performance.

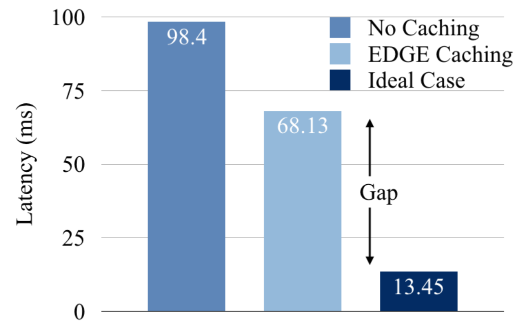

Figure 1: latency measurements in milliseconds for three different settings: no caching, edge caching, and the ideal case where it is always a cache hit.

Preliminary Work As an experiment setup, we used a *k*-ary PoP network topology with *k*=2 and a tree depth of 4, which is similar to the setup used in [1]. We used this setup to demonstrate the gap in Figure 1, and demonstrate that we might be able to close the gap by showing that data could be served by nearby caches more than 30% of the time (by assuming that a nearby cache is an edge cache that is less than two hops away). In short, although there are many questions to answer, the preliminary steps we have taken show great promise.

References

[1] S. K. Fayazbakhsh, Y. Lin, A. Tootoonchian, A. Ghodsi, T. Koponen, B. M. Maggs, K. C. Ng, V. Sekar, and S. Shenker. Less pain, most of the gain: Incrementally deployable ICN. In *Proceedings of the ACM SIGCOMM 2013 conference on SIGCOMM*, Hong Kong, China, August 2013.