Toward Energy-Efficient Content Dissemination

Uichin Lee, KAIST KSE
Ivica Rimac, Daniel Kilper, and Volker Hilt, Alcatel-Lucent Bell Labs

Abstract

A major role of today's Internet is to provide efficient content dissemination among users, such as distributing multimedia content and sharing user generated data. To meet the ever increasing demands, the Internet has been rapidly growing, and it now includes a web of tens of millions of networked devices ranging from content servers to core and edge routers to home gateways. Due to the sheer numbers, however, it is reported that these devices, such as those used for content delivery, consume a considerable amount of energy. While optimizing the energy efficiency of data centers is well studied in the literature, understanding the energy efficiency of various content dissemination strategies has received comparatively little attention thus far. In this article we review existing content dissemination architectures and survey the energy efficiency of various network devices used for content delivery. The energy efficiency comparison using simple trace-based simulations reveals that a change from a host-oriented to a content-centric networking model can substantially improve energy efficiency of content dissemination. Our preliminary results are encouraging and will stimulate further research in this direction.

The vast majority of current Internet usage consists of content being disseminated from a source to a number of users, ranging from distributing conventional multi-media data (e.g., IPTV, Hulu, and Netflix) to sharing user generated data over the web such as text, image, and video data (e.g., Facebook, Twitter, and YouTube). Growing demands for content distribution led content service providers like Google, Yahoo, and Microsoft to invest in large data centers with hundreds of thousands of machines distributed across different geographic regions [1].

Due to the sheer size, it is reported that data centers consume significant energy — the U.S. Environmental Protection Agency (EPA) estimates that servers and data centers could consume 100 billion kWh at a cost of $7.4 billion per year by 2011. A number of different approaches have been proposed to address this problem, such as introducing efficient cooling methods, dynamically provisioning servers to account for diurnal usage patterns, and scaling power consumption of servers proportional to their utilization (also known as energy-proportional computing) [2, 3].

Along with the expansion of data centers, backbone network providers have been increasing network capacity to meet the ever increasing demands, by deploying a large number of high-speed routers and fiber transmission systems. For more efficient content delivery, large content service providers also build their own private networks that interconnect their data centers [1]. The Internet has been rapidly growing and now includes a web of tens of millions of networking devices ranging from routers to home gateways, thus consuming considerable energy overall.

However, the energy efficiency of current networking devices is reported to be very poor, and their power consumption does not significantly vary according to utilization [4]. Recently much attention has been focused on improving the energy efficiency of networking devices (by reducing energy losses and overheads), and realizing so called *energy-proportional networking* where the energy consumption is proportional to utilization of network interfaces. Thus far, several techniques have been proposed toward this goal [5] such as dynamic voltage and frequency scaling (DVFS) of line cards, energy-efficient sleep modes to exploit intermittent idle durations, and coordination of intermediate routers for batch processing to prolong idle durations.

The capability of energy-proportional computing and networking could save a considerable amount of energy for Internet-scale content dissemination. Yet as user demand grows, overall traffic increases. Energy consumption will likewise increase unless the efficiency of the network improves proportionately. At present, backbone network traffic growth per year is in the range of 40–50 percent, whereas network equipment efficiency has only been improving by 10–20 percent [6]. Network energy consumption is largely determined by trends related to Moore's Law, which is slowing with regard to power [6].

One possible method of alleviating this problem is to push content to the network edge as in content distribution networks (CDNs) such as Akamai and Limelight. While CDNs were originally proposed to improve user-perceived performance in terms of delay and throughput, as a byproduct they help significantly reduce transit traffic in the network backbone, thereby saving energy used for data transport. An extreme case to this end would be a nano data center (NaDa) for which an Internet service provider (ISP) coordinates nano servers in users' home gateways to distribute content in a

0890-8044/11/$25.00 © 2011 IEEE

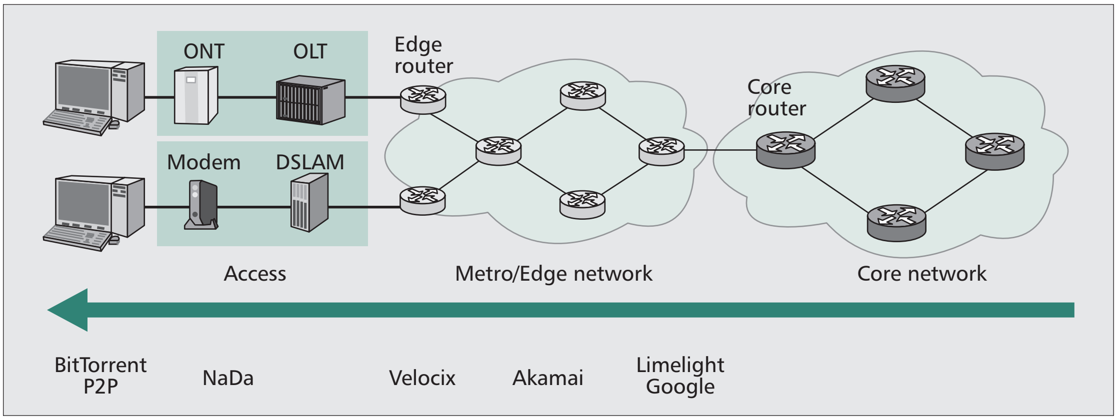

Figure 1. *Comparison of various content dissemination methods.*

peer-to-peer (P2P) fashion [7]. However, networking devices have a wide spectrum of energy consumption for packet forwarding (in terms of joules per bit), and surprisingly, it rapidly goes up as packets travel toward end users. For instance, when delivering the same amount of data, home gateways and desktop PCs via digital subscriber line (DSL) links could consume 100 and 1000 times more energy than core routers, respectively. Thus, distributing content from these less energy-efficient machines may result in higher energy consumption even with minimal transit traffic across networks.

Noting that routers are far more energy efficient than servers and access devices, a radical alternative is to introduce *content routers* into the network as proposed in content-centric networking [8]. Content-centric networking is a new network architecture for content dissemination that replaces conventional host-to-host conversations with named data oriented communications. The basic operation of a content router is very similar to that of an IP router, but the key departure is that it supports name-based routing and content caching throughout the network. This way, content-centric networking obviates the need of deploying separate, application specific mechanisms such as CDNs and P2P networks, which require dedicated infrastructure components and overlay mechanisms (e.g., for mapping named content to hosts).

In this article we begin by reviewing existing content dissemination methods and summarizing our evaluation results on energy efficiency of networking devices ranging from edge/core routers to home gateways. We then illustrate content-centric networking and present energy-efficient content router configurations for various networking devices. The energy efficiency of various content dissemination strategies is evaluated via simple trace-based simulations. By considering various incremental deployment scenarios, we show that content-centric networking can substantially improve the energy efficiency of content dissemination. Our preliminary results are encouraging and will stimulate further research on energy-aware content-centric networking.

Review of Content Dissemination Architectures

A content distribution network (CDN) maintains content servers located in multiple sites such as backbones and ISP points of presence (PoPs). When a user makes a request, the CDN chooses a server so as to improve user-perceived performance in terms of delay and throughput. The current design of CDNs can be differentiated based on where the content servers are located. One approach is to build large data centers at a few strategic locations close to the PoPs of many large ISPs and to connect them using private high-speed links. This method is employed by commercial CDN solutions such as Limelight, and by several content service providers such as Google and Amazon [1]. Another approach mainly used by Akamai is to deploy a large number of small content server clusters scattered across the Internet in multiple ISP PoPs and backbones. Since servers are highly distributed, locating a proper content server requires sophisticated algorithms such as real-time server monitoring and extensive network measurement.

As an alternative to relying on shared CDNs such as Akamai and Limelight, service providers could deploy their own ISP-level CDN solution in their networks (e.g., Velocix Digital Media Delivery Platform) to bring multimedia content closer to their subscribers, enabling them to offer differentiated video stream quality and faster file downloads. In this case, the content servers are located even closer than highly distributed CDNs such as Akamai, but it is more amenable to support optimized content delivery since ISPs have full knowledge of their networks.

The extreme case of a highly distributed approach is P2P content distribution where content servers are located at the customer premises. This includes file swarming (e.g., BitTorrent and eMule), and P2P video streaming (e.g., PPLive and Zattoo), and nano data center (NaDa), a distributed content distribution platform based on nano servers in home gateways. In P2P content distribution, content is divided into a number of small pieces, and peers cooperatively share whatever pieces they have. Peers sharing the same content can discover one another via a central server such as a tracker in BitTorrent, or a distributed hash table (DHT) such as Kademlia DHT in eMule. NaDa uses BitTorrent-like file swarming over nano servers in home gateways, yet the key difference is that an ISP manages and coordinates nano servers (called a managed P2P system). NaDa does not suffer from free-riding, node dynamics, and lack of awareness of underlying network conditions. The key benefit is that NaDa can efficiently provide traffic localization resulting from the collocation between nano servers and end users (e.g., within the same access or ISP network).

Energy Efficiency of Networking Devices

Delivering content from servers to end users involves various networking devices in between. We survey energy efficiency of network devices with an underlying assumption of energy-proportional computing and networking. We consider backbone routers, edge routers, and access network devices such as digital subscriber line (DSL) and gigabit passive optical network (GPON).

Given that a device has the nameplate power consumption of P_N and idle power consumption of P_0, we model power

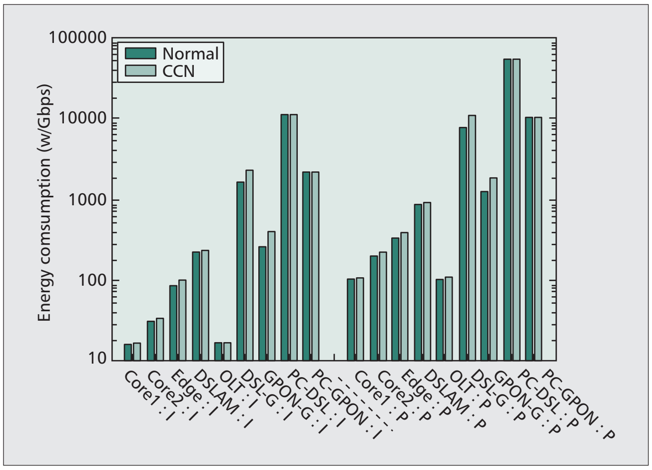

Figure 2. *Energy consumption of network devices (I: ideal scenario; P: practical scenario).*

consumption of a device under load ℓ as $P(\ell) = \alpha(P_0 + \ell \times (P_N - P_0)/R_{\max})$ where $R_{\max}$ is the maximum forwarding rate, and α represents additional overhead such as external power supplies and cooling requirements. Given a set of parameters (i.e., nameplate power P_N, idle power P_0, average load ℓ, and additional overhead α), we capture the energy efficiency of a device by dividing the power consumption by the load (i.e., $P(\ell)/\ell$). The resulting unit is Watts per gigabit per second — the amount of energy for forwarding a gigabit of data. In the following analysis, we first consider ideal scenarios of perfect energy-proportional networking where we have $P_0 = 0$ and $\alpha = 1$, and the energy efficiency is simply given as $P_N/R_{\max}$. We then show the results under pathological (yet practical) scenarios, such as high idle power consumption (95 percent of nameplate power) and lower utilization due to overprovisioning.

For the backbone, we consider a Cisco CRS 1 series router with 8 slots in a single shelf. There are 8 line cards each of which supports 40 Gb/s of data forwarding. The maximum forwarding rate is 320 Gb/s, and its nameplate power is 4834 W. CRS 1's ideal energy efficiency is 15 W/Gb/s. This is slightly above the 12 W/Gb/s industry average for 2008 for core routers compiled by Tamm *et al.* [6]. A less powerful backbone router Cisco GSR 12000 has 7 line card slots and supports rate up to 27Gb/s at the nameplate power consumption of 800 W [4]. GSR's energy efficiency is 28.6 W/Gb/s. Cisco 7507, an edge router has 5 line cards and supports forwarding rate up to 5Gb/s at the nameplate power consumption of 400 W. The 7507's energy efficiency is 80 W/Gb/s. This result shows that edge routers tend to spend more energy per bit than core routers. The reason is that while core routers are designed with one purpose in mind — to move traffic as quickly as possible — the requirements at the edge are diverse and complex, mandating power-hungry packet processing such as multiservice support (asynchronous transfer mode [ATM], frame relay, IP, multiprotocol label switching [MPLS], etc), quality of service, accounting and billing, and network management.

Both DSL and GPON need multiplexing equipment at the service provider's central office, which are called digital subscriber line access multiplexer (DSLAM) and optical line terminal (OLT), respectively. Moreover, users access networks via home gateways such as a DSL modem for DSL and an optical network terminal (ONT) for GPON. For a DSLAM, we consider the Zyxel IES-500M, which has 10 slots supporting 48 ADSL ports. Assuming a maximum speed of 10 Mb/s per line, the aggregated rate is 3.85 Gb/s at the nameplate power consumption of 800 W, and IES-500M's energy efficiency is 208.3 W/Gb/s. A typical DSL modem (e.g., D-Link DSL 2320B) consumes around 15 W, and its energy efficiency is 1536 W/Gb/s. Unlike a DSLAM, a GPON OLT can deliver a much higher aggregated rate. Fujitsu FA2232U has 16 fiber slots, each of which accommodates 32 subscribers (50 Mb/s per user). FA2232U's maximum rate is 16 Gb/s at the nameplate power of 400 W (as each fiber slot has a Gigabit Ethernet uplink), and its energy efficiency is 25 W/Gb/s. Due to the sheer aggregated traffic volume, ISPs (e.g., Verizon) connect OLTs directly to a set of collector rings, which is an edge optical network with reconfigurable optical add/drop multiplexers (ROADMs). Optical transmission hardware is an order of magnitude more efficient than a core router and therefore is neglected here [6]. The efficiency of a GPON gateway is much better than that of a DSL modem. For instance, Allied Data's GPON Gateway consumes 12 W at the speed of 50 Mb/s, and its energy efficiency is 245.8 W/Gb/s.

Content servers could be blade servers in a data center, home PCs, and nano servers in home gateways. Note that while a blade server in a data center could fully utilize its maximum server throughput (under no oversubscription scenarios), both home PCs and nano servers are limited by their uplink rates. A blade server with Xeon processors has an energy efficiency of 360 W/Gb/s [7], which is comparable to that of a nano server in a GPON gateway (245.8 W/Gb/s). An Intel dual core (2.4 GHz Core2Duo) OptiPlex 745 with 2 Gbytes RAM running Windows Vista consumes around 100 W in the normal idle state, and the energy efficiency of an access connection is at most 10 kW/Gb/s and 2 kW/Gb/s for DSL and GPON, respectively.

The aforementioned calculations can be extended to pathological scenarios with high idle power consumption and low utilization. For comparison, we consider a pathological case with the following assumptions. It is reported that the idle power consumption of current generation networking devices is as high as 95 percent of the actual maximum power draw [4]. Core, edge, and access nodes (OLT/DSLAM) are overprovisioned, and their utilization is 30, 50, and 50 percent, respectively, whereas end users are oversubscribed, and their maximum rate is 20 percent of the advertised rate [9]. The additional overhead such as cooling for network and perimeter devices is given as $\alpha = 2$ and $\alpha = 1$, respectively. Figure 2 shows the results. The overall trend of both cases is quite consistent; i.e., energy consumption exponentially increases as packets travel toward end users.

Toward Content-Centric Networking

Content servers (in data centers or in users' premises) are an order of magnitude less energy efficient than networking devices such as core/edge routers and optical multiplexers. This observation indicates that today's host-to-host based content distribution is inherently less energy efficient and brings home to us the value of exploiting the capability of energy-efficient networking devices using a radically different approach, namely content-centric/oriented networking where queries and data are routed based on content name, which can be either opaque or structured (as in URL). In recent years, a number of different approaches have been proposed in the research community such as TRIAD, Data-Oriented Network Architecture (DONA), and Content Centric Networking (CCN). Among these proposals, this article focuses on CCN proposed by Jacobson *et al.* [8] because it provides an ideal platform for network-wide content caching.

CCN uses user-friendly hierarchical names like URLs. As in BitTorrent-like file swarming, the original content is divided into multiple chunks. In CCN, the content publisher signs

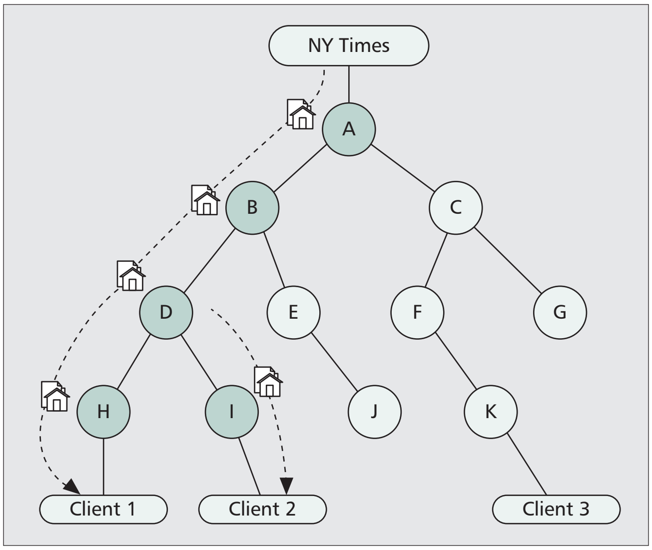

Figure 3. *Illustration of content caching in CCN: Client 1 fetches the front page of* NY Times *by sending an interest packet. Initially, there were no other interests issued earlier, and intermediate CCN routers do not have the requested chunk in their buffer. The interest packet is eventually delivered to the origin server of* NY Times *along the shortest path (i.e., H, D, B, A). The requested chunk is then delivered by traversing the reverse path (i.e., A, B, D, H), and each intermediate router keeps the forwarded chunk in its router buffer. If client 2 wants to access the same page, its interest packet will find a match at node D, and the requested chunk will be directly delivered from that node.*

each chunk along with its name to guarantee secure linkage between the name and the content chunk; intermediate routers can validate an incoming chunk using this per-chunk signature. A content publisher announces content availability via Border Gateway Protocol (BGP) prefix announcement. The basic operation of a CCN router is very similar to an IP router. An interest (or request) packet arrives on an interface, and then a longest-match lookup is done on *its name* to make a forwarding decision. Since CCN routers support caching within the network, any intermediate CCN routers on paths that have the requested chunk can answer the request. The request chunk is then delivered by following a reverse path created while forwarding the interest packet. In Fig. 3 readers can find an illustration of content caching in CCN.

We now discuss how to build energy-efficient CCN routers. The critical components are storage devices: DRAM, solid state drives (SSDs), and disks. Modern routers have a dedicated forwarding engine for each line card, a routing processor for control packets (e.g., routing, network management), and a set of switching fabrics that interconnect pairs of arbitrary interfaces. Forwarding engines can communicate with the route processor via a high-speed backplane. DRAM and disks are power-hungry (e.g., 2.5 W/Gbyte for DRAM, and 12 W for a Seagate ST3160021A 160-Gbyte disk) compared to SSD (e.g., 1 W for a 64-Gbyte Samsung SSD) [3]. Thus, it is desirable to implement a hierarchical storage structure spanning both forwarding engines and the route processor: i.e., forwarding engine can be equipped with additional DRAM and SSD (say, 4 Gbytes of DRAM and 64 Gbytes of SSD), and the route processor can have a large storage device. If a router has 8 line cards, for example, this configuration incurs additional power usage of 100 W. This extra overhead is reasonable in most network devices except the home gateways. In NaDa, we should only install a small amount of DRAM and a single SSD (say, 2 Gbytes DRAM and 64 Gbytes SSD) due to the energy efficiency concern. In Fig. 2 we incorporate the above mentioned configurations into the energy consumption calculation of networking devices. The results show that CCN capability slightly increases energy consumption when comparing watts per gigabit per second.

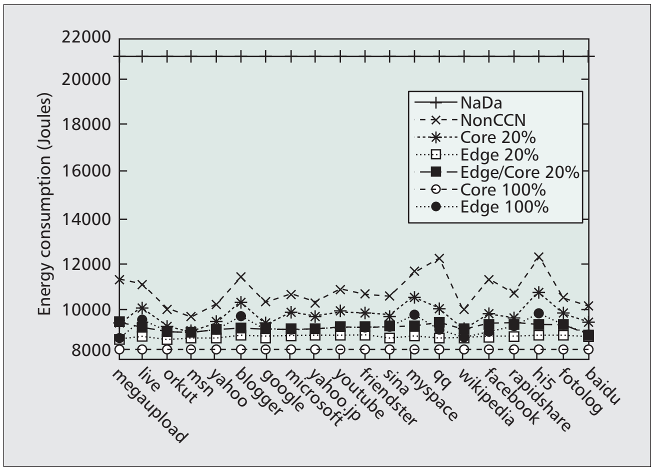

Figure 4. *Energy consumption comparison (DSL).*

The benefit of CCN from an energy perspective relies on its energy proportional computing and networking capability and deployment status. CCN can generally benefit from recent proposals on energy proportional networking such as rate adaptation and sleeping [2, 3]. However, CCN's additional logic such as content lookup/caching and cache replacement needs to be efficiently implemented to keep the energy consumption for content processing as low as possible; otherwise, the transport energy must be traded against the energy-hungry content processing. Another critical factor is that a sufficient number of content routers must be deployed throughout the network to reap the benefit of content caching.

Energy Efficiency Evaluation

Energy efficiency of the different content distribution strategies can be evaluated using the traceroute data set in [1], which captures a single snapshot of the network of the top 20 content providers ranked by Alexa in September 2007. It was collected by querying these providers from 18 different traceroute servers located in the United States. We identify which of the hops in the traceroute path belong to tier 1 ISPs using a publicly available tier 1 ISP list. We assume that tier 1 ISPs use core routers, and the other nodes in a path are either smaller core or edge routers. According to Gill *et al.* [1], large content providers such as Google and Microsoft have content servers located in a few large data centers that are interconnected using private high-speed networks (i.e., they tend to bypass tier 1 ISPs and have much smaller tier 1 hop counts).

The benefit of CCN under incremental deployment scenarios is evaluated by varying the fraction of CCN-enabled routers from 20 percent to 100 percent. For a given traceroute path, we randomly select a given fraction of core/edge routers as CCN routers. We assume that users download highly popular content; thus, any of the intermediate CCN routers (including NaDa servers in home gateways) can always serve the content. We measure the total energy consumption for downloading 1 Gb of video content (125 Mbytes).

The simulation is based on the energy consumption model with high idle power consumption and low utilization (Fig. 2). We consider different scenarios as follows. In a traditional content distribution network (denoted NonCCN), since content is fetched from the origin server, the energy model considers content server, core/edge routers, DSLAM/OLT, and gateways. Unlike NonCCN, the energy model for CCN may

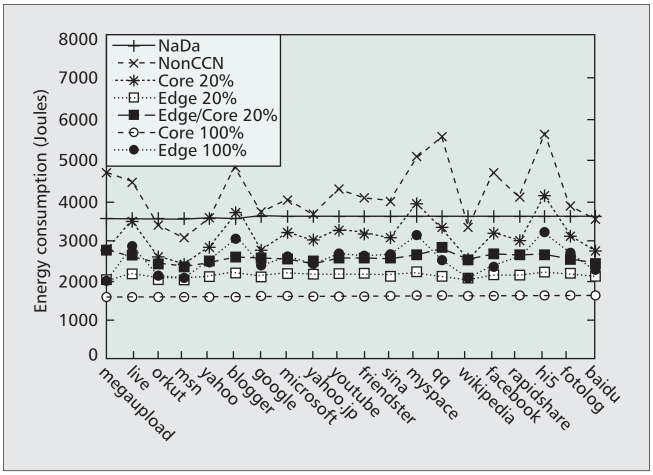

Figure 5. *Energy consumption comparison (GPON).*

not include the server due to content caching in CCN routers. In the incremental deployment scenarios, requests may not go through CCN routers, and content is fetched from the origin server. CCN capability is only enabled in the edge and core routers; the DSLAM/OLT and home gateway have the same energy efficiency as in the case of NonCCN. In NaDa content is always delivered from one gateway to the other via a local DSLAM/OLT in two hops (from source gateway to DSLAM/OLT to destination gateway). Since a gateway is augmented with content server capability as illustrated in the previous section, its power consumption is higher than that of NonCCN.

Energy consumption comparison: We consider the energy consumption of different deployment scenarios of CCN nodes (20 and 100 percent) and compare its performance with NaDa and the case without using CCN. We plot the results for DSL and GPON users in Figs. 4 and 5, respectively. The figures show that even with 20 percent deployment of CCN routers in the core, CCN can effectively reduce the hop length, thereby significantly reducing energy consumption. Also, it is more effective to deploy CCN nodes at the edge because requests travel across a shorter distance. Note that the total number of CCN routers will be much greater when the same fraction of content routers are deployed in the edge, as the number of edge routers tends to grow exponentially. Figure 4 shows that NaDa's energy efficiency for DSL users is much worse than that of non-CCN scenarios mainly because NaDa uses energy-inefficient DSL gateways twice (i.e., 2 × 10,080 J). When considering GPON-based gateways, NaDa's performance is better than non-CCN scenarios, and yet is still worse than CCN scenarios in most cases (Fig. 5).

In Fig. 6, we plot the percentage of total energy consumption of various components in the content delivery chain: server, core/edge router, DSLAM/OLT, and home gateway. It shows that home gateways contribute to a large fraction of the total energy consumption, and this fraction grows as the number of CCN routers increases. For instance, DSL users with NonCCN consume 66 percent of the total energy, whereas those users with 100 percent deployment of CCN routers in the edge consume 86 percent of the total energy in their gateways. The overall trend is quite persistent in both DSL and GPON. As the fraction of CCN-enabled routers increases, the server cost decreases, because requests can be served in the intermediate routers. In the CCN scenario with 100 percent deployment in the edge, edge routers can fully serve the content; thus, there is no server cost in the energy consumption breakdown. In the CCN scenario with 100 percent deployment in the core, however, core routers may not fully serve the content, since requests could bypass core routers. Recall that in the traceroute data set, most paths are composed of a sequence of edge-core-edge hops, and few paths do not have core hops [1].

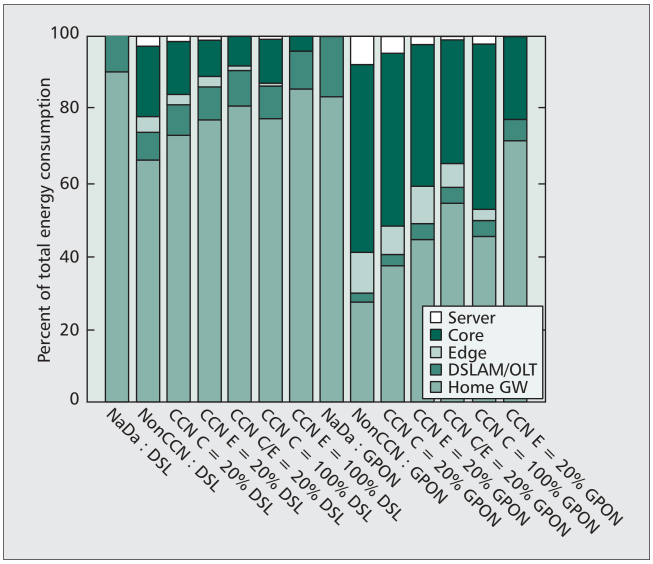

Figure 6. *Energy consumption breakdown.*

Conclusion

We have reviewed existing content dissemination methods and summarized our evaluation results on energy efficiency of networking devices ranging from edge/core routers to home gateways. After illustrating content-centric networking, we have considered energy-efficient content router configurations of different networking devices. The energy efficiency comparison via trace-based simulations shows that content-centric networking can substantially improve the energy efficiency of content dissemination. Our preliminary results are encouraging and will stimulate further research in this direction. CCN's content processing needs to be efficiently implemented to keep the energy consumption as low as possible. Given that energy saving mainly comes from reducing hop counts by serving content at intermediate nodes, the impact of cache placement (under incremental deployment scenarios) and local/cooperative content replacement strategies needs to be carefully investigated to better understand the energy efficiency of content-centric networking. Since optical transmission hardware is an order of magnitude more efficient than a core router [6], transparent optical connections at the physical layer (e.g., optical bypass) have similarly been considered as a method to reduce hop counts and energy consumption in content distribution [10]. Although providing dynamic content-level optical bypass is not practical today due to physical layer constraints on wavelength switching, it would be interesting to understand its interaction with CCN for energy saving.

References

[1] P. Gill *et al.*, "The Flattening Internet Topology: Natural Evolution, Unsightly Barnacles or Contrived Collapse?" *PAM '08*, Cleveland, OH, Apr. 2008.

[2] L. A. Barroso and U. Holzle, "The Case for Energy-Proportional Computing," *IEEE Computer*, vol. 40, 2007, pp. 33–37.

[3] D. Meisner, B. T. Gold, and T. F. Wenisch, "PowerNap: Eliminating Server Idle Power," *ASPLOS '09*, Washington, DC, Mar. 2009.

[4] J. Chabarek *et al.*, "Power Awareness in Network Design and Routing," *INFOCOM '08*, Phoenix, AZ, Apr. 2008.

[5] S. Nedevsch *et al.*, "Reducing Network Energy Consumption via Sleeping and Rate-Adaptation," *NSDI '08*, San Francisco, CA, Apr. 2008.

[6] O. Tamm, C. Hermsmeyer, and A. M. Rush, "Eco-Sustainable System and Network Architectures for Future Transport Networks," *Bell Labs Tech. J.*, vol. 14, 2010, pp. 311–27.

[7] V. Valancius *et al.*, "Greening the Internet with Nano Data Centers," *CoNEXT '09*, Rome, Italy, Dec. 2009.

[8] V. Jacobson *et al.*, "Networking Named Content," *CoNEXT '09*, Rome, Italy, Dec. 2009.
[9] X. Xiao *et al.*, "A Practical Approach for Providing QoS in the Internet Backbone," *IEEE Commun. Mag.*, vol. 40, 2002, pp. 56–62.
[10] J. Baliga *et al.*, "Energy Consumption in Optical IP Networks," *J. Lightwave Tech.*, vol. 27, 2009, pp. 2391–2403.

Biographies

UICHIN LEE (uclee@kaist.ac.kr) is an assistant professor in the Department of Knowledge Service Engineering at Korea Advanced Institute of Science and Technology (KAIST). He received a B.S. in computer engineering from Chonbuk National University in 2001, an M.S. degree in computer science from KAIST in 2003, and a Ph.D. degree in computer science from the University of California at Los Angeles (UCLA) in 2008. Before joining KAIST, he was a member of technical staff at Bell Laboratories, Alcatel-Lucent until 2010. His research interests include distributed systems and mobile/pervasive computing.

IVICA RIMAC (ivica.rimac@alcatel-lucent.com) is a member of technical staff in the Service Infrastructure Research Domain at Bell Laboratories, Alcatel-Lucent. He received his Master's and Ph.D. degrees in electrical engineering and information technology from Darmstadt University of Technology, Germany. He has conducted research in the areas of content distribution networks, content delivery technologies for mobile networks, peer-to-peer technologies, and content-centric networks.

VOLKER HILT (volker.hilt@alcatel-lucent.com) works for Bell Laboratories, Alcatel-Lucent, where he is headi of a department in the Service Infrastructure Research Domain. His research interests include content distribution networks, cloud computing, P2P technologies, and distributed multimedia applications. He is a leading contributor to several standards in the Internet Engineering Task Force (IETF), and chairs the IETF SOC WG and Internet Research Task Force P2P RG. He has published numerous papers in international conferences and journals and has served as co-chair and program committee member for many conferences. He has received his Master's degree in computer science and business administration in 1996 and his Ph.D. in computer science in 2001, both from the University of Mannheim, Germany.

DANIEL KILPER [SM] (dan.kilper@alcatel-lucent.com) is a member of technical staff at Bell Laboratories, Alcatel-Lucent. He received B.S. degrees in electrical engineering and physics from Virginia Tech, and M.S. and Ph.D. degrees in physics from the University of Michigan, Ann Arbor. He is an associate editor for *OSA/IEEE Journal of Optical Communications and Networking*. He currently serves as chair of the GreenTouch Consortium technical committee. While at Bell Laboratories he has conducted research on optical performance monitoring, network energy trends, and transmission and control systems for transparent optical networks.