# Energy Efficient Thresholds for Cached Content in Content Centric Networking


Sébastien Lafond
Department of Information Technologies
Åbo Akademi University
Turku, Finland
Email: sebastien.lafond@abo.fi

Tuan Anh Trinh
Department of Telecommunications and media Informatics
Budapest University of technology and Economics
Budapest, Hungary
Email: trinh@tmit.bme.hu



***Abstract*—In this paper we investigate the impact of caching content on Content Centric Networking (CCN) taking into account the energy efficiency of the caching strategy. We analyse the different threshold levels from which caching content starts to be beneficial based on two metrics, and take into consideration the used storage technologies and the location of the CCN router. Based on the presented analytical results we address practical issues regarding the possible deployment of CCN routers from an energy efficiency and performance trade-off point of view. This work lays down the foundation for the development of a model to design practical caching strategies in CCN taking into account the energy efficiency of different approaches.**


## I. Introduction

Due to the current explosion of the Internet traffic generated by content production and dissemination applications, like popular video and social networking services, we can observe a paradigm shift from host-centric networking towards content-centric networking. Host-centric networking is based on the assumption that content is centrally served by a reachable host assigned with a unique IP address. On the other hand content-centric networking is based on the principle that content is the focal point in the networking communication model (and not the serving host). In Content Centric Networking (CCN) content items have a unique identifier and can be transparently cached along the network topology. In CCN when content is cached in network elements closer to the network edge the communication delays to serve multiple requests to the same content as well as backbone network traffic will be reduced. This can enable, depending of the caching strategies, a somehow self-distribution of data along the network topology based on the content popularity.

Caching content can be done in CCN enabled router, and taking only into consideration the handling of data and requests, a CCN router is composed of three main elements: a forwarding Information Base (FIB), a Content Store (CS) and a Pending Interest Table (PIT)

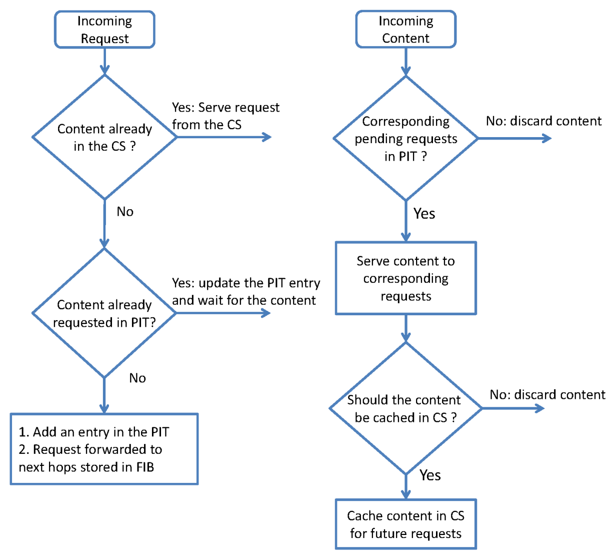


Fig. 1. Flowcharts describing incoming requests and content in a CCN router

[1]. The CS is used to buffer requested data, like in any other router, but can also store cached content for future requests. The PIT tracks the recently requested but not yet served content requests. The FIB is used to connect content name prefix with one or several following hop routers when the requested content is not present in the CS or in the PIT. Figure 1 describes the use of these elements when requests and data content arrive on the router interfaces.

So far much of the research effort has focused on architectural issues including naming and addressing as well as transport, caching, error and flow control in CCN. However when content is cached on network elements, the energy efficiency of such approach should also be evaluated. Basically, content need to be cached as long as it is cost efficient to keep it, i.e. as long as the cost of storing the data on a CCN router stays lower than the cost of re-transmission. The decision to add data into the cache is taken after serving the pending requests from the PIT as described on the left part of Figure 1.

Deciding which content need to be cached is usually based on the characteristics of incoming requests and content being requested. However, because a fast storage alternative is assumed to consume more energy per storage unit than slower memory technologies an optimal caching approach should be based on a trade-off between the memory technology, the level of content popularity, the energy consumption, and the total size of the CS. The position of CCN routers in a network has also a direct influence on the overall performance as the further the content is cached away from the user, the less duplication of content data is needed and overall less storage energy is required. On the other hand the closer the content is cached to the user the faster the content can be served and overall less transmission energy is required.

In this paper we analytically analyse the thresholds from which caching content starts to be beneficial based on two metrics considering the energy efficiency of content caching and taking into account the storage technologies and the location of the CCN router. In this paper CCN is based on the CCN model proposed by [2].

The remainder of the paper is organized as follow. In Section II we present the related work on energy efficiency of CCN. Section III introduces two possible metrics to evaluate the energy efficiency and performance trade-off of caching content in CCN. In Section IV we analyse the different thresholds from which caching content starts to be beneficial based on the metrics introduced in the previous Section, and put the results in the context of real access patterns to a popular video sharing service. Finally Section V provides a conclusion on the paper contributions.

## II. Related Work

Several CCN designs have already been proposed in the past decade [2]–[4] and few energy related aspects of CCN have been previously studied in the last few years [1], [5]–[8].

In [5] the energy efficiency of several content dissemination strategies is evaluated based on simple trace-based simulations and indicates that CCN can substantially improve the energy consumption of content dissemination. This paper looks at the energy efficiency on the overall network level for different fractions of CCN-enable routers and does not investigate the energy efficiency and related metrics for caching content into the CS.

In [1] the authors evaluate the suitability of existing software and hardware components in routers for the support of CCN. This work indicates that while CCN deployment is feasible at a Content Distribution Network and ISP scale, an Internet scale deployment is still impracticable due to technology limitations. In our work we re-use the results and assumption presented in [1] and focus on possible metrics to define thresholds for cached content in CCN.

In [6] and [7] the authors provide an assessment of the relative energy benefit of CCN over other content delivery architectures and the impact of different memory technologies on the overall energy efficiency. However this paper does not investigate when content should be cached based on the networking element locations and used memory technology.

This paper extends the preliminary results presented in [8] with a more extensive analysis of metrics to define thresholds from which caching content starts to be beneficial and takes into account the energy efficiency of the content caching strategy. We also put our results within the context of real access patterns observed on different network sizes.

## III. Energy efficiency and performance trade-off in CCN

One of the main benefits of CCN is the reduction of networking delay to serve content as fast as possible. Therefore the sole use of energy consumption perspectives to decide on the cached contend might lead to undesirable performance degradations and increased delays. To assess the overall system behaviour the trade-off between the energy efficiency and performance should be evaluated. The energy-delay-product (EDP) is a widely accepted metric to estimate a system performance when energy efficiency and delays are the two main system characteristics. In the following analysis we assume a set of clients retrieving content from a server through an intermediate CCN router as represented by Figure 2. All clients are $M$ hops away from the CCN router and $N$ hops away from the content server. For the sake of simplicity we assume that the network bandwidth and energy consumption per size unit is the same over all the links of the network.

Previous works on the energy efficiency of caching strategies in the content store of CCN routers [8] analysed the energy overhead of CCN storage compared to the data link energy consumption. This analysis leads to the estimation of the minimum average access rate of content objects in order to get an energy efficient caching policy. The minimum average access rate directly reflects the popularity of the data over a period of time. Because a) the power efficiency of data stored on different caching hardware technologies and b) the energy efficiency of data link are directly proportional to the data size [1], [9], the minimum average access rate is not relative to the size of the data and is defined as:

$$f > \frac{P_{storage}}{E_{link} \cdot (N - M)} \quad (1)$$

where $P_{storgare}$ represents the power dissipated per size unit to store the content in the cache (content store), and $E_{link}$ the energy to transport data per size unit across one link.

We extend the analytic analysis of the energy efficiency of content stored in CCN routers proposed in [8] and define the energy-delay-product for $K$ requests to the same non-cached, i.e. retrieved from the source, and cached content as:

$$EDP_{from\_source} = \overbrace{K \cdot N \cdot E_{link} \cdot DataSize}^{Energy} \cdot \underbrace{K \cdot N \cdot D \cdot DataSize}_{Delay} \quad (2)$$

$$EDP_{cached} = \overbrace{DataSize \cdot (K \cdot M \cdot E_{link} + T_{cached} \cdot P_{store})}^{Energy} \cdot \underbrace{K \cdot M \cdot D \cdot DataSize}_{Delay} \quad (3)$$

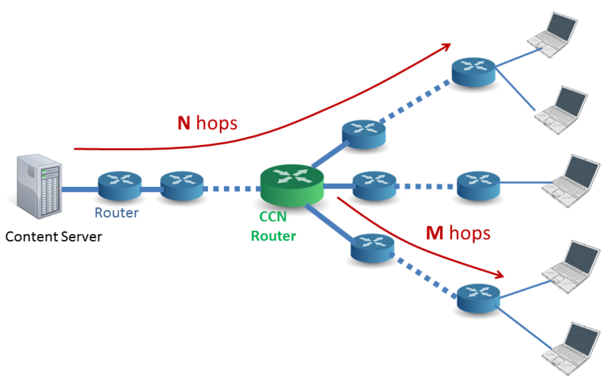


Fig. 2. Network scenario

where $DataSize$ represents the size of the content, $D$ the link delay per size unit to move data across it, and $T_{cached}$ the lifetime of the content in the cache.

Based on equations 2 and 3 caching content provides a lower energy-delay-product if:

$$\begin{gathered} EDP_{cached} < EDP_{from\_source} \\ M \cdot T_{cached} \cdot P_{store} < K \cdot E_{link} \cdot (N^2 - M^2) \\ \frac{T_{cached}}{K} < \frac{E_{link} \cdot (N^2 - M^2)}{M \cdot P_{store}} \end{gathered} \quad (4)$$

Based on the average time between two accesses to the same cached content $\frac{T_{cached}}{K}$ we can define the minimum average frequency access $f$ to the content such that caching provides a lower energy-delay-product as the following:

$$f > \frac{M \cdot P_{store}}{E_{link} \cdot (N^2 - M^2)} \quad (5)$$

## IV. Popularity Thresholds for Cached Content

The threshold from which content can be added on the content store of a CCN router can be based from a strictly energy efficiency perceptive or by taking into account the energy-delay-product. In both approaches the specified threshold defines the minimum number of accesses to one single content per time unit. This minimum average access frequency can be defined for different time windows depending of the typical content access pattern. Based on the differences in typical access patterns of user content generated in different online social networks [10] we use in this work an hourly based access frequency to define the popularity thresholds from which content can be cached. Using an hourly based average access frequency gives the opportunity to interleave popular contents having different, and sometime opposite, night and day access patterns. Nonetheless, we intent to extend our analysis with other time windows (like daily and weekly based average access frequency) as future work to evaluate its potential impact on the overall system performance and energy efficiency.

For different placement positions of the CCN router along the network described in Figure 2 we analysed the access frequency threshold from which caching content starts to be beneficial. This analysis is presented when a) only the energy efficiency of caching content and b) the resulting energy-delay-product is taken into account, based respectively on Equations 1 and 5. Additionally

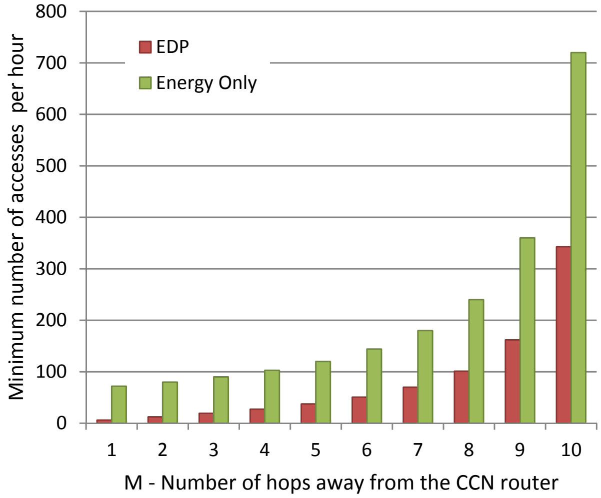


Fig. 3. Minimum access frequency depending on M for DRAM technology

TABLE I. Dissipated Power for Different Memory Technologies

| | DRAM | SRAM | TCAM |
|---|---|---|---|
| Dissipated power [W/MB] | 0.023 | 0.12 | 15 |

we also estimated the impact of the energy efficiency of different memory technologies on the frequency threshold. Table I presents the used memory technologies and their respective power dissipations per MB obtained from [1]. For all values we assumed the energy per size unit to transport data across a link, $E_{link}$, to be equal to 120 Joule per GB and the total number of hops $N$ in the network to be 11. The $E_{link}$ value is based on the energy efficiency of a Cisco CRS 1 series router [5].

Figure 3 presents the minimum access frequency per hour for different values of $M$ using the energy-delay-product and energy efficiency only approaches in case Dynamic Random Access Memory (DRAM) are used to store the cached content. In the same way Figures 4 and 5 represent the minimum access frequency per hour when using respectively Static Random Access Memory (SRAM) and Ternary Content-Addressable Memory (TCAM) as storage technology.

Figures 3, 4 and 5 exhibit the same patterns, that is: a) taking into account only the energy efficiency of cached content, the minimum number of accesses per hour is larger than when using as metric the energy-delay-product, and b) the further the CCN router is from the end-user the larger the access frequency threshold is. The fact that using as metric the energy-delay-product is less restrictive than taking into account only the energy efficiency of cached content can be clearly explained by Equations 1 and 5. In Equation 1 the frequency threshold depends on the distance between the CCN

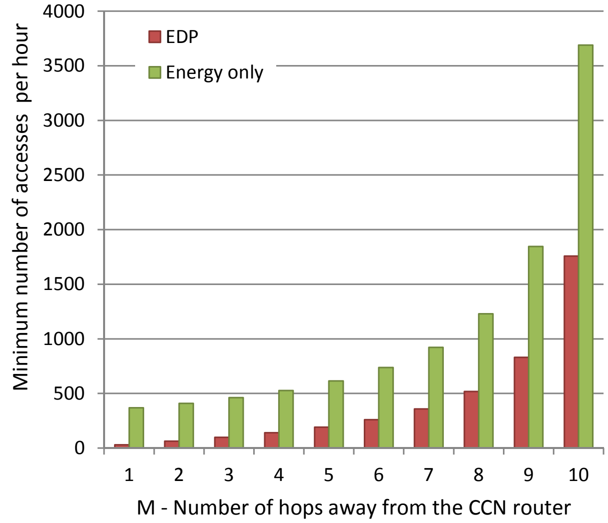


Fig. 4. Minimum access frequency on M for SRAM technology

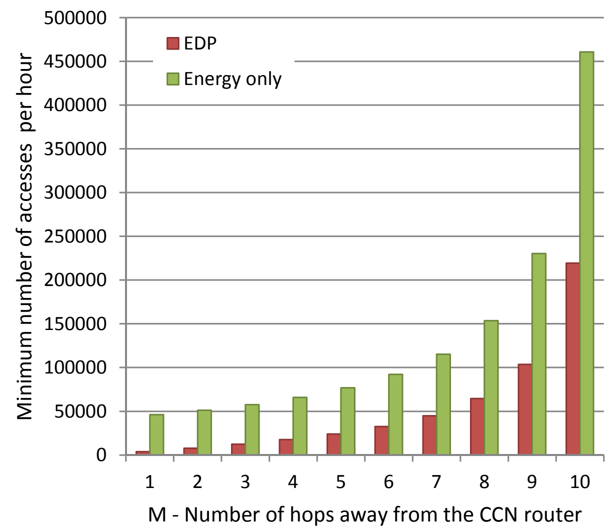


Fig. 5. Minimum access frequency depending on M for TCAM technology

router and the content server, while in Equation 5 the threshold depends on the difference between the squared number of total hops and the squared number of hops between the CCN router and user.

Moreover, the value of the frequency thresholds will depend to a large extent on the used memory technology. For example when placing the CCN router at one hop away from the user and using the energy-delay-product as metric the minimum access frequency per hour is 72 when using DRAM, 369 when using SRAM and 46080 when using TCAM. This can be explained by the relatively large difference in energy efficiency of the memory technologies. Furthermore, due to its preeminent energy efficiency, the use of Solid-State Drive (SSD) technology would lead to a corresponding minimum threshold of less than one access per hour for any position of the CCN router in the network and for both metrics. Nevertheless because of its large access time, a SSD disk could not sustain a large content store hit probability for any possible number of interest packet arrival rates and can only be used in edge routers [1]. In the same way DRAM based content store could be used on a link having up to 40 Gbps of bandwidth, and SRAM based content store could be used for links having up to 440 Gbps of bandwidth [1].

Taking into consideration the typical popularity of content on the Internet, we can analyse that the use of TCAM memory can never be energy efficient. For example on average the most popular video on YouTube has an overall access rate of up to 1 million per day, i.e. around 42000 accesses per hour [11], [12]. On Figure 5 this rate is even below the threshold when only energy efficiency is taking into account to cache content on an edge router ($M = 1$) using TCAM memory for its content store. As an edge router handles only a small fraction of the overall traffic, this observation completely ruled out the credibility of TCAM technology for energy efficient CCN router.

When looking at popular videos from YouTube at the level of a Canadian university campus consisting of 28,000 students and 5,300 faculty and staff [13] over a day the average number of view counts for the 100 most popular YouTube videos per hour is 875. Based on Figure 4 we can observe that it starts to be energy efficient to cache the 100 most popular YouTube videos on the campus if a CCN router using SRAM technology is at most 6 hops away from the users. Moreover, based on Figure 3 a CCN router using DRAM technology would be energy efficient to cache the 100 most popular YouTube videos on the level of the campus, up to 10 hops away from the users.

At the level of a Swedish municipality consisting of approximately 2200 households and 6800 end devices [14], the total number of requests to YouTube Video is around 600 per hour during its pick activity between 7 and 9pm. Therefore if we assume that half of the traffic originates from four equally popular video sources, based on Figure 4 all four videos can be energy efficiently cached in a CCN routed using DRAM for its content store only if it is one hop away from the user.

## V. Conclusion

In this paper we analysed the different threshold levels from which caching content starts to be beneficial based on two metrics and taking into consideration the used memory technologies and the location of the CCN router in the network. Based on the presented results we addressed practical issues regarding the possible deployment of CCN routers from an energy efficiency and performance trade-off point of view. We found out that based on the position of CCN routers and its traffic volume to popular content, different memory technologies and minimum access rate frequency should be selected in order to provide an energy and/or energy-delay-product efficient solution.
As future work we intent to more extensively analyse the overall energy characteristics of CCN routers taking into account the impact of caching strategies on not only the Content Store, but also all other elements when handling control and data requests. We also plan to extend this work taking into account multiple CCN routers and paths, along with the energy efficiency profiles of the routers.

## References


[1] D. Perino and M. Varvello, "A reality check for content centric networking," in *Proceedings of the ACM SIGCOMM workshop on Information-centric networking*, ser. ICN '11. New York, NY, USA: ACM, 2011, pp. 44–49. [Online]. Available: http://doi.acm.org/10.1145/2018584.2018596

[2] V. Jacobson, D. K. Smetters, J. D. Thornton, M. F. Plass, N. H. Briggs, and R. L. Braynard, "Networking named content," in *Proceedings of the 5th international conference on Emerging networking experiments and technologies*, ser. CoNEXT '09. New York, NY, USA: ACM, 2009, pp. 1–12. [Online]. Available: http://doi.acm.org/10.1145/1658939.1658941

[3] M. Gritter and D. R. Cheriton, "An architecture for content routing support in the internet," in *Proceedings of the 3rd conference on USENIX Symposium on Internet Technologies and Systems - Volume 3*, ser. USITS'01. Berkeley, CA, USA: USENIX Association, 2001, pp. 4–4. [Online]. Available: http://dl.acm.org/citation.cfm?id=1251440.1251444

[4] T. Koponen, M. Chawla, B.-G. Chun, A. Ermolinskiy, K. H. Kim, S. Shenker, and I. Stoica, "A data-oriented (and beyond) network architecture," *SIGCOMM Comput. Commun. Rev.*, vol. 37, no. 4, pp. 181–192, Aug. 2007. [Online]. Available: http://doi.acm.org/10.1145/1282427.1282402

[5] U. Lee, I. Rimac, D. Kilper, and V. Hilt, "Toward energy-efficient content dissemination," *Network, IEEE*, vol. 25, no. 2, pp. 14–19, 2011.

[6] K. Guan, G. Atkinson, D. Kilper, and E. Gulsen, "On the energy efficiency of content delivery architectures," in *Communications Workshops (ICC), 2011 IEEE International Conference on*, 2011, pp. 1–6.

[7] N. Choi, K. Guan, D. Kilper, and G. Atkinson, "In-network caching effect on optimal energy consumption in content-centric networking," in *Communications (ICC), 2012 IEEE International Conference on*, 2012, pp. 2889–2894.

[8] T. Braun and T. A. Trinh, "Energy efficiency issues in information-centric networking," in *Conference on Energy Efficiency in Large Scale Distributed Systems, EE-LSDS*, 2013.

[9] N. Choi, K. Guan, D. Kilper, and G. Atkinson, "In-network caching effect on optimal energy consumption in content-centric networking," in *Communications (ICC), 2012 IEEE International Conference on*, 2012, pp. 2889–2894.

[10] L. Guo, E. Tan, S. Chen, X. Zhang, and Y. E. Zhao, "Analyzing patterns of user content generation in online social networks," in *Proceedings of the 15th ACM SIGKDD international conference on Knowledge discovery and data mining*, ser. KDD '09. New York, NY, USA: ACM, 2009, pp. 369–378. [Online]. Available: http://doi.acm.org/10.1145/1557019.1557064

[11] M. Cha, H. Kwak, P. Rodriguez, Y.-Y. Ahn, and S. Moon, "I tube, you tube, everybody tubes: analyzing the world's largest user generated content video system," in *Proceedings of the 7th ACM SIGCOMM conference on Internet measurement*, ser. IMC '07. New York, NY, USA: ACM, 2007, pp. 1–14. [Online]. Available: http://doi.acm.org/10.1145/1298306.1298309

[12] A. Abhari and M. Soraya, "Workload generation for youtube," *Multimedia Tools Appl.*, vol. 46, no. 1, pp. 91–118, Jan. 2010. [Online]. Available: http://dx.doi.org/10.1007/s11042-009-0309-5

[13] P. Gill, M. Arlitt, Z. Li, and A. Mahanti, "Youtube traffic characterization: a view from the edge," in *Proceedings of the 7th ACM SIGCOMM conference on Internet measurement*, ser. IMC '07. New York, NY, USA: ACM, 2007, pp. 15–28. [Online]. Available: http://doi.acm.org/10.1145/1298306.1298310

[14] H. Wang, "Analyzing youtube content demand patterns and cacheability in a swedish municipal network," Master's thesis, KTH Royal Institute of Technology, 2012.