

Probabilistic In-Network Caching for Information-Centric Networks

Ioannis Psaras, Wei Koong Chai and George Pavlou

Dept. of Electrical & Electronic Engineering

University College London

WC1E 7JE, Torrington Place, London, UK

Email: {i.pсарas, w.chai, g.pavlou}@ee.ucl.ac.uk

Abstract—In-network caching necessitates the transformation of centralised operations of traditional, overlay caching techniques to a decentralised and uncoordinated environment. Given that caching capacity in routers is relatively small in comparison to the amount of forwarded content, a key aspect is the distribution of content among the available caches. In this paper, we are concerned with decentralised, real-time distribution of content in router caches. Our goal is to reduce caching redundancy and in turn, make more efficient utilisation of available cache resources along a content delivery path.

Our in-network caching scheme, which we call *ProbCache*, approximates the *caching capability* of a path and caches contents probabilistically in order to: *i*) leave caching space for other flows sharing (part of) the same path, and *ii*) fairly multiplex contents of different flows in caches along a shared path.

We compare our algorithm against universal caching and against schemes proposed in the past for Web-Caching architectures, such as Leave Copy Down (LCD). Our results show reduction of up to 20% in server hits, and up to 10% in the number of hops required to hit cached contents, but, most importantly, reduction of cache-evictions by an order of magnitude in comparison to universal caching.

I. INTRODUCTION

Naming content objects directly instead of their respective end-host machines gives the opportunity to identify content objects as they travel from source to destination [1], [2], [3]. In turn, given that the network transfers *named* objects (instead of unidentifiable *data containers*, *i.e.*, IP packets), these objects can be cached *in the network* and be forwarded to subsequent users interested in the same content [4], [5].

In-network caching has therefore emerged as a distinct research field in the context of *Information- or Content-Centric Networks* (ICN/CCN). In-network caching exhibits fundamental differences from overlay web-caching [6], [7], or hierarchical and co-operative caching approaches [8], [9] and poses new challenges [10], [11]. For instance, past research considered mainly caching of whole files (with a few exceptions [12]) as well as administration of their placement [13] and location [14] by centralised entities, *e.g.*, DNS and HTTP redirection. Centralised administration of content placement gives the opportunity to control and manage network resources better at the cost of: *i*) increased communication overhead to update the content location database, and *ii*) reduced flexibility in terms of available cache locations.

In contrast, Information-Centric Networking enables caching of addressable content *chunks* [2], [12] in *every* cache-equipped network device [15] and replacement of cached chunks at line-speed [16]. Although, this operation

increases the availability of cache locations [17] it renders prohibitive the process of updating logically-centralised content location databases with the exact location of cached contents. This decentralised, location-independent operation alters many of the basic features of past overlay caching techniques, *e.g.*, content-to-cache allocation [13], while it invalidates the applicability of some others, *e.g.*, content placement based on fixed overlay topologies of caches and servers [14].

In this paper, we are concerned with cache management operations that have to be adjusted to fit in a completely decentralised and uncoordinated environment. We focus on the fair sharing of the available cache capacity of *a path* among the content flows that use part of this path per unit time. In other words, we focus on the allocation of the available cache capacity along *a path* of caching entities among different content flows. As a starting point we intuitively observe that caching *every* content in *every* cache-enabled device along the delivery route (an operation implicitly supported in [2]), inherently causes huge *caching redundancy*. Subsequently, our goal is to reduce caching redundancy and make more efficient use of available cache resources, in order to reduce overall network utilisation and potentially increase user-perceived quality. Indeed, our initial investigation of selective caching policies based on node centrality metrics shows very promising results [18] on this direction.

To achieve our goal, we approximate the *caching capability* of a given path *per unit time* (Section II-B) and we design *ProbCache*, a *probabilistic algorithm* for distributed content caching along a path of caches (Section III). Our results suggest that there is indeed a lot of space for resource management optimisation of in-network caching policies, given that appropriate content multiplexing rules are in place (Section IV).

We use the terms “router” and “cache” interchangeably to refer to cache-enabled network devices [16]; it should be noted that our approach does not require every router to be cache-enabled, but it will work in hybrid architectures as well. Furthermore, we refer to content “packets”, “messages” and “chunks” interchangeably to refer to the *cacheable unit*, which is not necessarily of similar size to an IP packet. In fact, we leave open the actual size of the cacheable unit which is yet to be defined by the ICN research community. We highlight that the concepts and algorithms proposed in this paper are *cache unit-* as well as *architecture-agnostic* and would apply to almost any ICN environment [19], [2], [20].

We differentiate between two types of *redundancy*, namely,

network traffic redundancy [5] and *caching redundancy* [13], [9]. Caching has been traditionally used to reduce traffic redundancy [4], [6]. Assuming a set of available caches, as opposed to a single proxy cache, redundancy can still exist between cached contents in different locations, something that has also been investigated in the past for overlay caching schemes [13], [9]. However, and as mentioned earlier, overlay schemes require some form of co-operation between the caches themselves and/or between the caches and a central management entity. In contrast, in the case of in-network caching, management has to happen in an uncoordinated, uncooperative fashion. In this paper, we consider each path of caching entities as a pool of caching resources; we try to find optimal ways of distributing content in these caches in order to eliminate *caching redundancy* and in turn, reduce *traffic redundancy*. Our results show that careful content flow multiplexing in caches can achieve up to 20% more cache hits. Surprisingly, this translates to one order of magnitude reduction in terms of traffic redundancy.

II. SYSTEM MODEL AND ASSUMPTIONS

We argue that *in-network cache management has to take into account the approximate cache capacity of the path of caches and the estimated amount of traffic that these caches serve per unit time, in order to make decisions on whether to cache incoming contents or not*. In Section II-A, we make assumptions that help us approximate the cache capacity of a given path and in Section II-B, we present our design principles.

A. Assumptions on Caching Technologies

By definition, *caching* is different to *storage*, both in networks and in computer systems, in that caching keeps contents stored *for a specific amount of time and not indefinitely*, as in storage. Therefore, the size of a cache is a relative factor, which cannot stand on its own, but instead has to be linked to the amount of time that a given content is cached for. We, therefore, associate the cache size with the traffic that the corresponding router serves per second. Our cache size unit is the *number of seconds worth of traffic cached* in a given router and depends on the speed of the outgoing links of the router in question.

One important question then is: “*For how long can we afford to cache contents in a given router?*”. Furthermore, given that in this study we are concerned with *paths* of caches and not with single-caches only, another important question is: “*For how long do we need to cache contents in a given path in order to minimise redundant traffic and maximise gain?*”. Our reasoning for answering these questions is as follows:

- Today’s memory access technologies guarantee *line-speed* access to DRAM chips of up to 10GBytes at a reasonable price [16]. This means that a 5GByte-long

LINK NAME	LINK SPEED	1-SEC TRAFFIC	SECS OF TRAFFIC IN A 10GB CACHE
OC-24	1,2 Gbps	~ 0.15 GBs	~ 64 secs
OC-48	2,4 Gbps	~ 0.31 GBs	~ 32 secs
OC-192	9,9 Gbps	~ 1.25 GBs	~ 4 secs
OC-768	39,8 Gbps	~ 5 GBs	~ 2 secs
OC-1536	79,6 Gbps	~ 10 GBs	~ 1 sec
OC-3072	159,2 Gbps	~ 20 GBs	~ 0.5 secs

TABLE I
LINK SPEEDS AND RELATED CACHING PROPERTIES

cache behind a 40Gbps link¹ can safely be assumed to hold contents for one second (see also Table I). *Without loss of generality, we assume that each cache along a path has sufficient memory to cache contents in a DRAM chip for at least one second* (see third column in Table I).

- Authors in [5] show up to 60% bandwidth savings by redundant traffic elimination *within the first 10 seconds after the original transmission*, in some enterprise networks. We associate redundant traffic, *i.e.*, subsequent requests for the same content, with the aforementioned result. That is, we consider, without loss of generality, that any content should be kept in any one of the path’s caches for a *target time window*, T_{tw} , of 10 seconds.

Both the above settings are relatively arbitrary and can change in the future, but these values are a good starting point based on today’s technology. In addition, our concepts and algorithms presented next are still applicable should these values change.

B. System Model

We assume the topology of Fig. 1. The path from source to destination comprises n routers, where router r_i has N_i cache slots, each able to hold one addressable content chunk; based on the discussion above, we assume that N_i slots can hold one second worth of traffic. Our model notation is given in Table II.

Path Cache Capacity. The *caching capacity* of our path of caches is $\sum_{i=1}^n N_i$, in terms of memory, which amounts to n seconds worth of traffic cached along the path.

Path Cache Capability. Given that our target time window is T_{tw} seconds worth of traffic cached along a given path, the *caching capability* of an n -long path, as a fraction of the required capacity for T_{tw} seconds, is $\frac{\sum_{i=1}^n N_i}{T_{tw}N}$, where N is the average cache size along the path. We revisit the issue of average cache size in the next section.

Symmetric Paths. Request and Content messages follow the same route, according to recent proposals [1], [19], [2], [21].

Path Length Monitoring. Similar to the TTL field included in IP packets, our design requires that ICN *request message*

¹To the best of our knowledge, operators to date use links of up to 40 Gbps, while some operators plan to update a limited number of their links to 100 Gbps: <http://www.prnewswire.com/news-releases/verizon-first-service-provider-to-announce-100g-deployment-on-us-network-118891754.html>

SYMBOL	MEANING
n	Number of caches on the path
N_i	Cache memory in r_i that holds 1-sec worth of traffic
T_{tw}	Target Time Window (set to 10 secs here)
TSI, c	Time Since Inception (Header field - Request Message): Hop-Distance from Client, Value range: 1 to n
TSB, x	Time Since Birth (Header field - Content Message): Hop-Distance from Server, Value range: 1 to n

TABLE II
MODEL NOTATION

headers include the *Time Since Inception* (TSI) field and *content message* headers include both the TSI and the *Time Since Birth* (TSB) fields. Every router increases the TSI value of request packets by one. The content source attaches the TSI value that it sees on the request message to the content message. Every router increases the TSB value of the content message by one. Hence, during the content message's journey back to the client, the TSI value in the content message is a fixed value and denotes the path-length of this specific content flow, while the TSB value denotes the number of hops that the content message has travelled so far².




Fig. 1. Design Topology

III. BUILDING *ProbCache*

We approach the problem of content placement within a system of caches from the *path caching capability* point of view. In particular, *each router, based on the amount of traffic that it has to serve per unit time, indirectly approximates the number of copies of incoming contents that the (rest of the) path can accommodate*. This value is the *TimesIn* factor (see Section III-A). Based on the *TimesIn* indication and on the *router's distance from the user*, which we call *CacheWeight* (see Section III-B), each router probabilistically caches contents as they travel along the path (see Section III-C).

A. Estimating the Caching Capability of a Path

Consider two users shown, in Fig. 1, five and four hops away from the server, respectively. The total cache capacity of the path is $\sum_{i=1}^n N_i$, where $n_1 = 5$ for *Request₁* and $n_2 = 4$ for *Request₂*. Grey circles denote the caches that have to be shared between the two users, while white and black circles denote caches used exclusively by Users 1 and 2, respectively.

The number of times that the path can afford to cache this packet is reflected in the *TimesIn* factor, whose calculation takes place as follows:

²In case of a cache hit, the TSI and TSB values are treated as if the cache is the actual source, that is, the TSI value of the content packet is replaced by that of the *Request* packet, while the TSB is set to 1.

$$TimesIn(x) = \frac{\sum_{i=1}^{c-(x-1)} N_i}{T_{tw} N_x} \quad (1)$$

where c is the *Time Since Inception* (TSI) value and x is the *Time Since Birth* (TSB) value that the router sees in the header of the content message (Table II). As an example, consider content messages traveling through router r_2 to fulfil *Request₁*, in Fig. 1, will have $TSI = 5$ and $TSB = 2$, while contents for *Request₂* will have $TSI = 4$ and $TSB = 2$. The sum in Eq. 1 considers the result of the subtraction of TSI minus TSB (or $c - (x - 1)$ in Eq. 1) to account for the *remaining* caches only, instead of the total number of caches from the content source to the client.

B. Weight-based Caching

We argue that in order to achieve fair resource (in our case cache) allocation in a distributed environment, each content flow has to take into account other content flows sharing the same path (grey circles in Fig. 1). Hence, to decide *where* to cache the number of copies that *TimesIn* indicated, we use the *Cache Weight* factor:

$$CacheWeight(x) = \frac{x}{c} \quad (2)$$

where x is the TSB value of the packet header and c is the TSI value. We note that the TSI value is fixed during the content packet's journey from the source to the client, while the TSB value is increasing for each router the packet traverses; hence, $CacheWeight \rightarrow 1$ as the packet is getting closer to its destination. This is a desirable system property considering path-diversity, in terms of number of hops, between clients and sources of different content flows.

C. *ProbCache*: Probabilistic In-Network Caching

ProbCache is the product of *TimesIn* and *CacheWeight*. Each router along the path caches incoming chunks with probability *ProbCache*, depending on their TSI and TSB values.

$$ProbCache(x) = \underbrace{\frac{\sum_{i=1}^{c-(x-1)} N_i}{T_{tw} N_x}}_{TimesIn} \times \underbrace{\frac{x}{c}}_{CacheWeight} \quad (3)$$

The *Cache Weight* factor increases the probability of a content being cached closer to its destination. This way, we achieve *fair content flow multiplexing* between contents that travel to different destinations in terms of path length. For example, contents for User 1 in Fig. 1 should be cached inversely proportionally to User 1's distance from the server, *i.e.*, in (white) routers r_4 , or r_5 , in order to leave (grey) routers $r_1 - r_3$ for clients travelling shorter paths to cache their contents. This is in accordance with our previous findings in [17] that contents tend to be cached for longer towards the edge of the network.

NOTE: To calculate the *TimesIn* factor each router has to conjecture on the size of the rest of the routers on the path.

However, given that we do not know what amount of memory each router will have, or if backbone routers, for instance, will have bigger caches than edge-network routers, we make the following simplifying assumption. *Each router assumes that all other routers on the path have the same amount of cache as it has got.* Even in a random-size cache deployment scenario, this assumption serves our purposes well. That is, a router with a big cache, compared to the caches along the path, will be caching contents with higher probability, while a router with a small cache will experience the opposite effect (Eq. 1). This is a desirable system property which alleviates the effect of unknown cache sizes, but at the same time guarantees fair load distribution among nodes with diverse caches. We relax the assumption of homogeneous cache sizes later on and show that although our simplifying assumption does not harm the performance of *ProbCache* in heterogeneous cache size environments, it fails to exploit extra caching resources (Section IV-B).

IV. PERFORMANCE EVALUATION

We test our algorithm in a custom-built simulator, where we use *Least Recently Used* (LRU) caches. Given that the ultimate goal of *ProbCache* is to manage caching resources more efficiently, by reducing cache redundancy, the straightforward metric of interest is the *reduction of Server Hits*. Furthermore, the gain from serving user requests from intermediate caches, instead of travelling to the origin server, depends on the number of hops that the request travels before it eventually hits cached contents. Clearly, as the number of hops increases, the overall gain decreases. To measure this gain, we also monitor and present the *Hop Reduction Ratio*.

We use binary tree topologies; set the exponent of the Zipf distribution of content popularity to 0.8 to capture the worst case of fairly unpopular content [22]; and compare the performance of *ProbCache* against: *i*) the universal caching approach proposed in [2], a scheme that we call *Cache Everything Everywhere* (CE^2), *ii*) the Leave Copy Down (LCD) [23] algorithm proposed in the past for overlay caching topologies³, *iii*) a probabilistic algorithm that caches with probability $p = 0.7$ at every cache and *iv*) a probabilistic algorithm that caches with probability $p = 0.3$.


A. Scenario 1: Path Cache Capacity Management

We use a 6-level binary tree topology of 127 nodes in total. The root node represents the server node and we configure requests to come from the last two levels of the tree, assuming that individual users are not connected to backbone network routers; we note however that results are similar in case of users attached to all routers. We generate a total of 100,000 requests, to allow enough time for the system to reach a steady-state. As noted earlier, our cache-size unit is the *number of*


seconds worth of traffic cached in a given router. In our first experiment, we assume homogeneous caches and we test the performance of the algorithms as we increase each router's cache capacity from one to six seconds (Fig. 2).

The Server Hits performance difference balances around a reduction of approximately 12-15% (12,000-15,000 Server Hits less) for *ProbCache* against CE^2 and around 7-10% compared to the rest of the algorithms evaluated here (Figs. 2(a)).

The *Hop Reduction* difference is roughly 8-10% against the CE^2 scheme, while it is smaller compared to the rest of the algorithms and especially compared to LCD (up to 3%). However, considering short path lengths (up to 6 hops) and the operational properties of both CE^2 and LCD, whose design targets bringing content as close to the end-user as possible, we argue that even this small performance difference unveils better exploitation of storage resources in merit of *ProbCache*. To validate this claim further, we trace the average number of: *i*) Cache Hits, *ii*) Received Requests, and *iii*) Cache Evictions per tree level. The size of the cache is set to be equal to three seconds worth of traffic at each cache.



(a) Server Hit Saving



(b) Hop Reduction Ratio

Fig. 2. Scenario 1: Homogeneous Caches

We observe, although do not present the related results here due to space limitations, disproportionate differences between the number of Cache Hits and the corresponding number of Received Requests in case of *ProbCache*. That is, *ProbCache* has more Cache Hits and less Received Requests than the rest of the algorithms. This difference in Cache Hits and Received Requests owes to the efficient resource management of *ProbCache*, which results in contents staying in caches for longer. To substantiate our claims, we plot in Fig. 3 the average number of evictions per tree-level, where the smaller the tree level, the closer to the server this level is. The very big difference in terms of cache evictions validates our claim that per-path, distributed resource management results in more contents staying in the caches for longer and therefore, getting more cache hits. This huge reduction in cache evictions in Fig. 3 reflects the reduction in terms of traffic that flows through the network and subsequently, redundant traffic elimination.

³According to LCD [23], every request for a specific content causes the content to be copied one hop closer to the user, or one level down in the cache hierarchy. If a request for a content is received after the content has been evicted from the cache, according to the LRU policy, the content has to be retrieved from another cache up-the-hierarchy, or from the origin server.




Fig. 3. Scenario 1: Average Number of Cache Evictions per Tree Level

B. Scenario 2: Heterogeneous Cache Sizes

We proceed to relax the assumption of homogeneous caches. In particular, we expect that as research in the area matures, caches will be sized according to a specific norm or a set of guidelines. We conjecture three potential norms: *i*) larger caches are deployed towards the core of the network, where servers reside, *ii*) larger caches are deployed towards the edges of the network, where users are connected, *iii*) all caches have roughly similar sizes (homogeneous cache scenario presented in the previous section).

We use the same topology as before, set cache sizes according to the formula: $N_i \times TSB$ (see Table II) and observe the performance difference in case of heterogeneous caches. That is, if a core router caches one second worth of traffic and we assume larger caches towards the edges of the network (which we denote as $c(ore) \rightarrow E(dge)$), then the edge router of a six-hop path will cache six seconds worth of traffic. The capacity of the path in that case is $c \rightarrow E : \sum_{i=1}^{c-(x-1)} (c-i)N_i$. The opposite applies for deployments where larger caches are deployed in the core of the network (which we denote as $C(ore) \rightarrow e(dge)$); the capacity of the path in this case is $C \rightarrow e : \sum_{i=1}^{c-(x-1)} iN_i$.

We modify the *TimesIn* factor based on the above path capacities for each case in Eqs. 4 and 5, respectively and call the new algorithm *ProbCache+*. The corresponding cache capacity calculations along a six-hop path for Eqs. 1, 4 and 5 are shown in Fig. 4. We use the default algorithm (*i.e.*, Eqs. 1 and 3) as our benchmark for comparisons, and for clarity, we compare the performance of *ProbCache+* against CE^2 [2] and LCD [23] only.

$$TimesIn_{(C \rightarrow e)}(x) = \frac{\sum_{i=1}^{c-(x-1)} iN_i}{T_{tw}N_x} \quad (4)$$

$$TimesIn_{(c \rightarrow E)}(x) = \frac{\sum_{i=1}^{c-(x-1)} (c-i)N_i}{T_{tw}N_x} \quad (5)$$

Figure 5 summarises our results. We present the ratio of each protocol's cache hits performance in case of extra cache added in the core (Scenario 2.1) or the edge (Scenario 2.2) over the performance in case of homogeneous caches (Scenario 1).

As noted, *ProbCache* is the original algorithm in Eq. 3, hence this algorithm does not take into account the extra caching capacity; in contrast, *ProbCache+* incorporates Eq. 4 in Scenario 2.1, Fig. 5(a), and Eq. 5 in Scenario 2.2, Fig. 5(b). In both cases, we observe that both CE^2 and LCD fail to utilise the extra available cache capacity; that is, their performance increases by less than 5%, although the extra

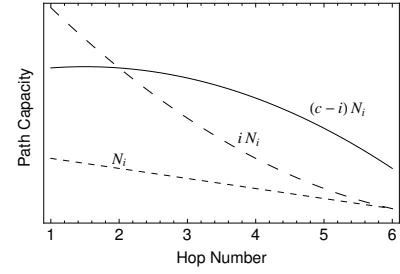




Fig. 4. Path Capacity calculation that takes place at each router along a 6-hop Path: the N_i curve is for homogeneous caches, hence the capacity is given by $c \rightarrow e : \sum_{i=1}^{c-(x-1)} N_i$; the iN_i curve is for larger core caches, hence $C \rightarrow e : \sum_{i=1}^{c-(x-1)} iN_i$; the $(c-i)N_i$ curve is for larger caches towards the edge, hence $c \rightarrow E : \sum_{i=1}^{c-(x-1)} (c-i)N_i$.

cache capacity we added is 30% compared to Scenario 1 (Section IV-A). We see that this is the case too with *ProbCache* and *ProbCache+* in case of extra cache capacity at the core of the network (*i.e.*, Fig. 5(a))⁴.

Interestingly, and in contrast to recent results reported in [24], where the authors show that cache size heterogeneity does not improve the caching performance in ICNs, we find that both *ProbCache* and especially *ProbCache+* utilise a good amount (up to 18%) of the extra caching capacity added towards the edge of the network in Scenario 2.2, Fig. 5(b). The overall performance difference in that case is between 20-25% compared to CE^2 and LCD, if we take into account the absolute results reported in Section IV-A, Fig. 2(a). We argue that this is due to more efficient resource utilisation of the proposed algorithm and better content multiplexing in caches, something that does not constitute a design goal for neither CE^2 nor LCD.



(a) Scenario 2.1: Larger Cache at the Core ($C \rightarrow e$)



(b) Scenario 2.2: Larger Cache at the Edge ($c \rightarrow E$)

Fig. 5. Scenario 2: Heterogeneous Caches – y-axis measures the performance difference ratio compared to the homogeneous scenario presented in Fig. 2 – the inner part of Fig. 5(a) is a zoom-in to the line-plots of the same figure

Summarising, we argue that these are the properties of *in-network caching* that the research community needs to investi-

⁴We note that in Fig. 5 we present the performance increase of the protocols compared to the previous scenario. This means that the performance of the algorithms in absolute terms is similar to the one presented in Fig. 2.

gate further, in order to exploit this inherent capability of ICNs for in-network caching. We have tested our algorithm in scale-free topologies with multiple servers and multiple replicas of contents. Our results show even greater differences in terms of redundant traffic elimination in favour of *ProbCache*. In case of CE^2 this owes to inefficient resource utilisation, while in case of LCD it owes to its inherent design to fit to directed graphs only, where the direction from servers to clients is a fixed hierarchy. In contrast, *ProbCache* exploits the knowledge provided by TSI and TSB and multiplexes content flows according to their respective path lengths.

V. RELATED WORK

Several recent studies have focused specifically on the properties of a network with in-network caches (e.g., [10], [22], [25]). In [22], the authors provide a comprehensive performance evaluation of in-network caching taking into account several parameters, such as content request distributions, the catalog size, and cache replacement policies. They conclude that content popularity is (by far) the most important parameter of all. Along the same lines, authors in [26] investigate the impact of traffic mix on the caching performance of a two-level cache hierarchy. They conclude that VoD content should be cached towards the edge of the network, while other types of content should be stored in large discs towards the core.

In a recent study [24], the authors show that having heterogeneously sized caches does not improve overall performance. We argue that this owes to the assumption of ubiquitous caching, that is, caching chunks in all routers along the content delivery path. As we have shown both in this paper (see Section IV-B, Fig. 5) and in our recent study [18], by caching a limited number of copies of a chunk in selected caches in the network, we achieve significantly higher gains. In particular, in [18], we show that the centrality of nodes in a given network topology gives valid evidence of which nodes are within the most number of paths. In turn, caching in those nodes increases the overall performance by up to 15% in realistic scale free topologies [18].

VI. CONCLUSIONS

We have argued that caching named chunks in network routers' DRAM memory, as opposed to caching large objects or files in proxy disks, calls for reconsideration of past approaches to caching. That is, in-network caching in ICNs has to happen in an uncoordinated and distributed fashion. We have proposed *ProbCache*, an algorithm that approximates the capability of paths to cache contents, based on path lengths, and multiplexes content flows accordingly. The ultimate goal of *ProbCache* is to utilise resources efficiently, reduce *caching redundancy* and in turn, *network traffic redundancy*. We have considered both homogeneous and heterogeneous cache sizes and have adjusted *ProbCache* to fit in both environments.

We report savings of up to 20% in server hits; 7-8% in the number of hops to hit cached contents; and reduction by an order of magnitude in cache evictions, which directly

translates to network traffic redundancy elimination by the same proportion.

Several issues in our design still remain open. For example, our formula can be applied to incentivise caching towards the edge of a domain and before an expensive transit link. We hope that our study will trigger further research on similar directions as the ICN research field matures.

REFERENCES

- [1] T. Koponen and et al., "A Data-Oriented (and beyond) Network Architecture," *SIGCOMM*, vol. 37, no. 4, pp. 181–192, 2007.
- [2] V. Jacobson and et al., "Networking Named Content," in *CoNEXT '09*. New York, NY, USA: ACM, 2009, pp. 1–12.
- [3] A. Ghodsi and et al., "Naming in content-oriented architectures," in *ACM SIGCOMM ICN Workshop*, 2011, pp. 1–6.
- [4] A. Anand, A. Gupta, A. Akella, S. Seshan, and S. Shenker, "Packet caches on routers: the implications of universal redundant traffic elimination," in *ACM SIGCOMM*, 2008, pp. 219–230.
- [5] A. Anand, C. Muthukrishnan, A. Akella, and R. Ramjee, "Redundancy in network traffic: findings and implications," in *SIGMETRICS '09*.
- [6] L. Breslau and et al., "Web caching and zipf-like distributions: Evidence and implications," in *In INFOCOM*, 1999, pp. 126–134.
- [7] L. Fan, P. Cao, J. Almeida, and A. Z. Broder, "Summary cache: a scalable wide-area web cache sharing protocol," *IEEE/ACM Trans. Netw.*, vol. 8, pp. 281–293, June 2000.
- [8] H. Che, Z. Wang, and Y. Tung, "Analysis and Design of Hierarchical Web Caching Systems," in *INFOCOM*. IEEE, 2001, pp. 1416–1424.
- [9] N. Fujita, Y. Ishikawa, A. Iwata, and R. Izmailov, "Coarse-grain replica management strategies for dynamic replication of web contents," *Comput. Netw.*, vol. 45, no. 1, 2004.
- [10] L. Muscariello, G. Carofiglio, and M. Gallo, "Bandwidth and storage sharing performance in information centric networking," in *ACM SIGCOMM ICN Workshop*, 2011, pp. 26–31.
- [11] U. Lee, I. Rimac, and V. Hilt, "Greening the internet with content-centric networking," *eEnergy*, p. 179, 2010.
- [12] E. J. Rosensweig and J. Kurose, "Breadcrumbs: Efficient, Best-Effort Content Location in Cache Networks," in *INFOCOM*, 09.
- [13] A. A. Jiang and J. Bruck, "Optimal content placement for en-route web caching," in *Proceedings of IEEE NCA '03*.
- [14] P. Krishnan, D. Raz, and Y. Shavitt, "The cache location problem," *IEEE/ACM Trans. Netw.*, vol. 8, no. 5, pp. 568–582, Oct. 2000.
- [15] D. Perino and M. Varvello, "A reality check for content centric networking," in *ACM SIGCOMM ICN Workshop*, 2011, pp. 44–49.
- [16] S. Arianfar, P. Nikander, and J. Ott, "On content-centric router design and implications," in *ReArch Workshop*, vol. 9. ACM, 2010, p. 5.
- [17] I. Psaras, R. G. Clegg, R. Landa, W. K. Chai, and G. Pavlou, "Modelling and Evaluation of CCN-Caching Trees," in *Proceedings of IFIP NETWORKING*, 2011, pp. 78–91.
- [18] W. K. Chai, D. He, I. Psaras, and G. Pavlou, "Cache 'Less for More' in Information-centric Networks," in *Proceedings of IFIP Networking, Prague, Czech Republic, May, 2012*.
- [19] W. K. Chai and et al., "Curling: Content-ubiquitous resolution and delivery infrastructure for next-generation services," *IEEE Communications Magazine*, vol. 49, no. 3, pp. 112–120, 2011.
- [20] N. Fotiou, G. C. Polyzos, and D. Trossen, "Illustrating a publish-subscribe internet architecture," *Journal on Telecommunication Systems, Springer, March*, 2011.
- [21] K. Katsaros, G. Xylomenos, and G. C. Polyzos, "Multicache: An overlay architecture for information-centric networking," *Computer Networks, to appear*, 2011.
- [22] D. Rossi and G. Rossini, "Caching performance of content centric networks under multi-path routing," *Technical Report*, 2011.
- [23] N. Laoutaris, H. Che, and I. Stavrakakis, "The lcd interconnection of lru caches and its analysis," *Perform. Eval.*, vol. 63, no. 7, Jul. 2006.
- [24] D. Rossi and G. Rossini, "On sizing CCN content stores by exploiting topological information," in *IEEE NOMEN Workshop*, 2012.
- [25] E. J. Rosensweig, J. Kurose, and D. Towsley, "Approximate models for general cache networks," in *IEEE INFOCOM*, 2010.
- [26] C. Fricker, P. Robert, J. Roberts, and N. Sbihi, "Impact of traffic mix on caching performance in a content-centric network," in *In Proceedings of IEEE NOMEN Workshop*, 2012.