

AN ENERGY-EFFICIENCY ASSESSMENT OF CONTENT CENTRIC NETWORKING (CCN)

Muhammad Rizwan Butt, Oscar Delgado, Mark Coates

Department of Electrical and Computer Engineering
McGill University, Montreal, QC, Canada

Email: {muhammad.butt, oscar.delgadocollao}@mail.mcgill.ca, mark.coates@mcgill.ca

ABSTRACT

Content-Centric Networking (CCN) is a recently proposed networking architecture that can potentially lead to reduced bandwidth usage and better scalability and security as compared to the current IP-based architecture. In this paper, we conduct an energy consumption analysis of content-centric networking and IP-based networking for a video streaming scenario. We consider two types of energy consumption: the energy required to manufacture the network devices and the energy required for operation. We perform simulations of content-centric networking over a general-tree topology to assess the traffic rate reductions achieved by CCN's insertion of caches at routers. Although CCN network devices have a higher intrinsic energy consumption compared to the IP-based devices because of the presence of additional memory, by exploiting their caching capabilities it is possible to reduce the overall energy consumption of the network. We consider both the incorporation of an online rate adaptation mechanism as well as a static network provisioning approach and observe that these approaches can lead to an energy consumption reduction of 10-20 percent.

Index Terms— CCN, Energy analysis, Energy gain.

1. INTRODUCTION

A major portion of Internet traffic is generated by video streaming and Video on Demand (VoD) applications. The host-oriented IP-based architecture is not well-matched to these types of applications, and substantial research efforts have been made in recent years to explore data-centric networking architectures. See [1] for a review of the developments in this field and a discussion of commonalities and differences of the proposals. One of the prominent approaches is content-centric networking (CCN) [2]. In this paper, we provide a comparative study of CCN and IP-based networks in terms of energy consumption. As it constitutes one of the more compelling applications for the CCN architecture, we consider a dedicated video streaming application. Our energy analysis includes both the consumption during operation and manufacture. We use the notion of energy [3] to embody the latter expenditure.

The network routers in the CCN architecture are provisioned with additional memory and have the capability to cache content. The manufacture and operation of this memory is an additional energy expenditure compared to the IP network. This implies that if both network architectures are operated in the same fashion, the CCN approach will consume more energy to deploy and operate. The caching of content reduces the traffic on many links, suggesting that links and routers in the CCN architecture can operate at a lower rate compared to the IP-based implementation. This can

lead to a decreased energy consumption [4], under the reasonable assumption that devices operating at lower rates require less energy. In previous work [5], it has been suggested that energy can be saved by periodically switching network devices into sleep mode. We consider that this is not a feasible option for a streaming network application, since it can introduce unacceptable delays. We investigate the impact of two alternative strategies on energy consumption. The first involves on-line rate adaptation, in which routers are capable of adjusting their operating rate according to the bandwidth demand. The other is a static approach, where links and routers are statically provisioned according to the anticipated load. We have implemented a simulator for content-dissemination in CCNs and we use this to estimate the traffic rate on each link during operation of the VoD application. This allows us to calculate energy consumption for the two proposed strategies.

The work that is most closely related to our work is that of Lee et al. [6], in which Lee et al. propose an architecture for an energy-efficient CCN router and claim that energy benefits can be reaped by deploying these routers incrementally throughout the Internet. The energy analysis in [6] only considers the energy cost when clients are downloading content and all the network devices are operating at their peak power ratings. There is no consideration of the manufacturing energy cost nor of the extra energy required to power the additional cache memory of the CCN routers. Lee et al. suggest that energy savings can be obtained by switching the CCN routers on or off according to the network traffic requirements.

Other work that assesses energy consumption does not involve CCNs. Seetharam et al. in [7] provide a green (carbon footprint and energy consumption) analysis comparing an IP-network movie streaming application with a postal-service movie distribution system. They provide a detailed discussion about methods for estimating the operational and manufacturing energy consumption of network devices. Raghavan et al. in [3] provide an approximate analysis of the Internet's energy consumption. They calculate both the operating energy consumption and the energy (embodied energy), the energy required to construct the Internet. This energy is equivalent to the manufacturing energy consumption discussed by Seetharam et al. in [7].

The rest of the paper is organized as follows. In Section 2 we develop our energy model. Section 3 describes the evaluation of our model via simulations and Section 4 concludes the paper.

2. ENERGY MODEL

2.1. Assumptions and notations

Consider a general tree-based network topology composed of multiple edge routers and a single server running a video streaming

service. Assume we have two similar networks with the above characteristics, one designed according to a CCN architecture and the other according to an IP-based architecture. The routers involved in these two networks have one important difference: CCN routers have an additional cache to store data passing through them.

We consider two main aspects of the energy consumption in a network: energy and the operating energy consumption. It is common to use the power-rating when analyzing the energy consumption during the operation of a network device [3]. However, energy is expressed in joules so it is unable to capture the time variation. To be consistent in terms of units we define the *embodied power* which is calculated by dividing the energy by the life cycle of a device. Table 1 summarizes the parameters employed in our analysis.

An important component of our analysis is the impact of link

Parameter	Description
K	Total number of routers
M_{IP}, M_{CCN}	IP and CCN router embodied power
S_M, S_O	Server embodied and operating power
C_S	Embodied power consumption for server storage
C_M	Embodied power consumption for router cache
O_{IP}, O_{CCN}	IP and CCN router operating power
$R_{IP}(k)$	Link rate at IP router k
$R_h(k)$	Link rate threshold at CCN router k

Table 1. Parameters used for the analysis

rate on energy consumption. In our study we assume that the power consumption of a router varies in proportion to the link rate [5], [4]. We define power reduction factor ϕ_k , $k = 1, 2, \dots, K$ as:

$$\phi_k = \frac{R_h(k)}{R_{IP}(k)} \quad (1)$$

2.2. Power consumption

Although embodied energy consumption is a one time cost, it is important to include this value in our analysis in order to have a fair comparison between the energy consumptions of IP and CCN networks for VoD streaming applications. Equation (2) and (3) describes the embodied power cost of IP and CCN networks respectively.

$$P_{IP}^M = KM_{IP} + S_M \quad (2)$$

$$P_{CCN}^M = KM_{CCN} + S_M \quad (3)$$

We assume that the manufacturing energy needed to produce a server is not related to the network architecture. We identify the power consumption for both IP and CCN based networks as:

$$P_{IP}^O = C_S + S_O + KO_{IP} \quad (4)$$

$$P_{CCN}^O = C_S + S_O + KC_M + O_{CCN} \sum_{k=1}^K \phi_k \quad (5)$$

Equation (5) introduces the use of the power reduction factor ϕ_k . A discussion on ϕ_k is presented in subsection 2.3.

The total power consumption is the sum of the embodied and the operational power consumption of the network. Equations (6) and (7) describe the total power consumption of the IP and CCN networks respectively.

$$P_{IP} = P_{IP}^M + P_{IP}^O \quad (6)$$

$$P_{CCN} = P_{CCN}^M + P_{CCN}^O \quad (7)$$

2.3. Rate adaptation (RA)

We consider three scenarios. In the first scenario, the IP and CCN use routers with no RA capability (*without RA*) so $\phi_k = 1$. In the second scenario, the IP and CCN networks are provisioned differently, with links provided with sufficient capacity to support the generated traffic demand (rounded to the nearest 500 Mbps in our analysis *practical RA*). In the third scenario, there is on-line power adaptation, with link rates being adjusted relatively rapidly in response to changes in the traffic demand (*with RA*). Due to strict delay requirements we cannot use the strategy of [4] for the last scenario so we incorporate a strategy developed by Chamara et al. for ethernet networks [8]. Each switch monitors the traffic flowing through it and this information is used to identify traffic thresholds. When the traffic exceeds or falls below these thresholds, the link rate is adjusted to a higher or lower value accordingly, and this in turn impacts the consumed power of the device. As already introduced in equation (1) the rate factor ϕ_k allows us to quantify the relation between the link rate and the power consumption.

3. PERFORMANCE EVALUATION

We consider a VoD scenario where a single server is providing service to the clients. For our analysis we use a video server with 10 TB storage capacity and M10i routers [9] (with additional cache). We calculate the approximate embodied and operating powers of the network devices. To the best of our knowledge, there is no study available in the literature which provides an exact analysis of the energy or power consumed in manufacturing storage devices, servers and routers. Therefore we estimate these costs from data given in [3], [7], [10] and making use of the available data-sheets for different devices which we are considering in our analysis.

3.1. Embodied power consumption

In the year 2000 the energy of a disk drive of 30 GB was 2926 MJ [10]. According to Kryder's law [11], the storage capacity of storage devices doubles every 18 months. Hence in 2011 the amount of storage which can be manufactured using 2926 MJ of energy comes out to be 5TB. The energy of a 10 TB storage device at the server can then be approximated as $2926 \times 2 = 5852$ MJ. The size of the CCN router cache memories should depend on the scale of the network. We consider four candidate sizes to assess the difference the choice makes on the overall power consumption of the CCN network. Hence the energies for 256, 128, 96 and 64 GB caches are 150, 55, 75 and 37 MJ respectively. We assume that the life cycle for all the devices we are using in our analysis is 3 years [7]. Hence the embodied power for a these different size caches can be calculated by dividing the energy by the life cycle.

The server energy estimated by Seetharam et al. in [7] using the study of Williams in [10] is 550 MJ ($= 6$ J/s). We use the same estimate for our analysis. We are considering a simplistic dedicated video streaming scenario so we can assume that the network is populated with the same type of routers throughout (Juniper M10i edge routers in this case). Seetharam et al. estimated the energy of a router in [7] by scaling the weight of the router relative to the weight of the PC. The weight of a desktop PC varies between 13 to 35lbs [12]. For this analysis we assume that it is approximately 25lbs. The weight of a M10i router is 79lbs [9] therefore the energy

of an edge router becomes $\frac{79}{25} \times 550 \times 10^6 \approx 1200\text{MJ}$, which is equal to 13 J/s (embodied power). The total embodied power for server and network routers is summarized in Table 2. We are considering the same M10i routers for both CCN and IP-based networks. The only difference is that the routers in CCN have an additional cache memory.

The difference between the values estimated in Table 2 for server

Device	Power (J/s)
S_M	68
M_{IP}	13
M_{CCN} 256, 128, 96, 64 GB	16, 15, 14, 13.5
S_O	731
O_{IP}	116
O_{CCN} 256, 128, 96, 64 GB	121, 120, 120, 196
C_M 256, 128, 96, 64 GB	0.053
C_S	20

Table 2. Power consumption estimates

and router’s embodied power is of the order of 4 to 6 times. These values are consistent with the values estimated by Raghavan et al. in [3]. In the next section we estimate the power consumed in transmitting a movie through the IP and CCN-based network.

3.2. Operation power consumption

We assume that the library size for the server is $M = 1,000$ movies. Each movie is compressed using DiVX codec and the average size is 700MB. The server we are using has a storage capacity of 10 TB consisting of an array of ten 1 TB devices (e.g., RAID configuration). The maximum downstream rate the server can provide is bottlenecked by the read capacity of memory storage (1385 Mbps approx 175 MB/s in this case). The power consumed by 10 TB storage is 480 J/s [13] during streaming operation. The operating power consumption for 256, 128, 96 and 64 GB cache memory is 4.6, 4.2, 3.8 and 3.2 J/s respectively [13]. The total energy spent by a server’s chip-set during video streaming can be calculated by following the model adopted in [7]. For a typical server the operating power for a multi-media streaming application is 251 J/s. The chip-set of edge routers we are using for this analysis (i.e., M10i) operate at 116 J/s. For the IP-based network the router operating power is the only power consumed in operating the M10i routers but for CCN the router operating power also includes power consumed by the cache memory during its operation (see Table 2). We also consider the power required to keep alive the router cache and server’s memory storage for the duration of the network deployment. These values are extracted from the available data-sheet of the memory devices [13], as shown in Table 2.

3.3. Simulation Model

We implemented the basic relevant functionalities (caching, routing, fetching) of content-centric networking in MATLAB. We generate a general tree topology with N levels, and extend the binary tree model developed in [14]. The simulator provides three main services: request generation, router caching and routing. The request arrival process is modeled through a Markov Modulated Rate Process (MMRP) that captures the behavior of the system at both the content level and the chunk level. The cache process

in each router employs a Least Recently Used (LRU) replacement policy. There is a fixed routing path between child and parent node.

We consider a set of M different contents equally partitioned into D classes of popularity. We assume a Zipf popularity distribution [14], hence $q_d = c/d^\gamma$, $d > 1$ and parameter $\gamma > 1$. Each content request in class d is generated according to a Poisson process with rate $\lambda_d = \lambda \cdot q_d$. This generation coincides with the request of the first chunk and the contents are uniformly chosen among the $m = M/D$ different content items in class d . The inter-arrival times between consecutive chunks are deterministic, and the number of chunks in a given content item is assumed to have a geometric distribution with average size σ . Each node in the network has a cache of size X .

3.4. Simulation Results

We consider that $M = 1,000$ content items are organized in $D = 40$ classes of popularity with parameter $\gamma = 2$, each one with $m = 250$ content items. Chunks are of size 1 MB and the average number of chunks is $\sigma = 700$. We are assuming that network is operating at peak load ($\lambda = 1$ content/sec) for the duration of its life time. The runtime for simulations is 1 hour. One important

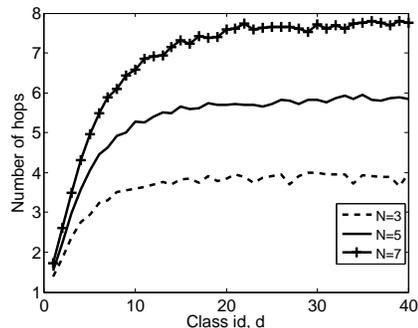


Fig. 1. CCN average number of hops per content request as a function of the popularity distribution class ID ($d = 1$ is the most popular content class), $X = 128$ GB, $\gamma = 2$, $\lambda = 1$ request/sec

aspect to consider in a CCN network is the average number of hops that a request has to travel in order to reach content. Fig. 1 shows the average number of hops as a function of the popularity class for a binary tree network with 3, 5 and 7 levels. It can be seen that content with the lowest class ID (very popular content) can be reached by traversing a smaller average number of hops (1.5 hops for $d = 1$). This indicates the fact that CCN link rates are smaller than IP link rates due to the caching process at each router.

To evaluate the energy performance we use the notion of rate adaptation introduced in Sec. 2.3. Fig. 2 confirms our intuition that at higher levels in a general tree topology, CCN with RA is energy efficient. The figure shows the power ratio for the cases when each router in the network has two or four children. When the nodes have more children, the energy savings at the higher levels of the tree are more pronounced. Even if we do not use RA the power ratio (P_{IP}/P_{CCN}) is very close to 1, indicating that the energy cost associated with manufacturing and powering the additional cache memory is relatively small. The practical RA implementation ensures energy gain for CCN at higher levels in the network but it has poorer behavior at lower levels.

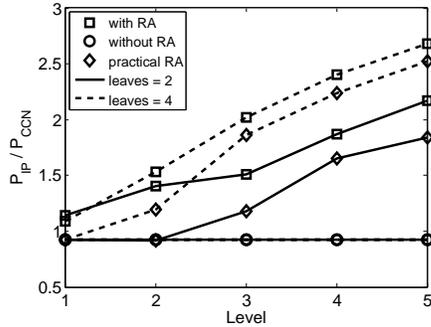


Fig. 2. Power ratio at each level of the general tree topology, $N = 5$, $X = 128$ GB

To better understand the energy behavior of CCN, we also present the power ratio as a function of network and cache size. In Fig. 3 we observe that as the network grows the power ratio improves when using RA or practical RA, because more routers benefit from lower link rates. In the absence of RA the energy consumption of CCN is slightly greater than IP as a consequence of energy cost associated with the CCN router's cache memory.

Larger caches reduce the link rates at higher levels in the network but they require more energy to manufacture and operate. Fig. 4 shows that it is possible to find a size of the cache that maximizes the power ratio. For the specific settings in our simulations 128 GB appears to maximize the power ratio.

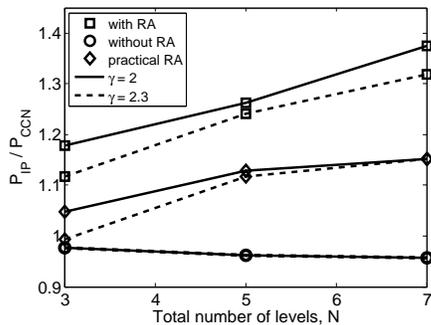


Fig. 3. Power ratio as a function of the size of the network (in terms of the number of levels in binary tree topology), $X = 128$ GB

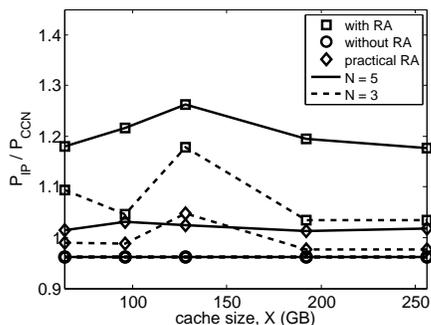


Fig. 4. Power ratio corresponding to different cache sizes, $N = 5$, 3

4. CONCLUSION

We have presented a comparative study of energy consumption for CCN and IP-based networks in a video on demand scenario over a tree topology. Our simulation results show that if we do not use any rate adaptation technique CCN networks require only slightly more energy to manufacture and operate than IP based networks. A simple implementation of rate adaptation can make CCN an energy-efficient architecture, leading to energy savings over the lifetime of the network of 10-20 percent.

5. REFERENCES

- [1] A. Ghodsi *et al.*, "Information-centric networking: seeing the forest for the trees," in *Proc. ACM Workshop Hot Topics in Networking*, Cambridge, MA, USA, Nov. 2011.
- [2] V. Jacobson *et al.*, "VoCCN: voice-over content-centric networks," in *Proc. Workshop Re-architecting the Internet*, Rome, Italy, Dec. 2009.
- [3] B. Raghavan and J. Ma, "The Energy and Emery of the Internet," in *Proc. ACM Workshop Hot Topics in Networks*, Cambridge, MA, USA, Nov. 2011.
- [4] S. Antonakopoulos, S. Fortune, and L. Zhang, "Power-aware routing with rate-adaptive network elements," in *Proc. GLOBECOM Workshops*, Miami, Florida, USA, Dec. 2010.
- [5] S. Nedeveschi *et al.*, "Reducing network energy consumption via sleeping and rate-adaptation," in *Proc. Symp. Networked Systems Design and Implementation*, Berkeley, CA, USA, Apr. 2008.
- [6] U. Lee, I. Rimac, and V. Hilt, "Greening the internet with content-centric networking," in *Proc. Int. Conf. Energy-Efficient Computing and Networking*, Passau, Germany, Apr. 2010.
- [7] A. Seetharam *et al.*, "Shipping to streaming: is this shift green?" in *Proc. ACM SIGCOMM Workshop Green Networking*, New Delhi, India, Aug. 2010.
- [8] C. Gunaratne *et al.*, "Reducing the energy consumption of ethernet with Adaptive Link Rate (ALR)," *IEEE Trans. Comput.*, vol. 57, pp. 448–461, 2008.
- [9] Juniper Networks, "M10i and m7i multiservice edge routers," [Online]. Available: <http://www.juniper.net/us/en/local/pdf/datasheets/1100032-en.pdf>.
- [10] E. Williams, "Energy intensity of computer manufacturing: Hybrid assessment combining process and economic input-output methods," *Int. J. Envr. Sci. and Tech.*, vol. 38, pp. 66–74, 2004.
- [11] C. Walter, "Kryder's Law," *Scientific American*, pp. 32–33, Aug. 2005.
- [12] H. Duan *et al.*, "Life cycle assessment study of a chinese desktop personal computer," *Science of the Total Environment*, vol. 407, no. 5, pp. 1755–1764, 2009.
- [13] Seagate, "Data-sheet Momentus 5400.4," [Online]. Available: http://www.seagate.com/docs/pdf/ds_m_5400.pdf.
- [14] G. Carofoglio *et al.*, "Modeling data transfer in content-centric networking," in *Proc. Int. Teletraffic Congress*, San Francisco, California, USA, Sept. 2011.